J Econ Interac Coord (2008) 3:25–41
DOI 10.1007/s11403-008-0037-3

REGULAR ARTICLE

On the plausibility of sunspot equilibria

Simulations based on agent-based artificial stock markets

Shu-Heng Chen · Chung-Chih Liao · Pei-Jung Chou

Published online: 3 April 2008

Abstract This study examines the plausibility of the emergence of sunspot equilibria in an agent-based artificial stock market. Using the agent-based model, we make the sunspots explicit so that we can test, e.g., by means of the Granger causality test, whether purely extrinsic uncertainty can influence price dynamics. In addition, through agent-based simulation, the coordination process, which is mainly driven by genetic programming, becomes observable, which enables us to analyze what agents perceive and whether they believe in sunspots. By manipulating different control variables, three series of experiments are conducted. Generally speaking, the chances of observing "sunspot equilibria" in this agent-based artificial stock market are small. However, the sunspot believers can never be driven out of the market. Nevertheless, they are always outnumbered by fundamental believers, which is evidence that the market as collective behavior is rational. We also find that lengthening the time horizon will make it difficult for sunspot believers to survive.

Keywords Sunspots · Agent-based artificial stock markets · Genetic programming · Granger causality test

S.-H. Chen (✉) · P.-J. Chou
Department of Economics, AI-ECON Research Center,
National Chengchi University, Taipei 116, Taiwan
e-mail: chchen@nccu.edu.tw

P.-J. Chou
e-mail: peijungchou@gmail.com

C.-C. Liao
Department of International Business, National Taiwan University,
Taipei 106, Taiwan
e-mail: chungliao@gmail.com

1 Motivation and introduction

While the theoretical plausibility of sunspot equilibria has been extensively addressed in different economic models, these models have been almost entirely studied within the context of the homogeneous rational expectations equilibrium, and have been derived with the device of representative agents.[1] This homogeneous-agent framework has also shaped the learning approach to sunspot equilibria that has later arisen. Although different kinds of adaptive learning have been proposed, they mainly deal with the learning of a representative agent Woodford (1990); Evans and McGough (2005). Dawid (1996) is probably the one of the few studies to address the learning of sunspot equilibria within the context of *heterogeneous agents*. As opposed to models of adaptive learning with representative agents, models of adaptive learning with heterogeneous agents enable us to explicitly tackle the *coordination issue*, specifically, the *coordination mechanism of expectations*. This is certainly desirable since sunspots are often used as a coordination device in relation to expectations.

In this paper, we continue the line of research of sunspot equilibria with the device of heterogeneous agents. As in Dawid (1996), we adopt an *agent-based computational approach* to study the plausibility of sunspot equilibria. The agent-based computational approach, now also known as *agent-based computational economics* (ACE), has been adopted to study many economic models which are related to the analytical models of sunspot equilibria. One of the most important classes of the ACE models is the *agent-based artificial stock markets* (AASMs). While AASMs have been extensively studied over the last decade, and different aspects have been enriched,[2] no one has ever seriously cast doubt on the plausibility of sunspot equilibria in this type of model. We, however, find that it would be extremely important to begin the inquiry.

First, the original inquiry into the plausibility of sunspot equilibria was very much inspired by stock-market phenomena. For example, in their celebrated work, “Do Sunspots Matter,” Cass and Shell (1983) raised this kind of question right at the very beginning of their thesis. Second, besides appearing in analytical papers, the same inquiry has also been clothed with empirical investigations. For example, in the first experimental study on sunspot equilibria, Marimon et al. (1993) questioned the relevance of sunspots to excess volatility in the stock price. Thirdly, when reflecting upon the policy implications of sunspot equilibria, we are once again concerned with the stock market Duffy and Fisher (2005). Finally, there are a number of studies that work with sunspot equilibria directly in the domain of asset and derivative markets, e.g., Kajii (1997). For this literature, the stock market plays an important role.

Given such a great interest in exploring the potential implications of sunspot equilibria for stock markets, it would be a little surprising that none of the theoretical, empirical, experimental, or even simulation models of sunspot equilibria directly

[1] We do acknowledge the recent increasing efforts devoted to the study of heterogeneous expectations, for example, Evans and Honkapohja (1996); Negroni (2005). However, few efforts have been applied to the study of sunspot equilibria.

[2] For a survey, see LeBaron (2006).

capture sunspots within a stock market composed of heterogeneous agents. This paper aims to break this silence. We bring the original Cass-Shell issue Cass and Shell (1983) back into the domain of the agent-based artificial stock market, and directly answer the question: *do investors concentrate on "fundamentals," or do they instead focus on the extrinsic variables?*

1.1 Agent-based artificial stock markets

The features and advantages of the ACE approach have become clear, given the intensive efforts made over past years.[3] In this specific case, the agent-based artificial stock market, like the theoretical or experimental environments, enables us to directly control *what exactly sunspot variables are*, which are more difficult to identify in empirical studies. In addition, as in experimental models, the agent-based approach enables us to observe explicit "*coordinating processes*" of expectations, from which sunspot equilibria may or may not arise. These two advantages enable us to collect *direct "evidence"* pertaining to sunspot equilibria.[4]

In this paper, we study the plausibility of sunspot equilibria using a variant of the SFI-type (Santa Fe Institute) artificial stock market, known as AIE-ASM, which is initiated by Chen and Yeh (2001, 2002). As in many other similar models, the stock price in this model is determined by agents' portfolio decisions, which in turn depend on their expectations of the stock price. Chen and Yeh (2001, 2002) modeled agents' expectations via genetic programming (GP), which basically allows agents, at any point in time, to "choose" any forecasting function of any variables "spanned" by the prespecified primitives, which are composed of a function set and a terminal set.

The terminal set is a set of constants and variables. In the past, only the "fundamentals," such as prices, dividends, and trading volumes, were included in the terminal set. *The key departure of this paper is to add sunspots to the terminal set.* By doing so, agents in this market will observe the sequence of past fundamentals and sunspots, and will form their expectations by selecting what they believe to be relevant to predict the future prices. Sunspots, by definition, are extrinsic uncertainties, and are supposed to be irrelevant, but agents do not know which variables are sunspots at the beginning. The learning process driven by GP may or may not help agents to learn which variables are sunspots. In the former case, agents do not learn to believe in sunspots, while in the latter case, agents learn to believe in sunspots.

1.2 Sunspot equilibria

When agents successfully distinguish sunspots from fundamentals, sunspots are merely external random variables, and will have no influence on prices. In this case, there is, of course, no sunspot equilibrium. When agents fail to distinguish sunspots from

[3] Tesfatsion and Judd (2006) provided a comprehensive review of the development of the ACE and its extensive applications to various economic domains.

[4] Based on Duffy and Fisher (2005), the evidence is *indirect* if the sunspot variables and the coordination processes are not identified.

fundamentals, agents' expectations will then be partially driven by sunspots. Nonetheless, due to the heterogeneity in beliefs, it is not immediately clear whether sunspots will influence price dynamics. This will depend on how many agents believe in sunspots, how "sincere" they are or how they interpret sunspots. As a result, the sunspot equilibrium may happen or may not happen.

Having said that, we have to remark on the sunspot equilibria used in this paper as opposed to the familiar ones. Sunspot equilibria are generally defined as *rational expectations equilibria in which purely extrinsic uncertainty affects equilibrium prices and allocations* Woodford (1990). Our notion of sunspot equilibria only retains the second part of this definition, i.e., non-fundamental stochastic disturbances influence model dynamics, particularly, price dynamics, but not the first part, *rational expectations equilibria*. Defining rational expectations equilibria is not simple when the homogeneous rational expectations are not necessarily attainable, and may not even be useful when the system keeps on evolving and changing. Therefore, we have to relax the original notion, while still keeping its essential ingredient. As we shall see later, this extended notion allows us to use the *Granger causality test* as a technical counterpart of our notion of sunspot "equilibria."

1.3 Forecasting with sunspots

The discussion above leads to another important feature of the paper, i.e., the explicit addressing of the forecasting with sunspots. While some studies on sunspot equilibria also share this feature, we work in a quite different manner. In our setting, the nature of extrinsic uncertainty is closely associated with boundedly-rational agents, who have only limited knowledge of the system, and cannot tell the intrinsic variables from the extrinsic variables. These boundedly-rational agents, however, are able to learn, but it is not guaranteed that they can eventually distinguish between the two. What is even more intriguing is that when during the course of learning the agents may mistakenly regard some extrinsic variables as intrinsic variables, even before they discover them, the extrinsic variables, due to a kind of *self-fulfilling prophecy*, have already been successfully "planted" into part of the system, and hence are no longer extrinsic. In this manner, we work with the entire coordination process of believing in sunspots, and the legitimate issue of spontaneous coordination of all agents on a given equilibrium, frequently raised in the conventional sunspot literature, does not occur here.

This trial-and-error learning process also sharply contrasts with the rational expectations approach, which usually assumes that agents *are aware of* the irrelevance of the sunspot process in the determination of the fundamentals. This is not surprising since rational-expectations agents, by definition, expect sunspots to affect prices and make their decisions according to such an expectation. Our agents, however, do not have this great awareness. The learning process for them is *not* a process to select the sunspots upon which they should coordinate their expectations. Right from the beginning, they simply cannot even tell which variables are sunspots, and which are not, and they are not even sure of the existence of sunspots. In a sense, they learn from scratch. The learning process for them is merely *a process to distinguish extrinsic variables from intrinsic variables*, i.e., a tendency to get rid of the irrelevant.

1.4 Coordination

Experimental approaches to sunspot equilibria always involve an explicit coordination process, such as a common historical experience, a common understanding of the sunspot language, or a common interpretation of sunspots Marimon et al. (1993); Duffy and Fisher (2005). When sunspots are observable, such as in the case of Duffy and Fisher (2005), the coordination process even involves some degree of rational learning Bray and Kreps (1987), in which agents will speculate on what other agents do with these sunspots.[5] If we consider this to be a formal coordination process, then our ACE approach is not explicit to this degree. Our agents, basically, will not explicitly care what other agents do with the sunspots. From beginning to end, they only care about how the historical data speaks to them.

At first sight, our agents might appear a little "dumber," but they are not. This has to do with how we introduce sunspot variables, and how many agents receive this "information." If sunspots are Markovian, and have only two states, say "high" and "low," and if this information is released to a very small society of agents, as often happens in the laboratory with human subjects, then it seems more natural for each agent to have a model of the minds of other agents. However, when sunspots are continuous and have an infinite number of states, and they are released to a rather large society of agents, then it is also natural for agents to be more "humble," given this increase in computational complexity.

In this regard, this paper, therefore, shares the views of Hens (2000) in regard to the role of sunspots. "Clearly sunspots can be used as a coordination device of expectations, but from a more general perspective sunspot equilibria are any equilibria in which the equilibrium allocation depends on some exogenous random event." (Ibid, p. 435).

The remainder of the paper is organized as follows. Section 2 formally introduces our definition of the sunspot equilibria of the agent-based artificial stock markets via the Granger causality test. Since the focus of the paper is neither the agent-based artificial stock market, nor the causality test, which are available in other studies, we shall dispense with the usual detailed description of both of them. Nonetheless, to make this paper as self-contained as possible, a brief introduction to the agent-based artificial stock markets and the causality test used in this paper is provided in Appendix A and Appendix B. We then proceed directly to the details of the experimental designs (Sect. 3). Three series of experiments will be conducted in this paper, and the results are presented in Sects. 4–6.

2 Sunspot equilibria and Granger causality

Based on the discussion in Sect. 1.2, we shall formally propose the definition of sunspot equilibria used in this paper. Consider an economic system, Ω, whose dynamics is represented by a stationary time series $\{X_t\}$. Let $\{Z_t\}$ be the other stationary time

[5] For example, by taking an example from Duffy and Fisher (2005), agents may speculate on how other agents interpret "sunshine" or "rain."

series, which is exogenously generated, and is extrinsic to Ω. Hence, $\{Z_t\}$ cannot determine and cannot help forecast $\{X_t\}$ unless there are agents *who believe that it can*, and the expectations self-fulfill. A sunspot equilibrium in Ω is then defined as a pair of $\{X_t, Z_t\}$, where $\{Z_t\}$ can *Granger-cause* $\{X_t\}$. With this definition, whether the sunspot equilibrium exists or not can be tested by the familiar Granger causality test. The null hypothesis is that the sunspot equilibrium does not exist, i.e., $\{Z_t\}$ fails to Granger-cause $\{X_t\}$. In this paper, Ω corresponds to the artificial stock market, and $\{X_t\}$ refers to the stock return series which are endogenously generated from the agent-based artificial stock market, of which participants can observe a sunspot series $\{Z_t\}$, which is to be detailed in Sect. 3.1. Therefore, the artificial stock market is said to observe a sunspot equilibrium if $\{Z_t\}$ Granger-causes the return series $\{X_t\}$. In this paper, in addition to the usual linear causality test Granger (1969), we also consider the non-linear causality test Hiemstra and Jones (1994).

3 Experimental designs

The paper is comprised of three series of experiments, which are motivated sequentially. In this section, we shall briefly summarize what we find from each series of experiments, and how that determines the design of the next series. All detailed results are given in Sects. 4–6. Before proceeding further, we notice that there is a fundamental design which is the same for all three series of experiments. Since they are the same, we will leave all the information about them to Appendix A.

3.1 Series I

The following series of experiments hinges upon how sunspots are designed and introduced into the system. We shall then first describe the design of sunspots. Sunspots are extrinsic (non-fundamental) random variables. Throughout the paper, we assume that the sunspot is *i.i.d.* uniform. Denote the sunspot series by $\{Z_t\}$, $Z_t \sim Uniform$ $[U_l, U_h]$. To make Z_t comparable to the scale of the stock price generated by the fundamentals, we set $U_l = 50$, and $U_h = 150$ so that $E(Z_t) = 100$.

As in the standard sunspot literature, agents in the markets can observe a sequence of past sunspot variables. While an ideal situation is to make them able to observe the infinite sequence of the past $\{Z_{t-j}\}_{j=1}^{\infty}$, we, however, have to put an upper limit on the number of lags so as to facilitate our implementation of GP. For example, a typical terminal set in Chen and Yeh (2001, 2002) is

$$\mathcal{T} = \{P_{t-1}, P_{t-2}, \ldots, P_{t-10}\}, \tag{1}$$

where P_{t-j} is the stock price with a lag of j. Since the price dynamics is endogenously generated, and the past price can in general impact the future price, there is little doubt that this series is fundamental and not extrinsic. If we start with a terminal set like (1), how should we then add the sunspot variables into the terminal?

Table 1 Experimental designs

		Case	τ	H	Terminal set
I	I-A	A01–A50	1	10	$\{P_{t-1}, Z_{t-1}\}$
	I-B	B01–B50	1	10	$\{P_{t-1}P_{t-2}, Z_{t-1}, Z_{t-2}\}$
	I-C	C01–C50	1	10	$\{P_{t-1}, \ldots, P_{t-3}, Z_{t-1}, \ldots, Z_{t-3}\}$
	I-D	D01–D50	1	10	$\{P_{t-1}, \ldots, P_{t-4}, Z_{t-1}, \ldots, Z_{t-4}\}$
	I-E	E01–E50	1	10	$\{P_{t-1}, \ldots, P_{t-15}, Z_{t-1}, \ldots, Z_{t-15}\}$
II	II-A	A01–A50	0.5	10	$\{P_{t-1}, P_{t-2}, Z_{t-1}, \ldots, Z_{t-4}\}$
	II-B	B01–B50	2	10	$\{P_{t-1}, \ldots, P_{t-4}, Z_{t-1}, Z_{t-2}\}$
III	III-A	A01–A50	1	30	$\{P_{t-1}, \ldots, P_{t-4}, Z_{t-1}, \ldots, Z_{t-4}\}$
	III-B	B01–B50	1	50	$\{P_{t-1}, \ldots, P_{t-4}, Z_{t-1}, \ldots, Z_{t-4}\}$

τ refers to sunspot density, whereas H refers to time horizon

A natural start is to think of an equally-long series of sunspots, namely,

$$\mathcal{T} = \{P_{t-1}, P_{t-2}, \ldots, P_{t-k}, Z_{t-1}, Z_{t-2}, \ldots, Z_{t-k}\}. \tag{2}$$

By doing so, we can avoid the initialization bias either because of too many past observations of the fundamental (price) or too many past observations of the sunspot. This is exactly how we design the first series of experiments.

The first series of experiments comprises five experiments, coded by I-A, -B, -C, -D, and -E, respectively. Over these five experiments, we make the size of the fundamental set, denoted by $|\,\mathbf{F}\,|$, and the size of sunspot set, denoted by $|\,\mathbf{S}\,|$, identical ($|\,\mathbf{S}\,| = |\,\mathbf{F}\,| = k$) so that the sunspot density, defined as

$$\tau = \frac{|\,\mathbf{S}\,|}{|\,\mathbf{F}\,|}, \tag{3}$$

is fixed to one ($\tau = 1$). From experiment I-A to I-E, we, however, vary the size of the sunspot set, k, by increasing it from 1, 2, 3, 4, and then to 15 (see the last column of Table 1). The purpose of this experiment is to examine whether the absolute size of the sunspot set can have an effect on the probabilities of the sunspot equilibria. For each experiment, we conduct 50 runs, and each run lasts for 3,000 periods (trading days). The results of this series of experiments are detailed in Sect. 4.

3.2 Series II

From the five experiments for Series I, we find that the absolute size of the sunspot set plays little role in the determination of the probability of sunspot equilibria; hence, we move to the second series of experiments, which is to vary the sunspot density.

We consider three possible sunspot densities: sunspots are weak ($\tau = 0.5$), sunspots are medium ($\tau = 1$), and sunspots are strong ($\tau = 2$). In the medium case, we set the size of both the fundamental and sunspot sets to be 4, $|\mathbf{S}| = |\mathbf{F}| = 4$, and adjust them to fit different densities accordingly. So, for the case of $\tau = 0.5$, $|\mathbf{S}| = 2$, and $|\mathbf{F}| = 4$; for the case of $\tau = 2$, $|\mathbf{S}| = 4$, and $|\mathbf{F}| = 2$. With this setting, we only have to run two more experiments, since the case of medium density is just the same as Market I-D. We shall then code these two additional experiments by II-A ($\tau = 0.5$) and II-B ($\tau = 2$). This has also been summarized in Table 1 (second block). The results of this series of experiments are not presented in this paper since we do not find the significance of sunspot density. The probability of sunspot equilibria remains low even though we increase the sunspot density. Therefore, we conclude that it is difficult to have sunspot equilibria in this agent-based artificial stock market.

We then look at the individual behavior of both Series I and II. Specifically, we are addressing the following question. Would the low probability of sunspot equilibria, the lack of the causal relation, be caused by the agents' ability to distinguish sunspots from fundamentals? From Series I and II, we observe that, while the market participants, as a whole, can tell the difference between the fundamentals and sunspots, a large number of *sunspot believers*, about 60% to almost 80% of market participants, still remain. At first sight, this phenomenon is puzzling since the absence of Granger causality in the aggregate dynamics might easily lead agents to see the irrelevance of sunspots, and disregard sunspots accordingly. So, the large number of surviving sunspot believers requires an explanation. This gives rise to the third series of experiments.

3.3 Series III

One possible explanation, motivated by LeBaron (2001), is that the survival pressure we put on the agents tends to drive agents to search for any possible short-term signals rather than a long-term relationship. As a result, we conducted another series of experiments in order for us to figure out the survival of a large number of sunspot believers. In the vein of LeBaron (2001), we choose the *time horizon* as another control variable, and simulate three time horizons, namely, $H = 10, 30,$ and 50. As in Series II, we choose Market I-D as the starting point, which has a H of 10. Two additional experiments, Market III-A and III-B, have the same design except that H is extended to 30 and 50, respectively. This setting is summarized in Table 1, the third block.

In AIE-ASM, during the course of evolution or learning, each forecasting rule has to be evaluated. Genetic programming is then operated based on the fitness of each forecasting rule. The time horizon serves as a time frame by which a forecasting rule is evaluated. A short time horizon implies that the evaluation will be made based on very recent periods, whereas a long time horizon implies that it will be done based on a longer recent past. For example, when $H = 10$, the performance is based on the in-sample forecasting errors of the 10 most recent periods, while when $H = 50$, it is based on the 50 most recent periods. It is conjectured that the longer the time frame, the easier it is for the agents to see the irrelevance of sunspots. The purpose of the third series of experiments is then to test whether this is the case, and to see how the time horizon impacts the microstructure and the macro-behavior. The results are detailed in Sect. 6.

Table 2 Granger causality test: series I

Experiment	Time horizon	Sunspot density (τ)	$z_t \nrightarrow r_t$ p-value	$z_t \nRightarrow r_t$ p-value
I-A	10	1	0.4207	0.5714
I-B	10	1	0.4696	0.5817
I-C	10	1	0.4331	0.5122
I-D	10	1	0.4343	0.5339
I-E	10	1	0.4799	0.6029

$z_t \nrightarrow r_t$ corresponds to the linear causality test, whereas $z_t \nRightarrow r_t$ corresponds to the non-linear causality test. Note that the p-values presented here are the simple averages taken over 50 runs for each experiment, and the null is that sunspots fail to Granger-cause stock returns

4 Experimental results: fixed sunspot density

4.1 Results from the top

Our results can be presented in two separate parts. First, at the aggregation level, we want to know the probability of sunspot equilibria, i.e., sunspot variables are *endogenous* in the sense of Granger causality. From our simulation results, we find that there is a chance of observing the sunspot equilibrium, but only with a small likelihood. Out of 250 runs, only 25 runs reject the null of no sunspot equilibrium at a significance level of 0.05, i.e., an estimate of 0.1 for the probability of sunspot equilibria. The 25 rejections are not distributed monotonically: 6 belong to Market I-A, 4 to Market I-B, 6 to Market I-C, 7 to Market I-D, and 2 to Market I-E. There is no indication that the absolute size of the sunspot ($|\mathbf{S}|$) variables will have an impact on the frequency of observing a sunspot equilibrium, given a fixed sunspot density (τ).

In Table 2, we present the results of the non-linear Granger causality test. Depending on the number of lags, there are different results for the non-linear Granger causality test.

What we do in Table 2 is, therefore, to take a simple average of all p-values up to lag 10. By so doing there is no single run that is able to reject the null. So, by combining the linear and non-linear part together, there is only a total of 25 runs which exhibit a causal relationship from sunspots to returns.

4.2 Results from the bottom

While the sunspot variables do not seem to have a real effect on the return dynamics, it would still be interesting to see whether individual agents have the capability to distinguish the fundamental from the sunspot variables. One way to do this is to count *the number of sunspot believers*. We consider two definitions of believers, namely, the *sunspot believers* and the *pure sunspot believers*. The former is characterized as the agent whose forecast is a function of sunspot variables, whereas the latter is characterized as the agent whose forecast is a function of sunspot variables *only*.[6] Call the number N_t^Z and N_t^{PZ}, respectively. In a similar way, we can define and count the number of *fundamental believers* and the number of *pure* fundamental believers, and denote them by N_t^X and N_t^{PX}, respectively. Figure 1 shows the time series plot of N_t^X and N_t^Z (the left half), and the time series plot of N_t^{PX} and N_t^{PZ} (the right half) of the five experiments. Each time series plotted here is not based on any single run, but on the average over 50 runs.

From Fig. 1, we can see that there are quite a number of sunspot believers, and even pure sunspot believers, that remain to the very end of the simulation. These believers of extrinsic uncertainty do not go away with time, even though their boundedly-rational behavior is designed to be able to learn, via genetic programming. However, if the number of sunspot believers is placed together with the number of fundamental believers, we can still see the difference: both N_t^X and N_t^{PX} are significantly greater than N_t^Z and N_t^{PZ}, respectively. Almost starting from the very beginning, we have already seen the dominance of the population of the fundamental believers over the population of the sunspot believers (also see Table 3 for the average of each series).[7] Therefore, while sunspot believers are not eliminated via learning, their population is dominated by that of the fundamental believers during almost the entire course of evolution. As a result, our agents as a collection, through the learning mechanism driven by genetic programming, are able to distinguish fundamentals from sunspots; at least, they learn that fundamentals are more pertinent than sunspot variables.

5 Would sunspot density matter?

It may be anticipated that sunspot density can have a positive influence on the probability of sunspot equilibria. However, the Granger causality test results seem to go against this intuition. Out of all 150 runs (50 runs for each market), there are a total of 16 runs supporting sunspot equilibria, which is again about an estimate of 10% for the probability of sunspot equilibria. Nevertheless, from the distribution of the frequencies of rejection, there is no evidence to indicate that the probability of sunspot equilibria will increase with sunspot density.

The general result from these two series of experiments is that Granger causality tests generally do not support the causal relationship from sunspots to returns, which indicates that sunspot variables largely remain *exogenous* to the system. This result is

[6] This idea is very similar to Chen and Yeh (2001) and Chen and Liao (2005), where they define the *martingale believer* and the *volume believer*. However, this way of defining the sunspot believer has its shortcomings since simple GP is notorious for creating many redundant parts of the forecasting function, e.g., multiplying a sunspot variable by zero, subtracting the included sunspot variable by itself, etc. We, nonetheless, consider this problem to not be too harmful because the problem also appears in the counting of the *fundamental believers*, to be discussed below. In this case, by merely comparing their relative differences, one can still draw some sensible conclusions.

[7] The last column of Table 3, N^C, refers to the number of agents who use neither sunspots nor fundamentals in their forecasting functions; they only use *constants* to form their forecasts. These agents believe that price plus dividends has a constant mean, although their expectations of this constant mean may be constantly revised. Since sunspots are not involved in their expectation behavior, believers of "mean regression" can broadly be regarded as fundamental believers.

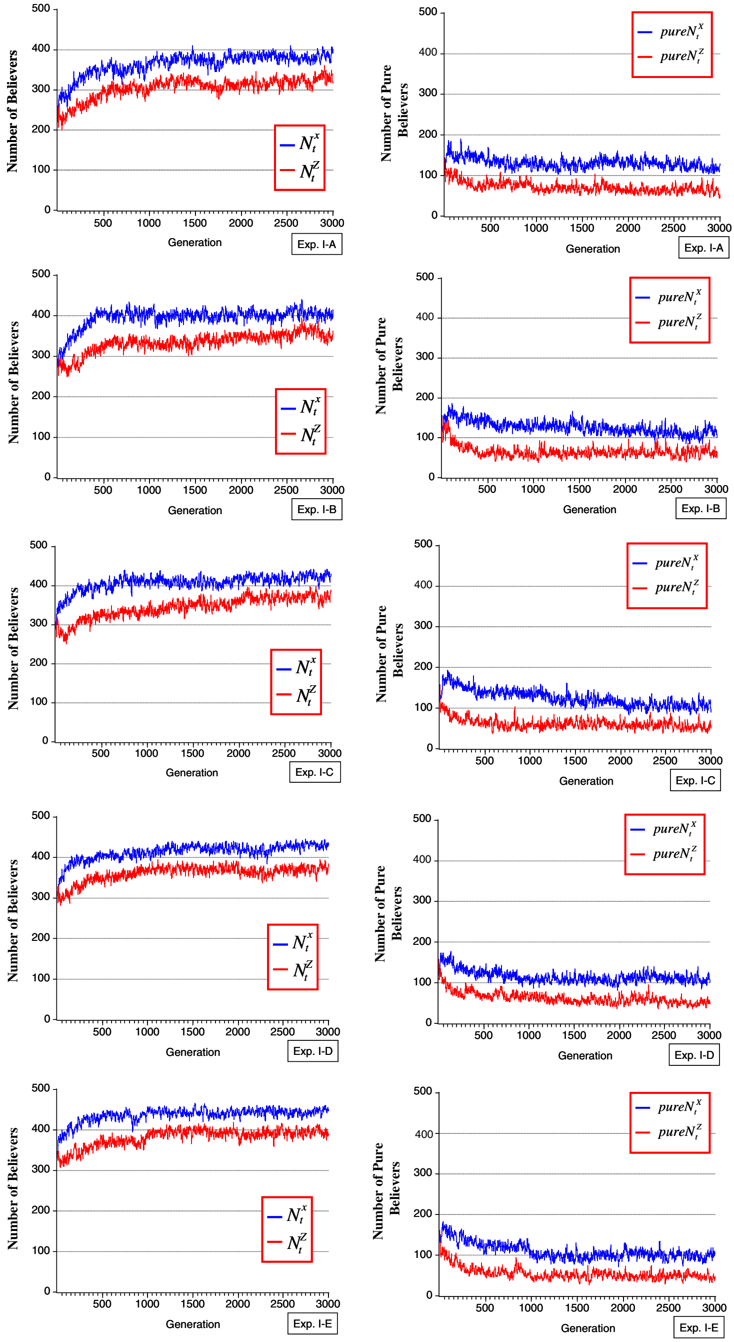

Fig. 1 Time series plots of sunspot and fundamental believers: series I in the *left panel* are the time series plots of the sunspot believers (N_t^Z) and fundamental believers (N_t^X), whereas in the *right panel* are the time series plots of the pure sunspot believers (N_t^{PZ}) and pure fundamental believers (N_t^{PX}). Notice that each plot is performed by taking the average of the respective number over 50 runs for each experiment

Table 3 Summary statistics of the number of various kinds of believers: series I

Market	Statistic	Number of believers					
		N^Z	N^X	$N^Z \bigcap N^X$	N^{PZ}	N^{PX}	N^C
I-A	Mean	314.33	373.61	245.53	68.81	128.08	57.58
	Std. Dev.	12.97	15.05	17.59	9.53	9.93	9.99
I-B	Mean	342.20	402.15	294.24	47.96	107.90	49.89
	Std. Dev.	14.19	10.45	22.01	17.04	21.48	18.81
I-C	Mean	352.96	413.80	294.25	58.72	119.56	27.48
	Std. Dev.	18.34	11.19	22.01	9.08	15.27	6.80
I-D	Mean	360.69	419.60	308.48	58.69	111.13	21.71
	Std. Dev.	19.05	10.93	16.48	9.46	9.50	4.86
I-E	Mean	388.61	441.75	337.86	50.75	103.89	7.49
	Std. Dev	12.50	8.94	18.35	8.71	12.56	2.65

The summary statistics presented here are calculated over the 50 runs for each experiment. N^Z and N^X denote the number of the sunspot and fundamental believers, whereas N^{PZ} and N^{PX} denote the number of the pure sunspot and pure fundamental believers, respectively. $N^Z \bigcap N^X$ refers to the number of agents who believe in both sunspots and fundamentals. For N^C, see footnote 7

strong to the extent that it is independent of sunspot density. The result that sunspots can hardly find their way into the market shows that sunspots as signals to coordinate agents' expectations are quite implausible. This finding is basically in line with the basic tone of the existing literature: "While there is a large theoretical literature on when sunspots may matter, empirical evidence that expectationally driven randomness is at work in real-world markets has been scarce." (Marimon et al. 1993, pp 74–75).

6 Would the design of learning matter?

In our previous two series of experiments, the population of sunspot believers is uniformly dominated by the population of the fundamental believers, and the market, as a whole, can successfully distinguish the fundamentals from the sunspots. However, given the empirical fact that the sunspot does not Granger-cause returns, it remains a puzzle why the proportion of sunspot believers, on average, ranges from 60 to 80% in the two previous series of experiments. Why are there so many agents who would like to use the sunspot to forecast returns, and why do 8 to 16% of agents use the sunspot *only*, even though the causality test shows that sunspots do not help forecast returns? Does genetic programming, as a selection mechanism to decide who shall survive and who shall become extinct, really work?

As mentioned in Sect. 3.3, one possible solution to get out of this conundrum is to make a distinction between the "insiders" (artificial agents) and "outsiders" (us as econometricians) of the model. What concerns "insiders" is how well their forecasting models perform over the very recent past, which has a very short time horizon, whereas an econometric test is conducted with a quite a long time horizon. Therefore, the

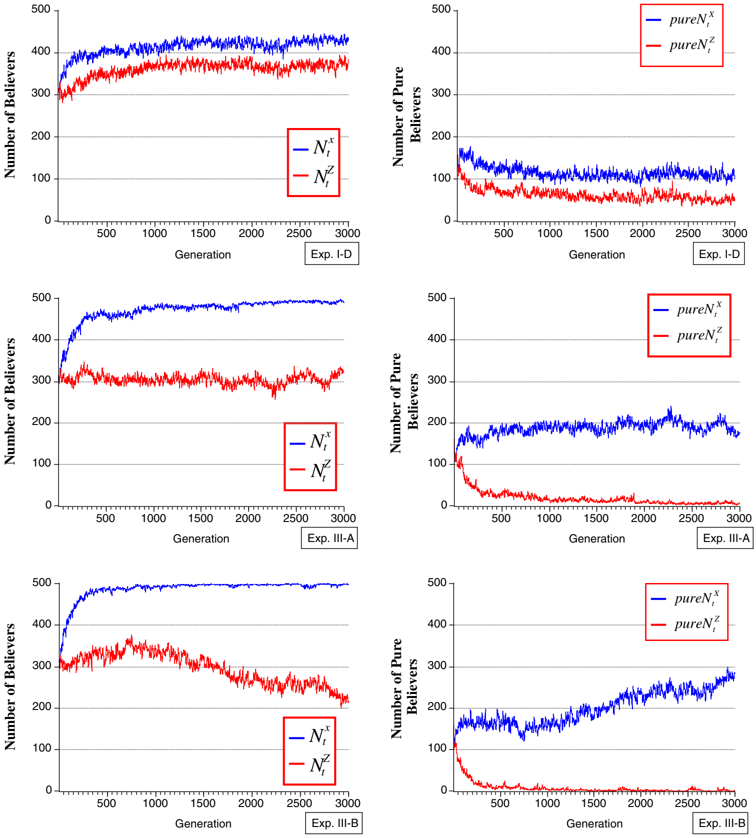

Fig. 2 Time series plots of sunspot and fundamental believers: series III Interpretation is the same as that of Fig. 1

survival pressure concerning agents and the causality test concerning econometricians can be two quite different things. To reconcile the difference, we increase the time horizon in the third series of experiments.

From Fig. 2, we can see how the time horizon significantly impacts the number of sunspot believers and fundamental believers. The left half of Fig. 2 shows both N^Z and N^{PZ} decreases with the increase in the time horizon H. In addition, from Table 4, the mean number of sunspot believers decreases from 360 down to 290 when H increases from 10 to 50, while, in the meantime, the mean number of fundamental believers increases from 419 to 495. The right half of Fig. 2 gives an even better picture. When the time horizon increases, we can see that, right from the beginning of the evolution, the number of pure sunspot believers declines at an even faster rate, which clearly indicates that agents who rely on sunspots only will find it increasingly difficult to survive. On the other hand, the number of agents who completely abandon sunspots increases from 111 to 207, when H is extended from 10 to 50. In other words, two fifths of the whole population become "rational" in the sense that they correctly identify what the extrinsic uncertainties are.

Table 4 Summary statistics of the number of various kinds of believers: series III

Case	Statistic	Number of believers					
		N^Z	N^X	$N^Z \bigcap N^X$	N^{PZ}	N^{PX}	N^C
I-D	Mean	360.69	419.60	308.48	58.69	111.13	21.71
	Std. Dev.	19.05	10.93	16.48	9.46	9.50	4.86
III-A	Mean	303.00	481.65	290.32	12.68	191.33	5.67
	Std. Dev.	11.27	10.72	13.87	7.33	11.75	4.27
III-B	Mean	290.53	495.22	287.25	3.28	207.96	1.51
	Std. Dev.	39.24	4.23	37.73	3.23	40.00	1.67

Notations and remarks are the same as those in Table 3

7 Concluding remarks

In this paper, we address the plausibility of sunspot equilibria in agent-based artificial stock markets. Agent-based modeling provides us with a control laboratory, while software agents provide us with a means by which both sunspot variables and coordination processes are identified. Hence, according to Duffy and Fisher (2005), the "evidence," if there is any, would be *direct*. The three series of experiments conducted in this paper show that the plausibility of observing sunspot equilibria is low. In an environment of stochastic simulation, the probability of observing sunspot equilibria in the artificial stock market is only about 10%. Nevertheless, the population of sunspot believers is much bigger than this figure: 60 to 80% of agents are sunspot believers, and 8 to 16% are even stronger ones. These two inconsistent figures indicate that what matters is not the number of sunspot believers, but how sunspots are "interpreted" by these believers. Our finding does show that different interpretations exist among sunspot believers, for example, the strong ones and the weak ones. Therefore, even though there is such a great population of sunspot believers, it does not necessarily mean that sunspot equilibria as an aggregate phenomenon will emerge from this micro-structure. In terms of Duffy and Fisher (2005), it is the semantics of sunspots that matters. By explicitly taking the coordination processes of heterogeneous agents into account, we confirm this finding, although with a quite different set-up.

Acknowledgments The authors are grateful to one anonymous referee for the helpful suggestions. NSC research grant No. 95-2415-H-004-002-MY3 is gratefully acknowledged.

Appendix A: AIE-ASM

The entire operation of AIE-ASM is controlled by three sets of parameters. The first part, as shown in the first block of Table 5, provides the basic description of the stock market. It includes the number of outstanding shares, the initial size of the riskless asset, the return on the riskless asset (the fixed interest rate), and the stochastic process of the return (dividend) on the risky asset. It also specifies how the price is determined. The use of GP makes the forecasting function generally non-linear. Therefore, deriving

Table 5 Parameters of the stock market

The stock market	
Shares per capita (h)	1
Initial money supply per capita (m)	100
Risk-free interest rate (r_f)	0.1
Stochastic process (D_t)	*i.i.d.* Normal($\mu = 10, \sigma^2 = 4$)
Price adjustment function	tanh
Price adjustment (β_1)	10^{-4}
Price adjustment (β_2)	0.2×10^{-4}
Traders	
Number of traders	500
Degree of ARA (λ)	0.5
θ_1	0.5
θ_2	10^{-4}
θ_3	0.0133
Sample size of $\sigma^2_{(t\|n_1)}$ (n_1)	10
Criterion of fitness (traders)	`Increments in wealth`
Evaluation cycle (n_2)	1
Sample size (n_3)	10
Search intensity	5
Business school	
Number of faculty members	500
Proportion of trees initiated	
by the full method	0.5
by the grow method	0.5
Function set	$\{+, -, \times, \div, \sqrt{}$, sin, cos, exp, Rlog, abs$\}$
Terminal set	see Table 1
Selection scheme	`Tournament selection`
Tournament size	2
Proportion of offspring trees created	
by reproduction (p_r)	0.1
by crossover (p_c)	0.7
by mutation (p_m)	0.2
Probability of mutation	0.0033
Mutation scheme	`Tree mutation`
Replacement scheme	`Tournament selection`
Maximum depth of tree	17
Number of generations	5,000
Maximum in the domain of RExp	1,700
Criterion of fitness (faculty)	`MAPE`
Evaluation cycle (m_1)	20
Time horizon for accuracy measure (H)	10 (Experiment Series I, II)
	10, 30, 50 (Experiment Series III)

the aggregate demand can be a daunting task, not to mention the solution to the market-clearing condition. Hence, the Walrasian tatonnement scheme is replaced by the rationing scheme in the price determination. The last three parameters in the first block specify how the price will be adjusted given the current size of the market disequilibrium.

The second block in Table 5 provides the basic characteristics of traders, including their risk attitude, expectation formation and adaptive behavior. We assume that all traders have a CARA (constant absolute risk aversion) utility function with an ARA (absolute risk aversion) coefficient of 0.5. The next four parameters, θ_1, θ_2, θ_3 and n_1, are the parameters used to determine the traders' own perceived excess risk-adjusted return, and hence their demand for the stock. Among these four, the first two are used to shape the traders' perceived excess return, whereas the next two are used to shape the traders' perceived risk.[8] The next two parameters are associated with the adaptive behavior of traders. The driving force behind the traders' adaptation comes from peer pressure as well as the traders' self-expectations. When traders are not satisfied with their current performance in terms of incremental wealth, they will consider getting a better forecasting model by searching in the *business school*. The last two parameters of this block determine how intensively the trader will search in the business school.[9]

We have mentioned that traders' adaptations are conducted via the business school. The third block in the Table 5 is related to the behavior of the business school. Basically, the business school is composed of a number of faculty members, say, 500 members. They are competing for the best forecasting model, and of course, their behavior is also adaptive and is driven by genetic programming. Therefore, all of the parameters listed here are the control parameters involved in running GP. Their details can be found in Chen and Yeh (2001).[10]

In this paper, we set the values of all parameters almost identically to those in Chen and Yeh (2001) except the terminal set and the time horizon due to different research focuses. To be able to study the possible influence of sunspots, we add sunspot variables to the terminal set, while in the meantime we simplify the included fundamental variables, as shown in Table 1. In addition, as mentioned in Sect. 3.3, in the third series of experiments, we would like to inquire into the possible connections between the time horizons and survival pressure, and their further impact on the sunspot believers. We, therefore, consider three different time horizons in Experiment Series 3.[11]

Appendix B: Causality tests

The Granger causality tests conducted in this paper are the same as the one used in Chen and Liao (2005), where a short but comprehensive review of the concept of

[8] For the details, see Chen and Yeh (2001), p 373.

[9] For the details, see Chen and Yeh (2001), Sect. 2.5.

[10] A menu-like introduction to **AIE-ASM** *Ver. 2* can also be found in Chen et al. (2002).

[11] In Chen and Yeh (2001), we do not use the term "time horizon;" instead, we use the sample size for calculating the MAPE (mean absolute percentage error). In addition, the notation m_2 is used there rather than the H used in this paper.

causality and the development of the causality test is also provided. Basically, we apply two versions of the causality tests in this paper. The version based on Granger (1969) can only be applied to test the *linear* causal relationship. Hence, an extension to the *non-linear* case, given by Hiemstra and Jones (1994), is also applied here.

References

Bray MM, Kreps DM (1987) Rational learning and rational expectations. In: Feiwl GR (ed) Arrow and the ascent of modern economic theory. NYU Press, New York

Cass D, Shell K (1983) Do sunspots matter? J Econ Theory 25:380–396

Chen S-H, Yeh C-H (2001) Evolving traders and the business school with genetic programming: a new architecture of the agent-based artificial stock market. J Econ Dyn Control 25:363–393

Chen S-H, Yeh C-H (2002) On the emergent properties of artificial stock markets. J Econ Behav Organ 49:217–229

Chen S-H, Liao C-C (2005) Agent-based computational modeling of the stock price-volume relation. Inf Sci 170:75–100

Chen S-H, Yeh C-H, Liao C-C (2002) On AIE-ASM: software to simulate artificial stock markets with genetic programming. In: Chen S-H (ed) Evolutionary computation in economics and finance. Physica-Verlag, Heidelberg, pp 107–122

Dawid H (1996) Learning of cycles and sunspot equilibria by genetic algorithms. J Evol Econ 6:361–373

Duffy J, Fisher E (2005) Sunspots in the laboratory. Am Econ Rev 95(3):510–529

Evans GW, Honkapohja S (1996) Least squares learning with heterogeneous agents. Econ Lett 53:197–201

Evans GW, McGough B (2005) Stable sunspot solutions in models with predetermined variables. J Econ Dyn Control 29:601–625

Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37:424–438

Hens T (2000) Do sunspots matter when spot market equilibrium are unique. Econometrica 68(2):435–441

Hiemstra C, Jones JD (1994) Testing for linear and nonlinear Granger causality in the stock price—volume relation. J Finance 49:1639–1664

Kajii A (1997) On the role of options in sunspots equilibria. Econometrica 65(4):977–986

LeBaron B (2001) Evolution and time horizons in an agent-based stock market. Macroecon Dyn 5:225–254

LeBaron B (2006) Agent-based computational finance. Forthcoming In: Tesfatsion L, Judd KL (eds) Handbook of computational economics, vol 2, Elsevier, North Holland

Marimon R, Spear S, Sunder S (1993) Exceptionally driven market volatility: an experimental study. J Econ Theory 61(1):74–103

Negroni G (2005) Eductive expectations coordination on deterministic cycles in an economy with heterogeneous agents. J Econ Dyn Control 29:931–952

Tesfatsion L, Judd KL (ed) (2006) Handbook of computational economics, vol 2. Elsevier, North Holland

Woodford M (1990) Learning to believe in sunspots. Econometrica 58:277–307