An Exploration of Neo-Austrian Theory Applied to Financial Markets
Author(s): Harald Benink and Peter Bossaerts
Source: *The Journal of Finance*, Vol. 56, No. 3 (Jun., 2001), pp. 1011-1027
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/222540
Accessed: 24/11/2009 12:04

THE JOURNAL OF FINANCE • VOL. LVI, NO. 3 • JUNE 2001

An Exploration of Neo-Austrian Theory Applied to Financial Markets

HARALD BENINK and PETER BOSSAERTS*

ABSTRACT

We attempt to translate Neo-Austrian ideas about the workings of financial markets, as originally advanced by F. A. Hayek, into the standard probabilistic language of modern finance. We focus on an apparent paradox, namely the insistence of Neo-Austrians on *order* (i.e., stationarity) together with *ever-reemerging inefficiencies*. The paper's findings have implications beyond Neo-Austrian theory: They demonstrate how easy it is to reject market efficiency, but how much more difficult it is to discern the *nature of the inefficiency*. We illustrate our findings with price data from the U.S. Treasury bill market over the period 1962 to 1999. There is ample evidence that the price of a three-month Treasury bill is not a random walk, yet the sign of the average price change is erratic, so that inference about the nature of the inefficiency is unreliable.

THE NEOCLASSICAL RATIONAL EXPECTATIONS view of financial markets is one of continuous equilibrium with informationally efficient prices. Empiricists have recently questioned the validity of this model, pointing to evidence of inefficiencies. Alternative views have been presented to better match the empirical evidence. One of these considers the market to be continuously evolving from one inefficiency to another, never attaining the perfect, efficient equilibrium, yet strongly attracted towards it. Creative investors track and exploit profit opportunities generated by continuous shocks in a never-ending cycle. The result would be a stable process with pronounced regularities.

Such a view follows from the Neo-Austrian theory of market processes, a recent rereading of the ideas of Friedrich Hayek (see, e.g., Hayek (1937, 1945, 1948, 1978), Littlechild (1982), Rizzo (1990), Kirzner (1997)). According to Neo-Austrian theory, a competitive market provides a systematic set of forces, put in motion by entrepreneurial alertness (eagerness to make money), which tend to reduce the extent of ignorance among market participants. The resulting knowledge is not perfect; neither is ignorance nec-

* Harald Benink is from the Rotterdam School of Management, Erasmus University, Rotterdam, The Netherlands. Peter Bossaerts is from the California Institute of Technology and Center for Economic Policy Research. This paper is based on an earlier one written by the second author, entitled "Time Series Analysis Of Inefficient Financial Markets." That paper benefited from comments during seminars at the University of New South Wales (Australia), at LIFE of Maastricht University (The Netherlands), and at the Weierstrass Institute of the Humboldt Universität (Germany). Many thanks to Israel Kirzner, Harald Uhlig, the editor and an anonymous referee for comments and criticism.

essarily invincible. *Equilibrium* (read: market efficiency) *is never attained*, yet the market does exhibit powerful tendencies towards it. The fact that equilibrium is never attained is attributed to an erratically changing world where traders realize that their *knowledge is imperfect*. At the same time, the changes are never so extreme as to frustrate the emergence of powerful and pervasive *economic regularities*.

For the reader who is not familiar with Neo-Austrian thinking, let us be more specific. According to Hayek, the problem of economic choice and ultimately the analysis of economic behavior in neoclassical theory is oversimplified, because it has been reduced to optimal behavior under constraints that agents are supposed to be very familiar with. These constraints concern (1) preferences, (2) production and market technology, and (3) resources. In contrast, the Neo-Austrian view stresses that fundamental uncertainty and ignorance exist regarding these constraints. This uncertainty and ignorance is claimed to lead to disequilibrium, and disequilibrium itself generates further uncertainty and ignorance regarding the constraints. Disequilibrium thereby becomes self-enforcing and permanent.

However, alert participants in the market process, whom the Neo-Austrians define as "entrepreneurs," do try to get a (necessarily incomplete) picture of the nature of the disequilibrium in the marketplace, because disequilibrium generates profit opportunities. The actions of these entrepreneurs produce the very signals that are needed to reduce disequilibrium. However, due to continuous change in the constraints, equilibrium is never achieved.

The Neo-Austrian view lacks the mathematical rigor that has become characteristic of neoclassical economics. Certainly when applied to financial markets, the Neo-Austrian approach seems to be vague at first. This paper should be viewed as an attempt to come to an understanding of it in the standard probabilistic language of finance. Some may disagree with our translation. In that case, the reader should consider this paper as providing an alternative view of the operation of financial markets, inspired by, but not necessarily adequately reflecting, Neo-Austrian economic thinking. As will be discussed shortly, the paper carries a distinct message about the empirical analysis of inefficient financial markets, independent of its goal to translate Neo-Austrian theory into standard financial language.

The most intriguing aspect of Neo-Austrian financial economics is its insistence on regularities, that is, on *order*, while at the same time *rejecting market efficiency*. We take the position that order must be translated into the probabilistic concept of *stationarity*. Loosely speaking, this means that once an event occurs, it is certain to recur at one point in the future, even if that point is unpredictable. Such events are regularities. But, if our translation of order into the notion of stationarity is correct, the Neo-Austrian view does present a puzzle when applied to financial markets. It is a simple implication of the ergodic theorem that stationarity enables one to consistently estimate all patterns in the data. This means that investors could uncover all inefficiencies, and exploit them, only to eventually lead the market towards an efficient equilibrium, something that Neo-Austrians claim will never be attained.

This puzzle is all the more forceful if one appreciates Neo-Austrians' insistence that the failure of markets to reach the classical equilibrium ought not to be attributed to costs of any nature (adjustment costs, information costs, trading costs, etc.). As mentioned before, the nonconvergence has its origin in limitations of knowledge, however vague this expression may seem. By this Neo-Austrians do not mean that certain parameters of the price processes are unknown; stationarity would contradict that.

We propose the following resolution for the puzzle. A fully rational, risk-averse investor, having detected evidence against market efficiency, will not exploit this if he/she fails to understand the nature of the inefficiencies. To determine the latter, he/she wants reliable confidence statements about the inference she draws from the data. There are situations where these statements cannot be made, namely, when the assumptions behind classical statistics are violated.[1] To obtain correct confidence intervals, more is needed than just stationarity. In other words, the return process may still lack the necessary conditions to attach precise probabilities to statements about the true nature of the inefficiencies. Faced with a situation where inference is unreliable from a classical point of view (yet the market is clearly inefficient), our investor may as well use rules of thumb to exploit the inefficiency. This situation is precisely the one that Neo-Austrians specify as being typical of financial markets.

Our investor resorting to using rules of thumb corresponds to the Neo-Austrian view that agents are not ignorant, but that they appreciate that knowledge cannot be perfect. Successful trading rules may be myopic and simple, and could be devoid of comprehensive economic reasoning.[2] They are, however, examples of reasonable reactions of a rational investor who realizes that inference may at times be limited even if the return process is stationary.

In this paper, we construct an example of an economy with a continuously inefficient financial market. We adjust the *memory* of investors' trading rules in order to generate a market process that can be characterized as stable, cycling from one inefficiency to another. Longer memory of trading rules affect equilibrium returns, to the extent that they cause violations of the usual mixing conditions needed for correct inference. Based on an analysis of average price changes, an investor will with high likelihood reject efficiency,

[1] We focus on classical inference because we ultimately want to say something about standard empirical analysis in finance, which, with few exceptions, followed the classical route.

[2] Charles Plott often likens traders to fish, which are generally extremely good swimmers without knowing any hydrodynamics. In fact, fish are not even interested in hydrodynamics, just like many traders often show a lack of interest in academic finance. We provide a foundation for this attitude of indifference.

yet the sign of the average is unreliable in predicting the sign in independent future replication. As a consequence, classical statistics cannot reliably assess the inefficiencies, and investors may as well use rules of thumb (which they do in our economy).

There are other ways to adjust investors' trading rules that would generate the type of long memory that is required to violate the usual mixing conditions on which classical asymptotic inference is based. We chose to focus on the memory of the trading rule, because its link with the memory of equilibrium returns is straightforward.

The main tool of analysis is simulation. Compared to analytical proofs (which, if possible, would involve advanced probability), simulation makes it easier for our message to be understood and, hence, to be appreciated by a wider audience. We study distributions of standard test statistics that have often been used to verify the nature of alleged inefficiencies in financial markets. We illustrate when these statistics are ill behaved, thus leading to erroneous inference. The reader who is interested in analytical examples of failures of large-sample theoretical results is referred to Bossaerts (1995).

Our model is an illustrative, mathematical example of what we interpret Neo-Austrian financial economics to be about. Alternatively, one can view it as an appropriate vehicle with which to evaluate the statistical implications of market inefficiency.

Indeed, it has recently become fashionable not only to question market efficiency, but also to explain what is wrong (e.g., Haugen (1995)). The evidence is interpreted by means of classical inference. The present study, however, illustrates that classical statistical inference may be erroneous when applied to inefficient financial markets.

As far as recent empirical work on the nature of market inefficiency is concerned, this paper could be interpreted as raising a paradox: *we know how to reject market efficiency, but, once we reject it, we may not necessarily be able to infer why*. It is unnecessary to add that trading strategies derived from an econometric analysis of inefficient financial markets may therefore be ill advised. The same can be said about public policy prescriptions.

A final note on related work. Mordecai Kurz has recently been working on a notion of rational beliefs equilibria that allows for differences in beliefs as long as these cannot be reconciled because of nonstationary (i.e., nonrecursive) features of the data (see, e.g., Kurz (1994)). This idea is closely related to ours: Our approach is also based on features of the data that cannot be verified. Neo-Austrians seem to insist on stationarity, however, so we had to rely on other nonverifiable features of the data.

Marcet and Nicolini (1997) have recently determined lower bounds on efficiency of boundedly rational mechanisms. Some of these lower bounds are quite capable of explaining hyperinflationary periods. Arthur et al. (1997) and Le Baron (1999) investigate the dynamics of a simulated market populated with robots who use trading algorithms that are adapted based on past performance. Their aim is to document the time-series features of such a market, in particular, predictability, excess volatility, volatility persistence, and leptokurtosis. Unlike the present paper, however, they do not ask to what extent some of the features they observe can be objectively learned without knowing how the data were generated in the first place.

The remainder of the paper is organized as follows. The next section describes the economy. Section II discusses the case where traders' investment rules have short memories. Section III investigates longer-memory trading rules. In Section IV, we illustrate how z-statistics computed from average daily price changes of the three-month Treasury bill over the period 1962 to 1999 fail to conform to a central limit theorem. Section V concludes.

I. The Economy

Consider the market for an infinite-lived asset that pays a periodic, exogenous dividend d_t at time t ($t = 1,2,\ldots$). The asset's time-t price is p_t, and its return is x_t ($= p_t + d_t$). Time is split in *epochs*, indexed $e = 1,2,\ldots$ An epoch lasts τ periods. During an epoch, investors follow fixed trading rules. Trading rules, however, may differ across epochs. The *beginning* of each subperiod within an epoch is indicated with σ. At the start of the first subperiod, $\sigma = 0$. The *end of the last subperiod* is also indicated with σ, namely $\sigma = \tau$. When an outsider ("econometrician") contemplates a return history, it will have length T, where T is a multiple n of τ, that is, $T = n\tau$.

The dividend follows a simple ARMA(1,2) process with a $N(0,1)$ error. We arbitrarily set the autoregressive coefficient equal to 0.8, and the moving average coefficients equal to 0.4 and 0.3. Investors do not know these parameters. Even if they knew them, they still would not know an important piece of information that is needed to determine the return process, namely the time-series characteristics of the equilibrium price process.

We do not specify the trading process in detail. Instead, we merely assume that prices are set in the following way. For each epoch, there exist parameters a_e and b_e such that

$$p_{e,\sigma+1} = p_{e,\sigma} + a_e + b_e(x_{e,\sigma} - p_{e,\sigma-1}), \tag{1}$$

where

$$\begin{aligned} p_{e,1} &= \begin{cases} p_{e-1,\sigma+1} & \text{if } e > 1 \\ p_{1,0} + a_1 + b_1 x_0 & \text{if } e = 1 \end{cases} \\ p_{e,0} &= \begin{cases} p_{e-1,\sigma} & \text{if } e > 1 \\ N(2,1)^2 & \text{if } e = 1. \end{cases} \end{aligned} \tag{2}$$

The parameters a_e and b_e are determined as a weighted average of (1) a_{e-1} and b_{e-1}, respectively, and (2) $\tilde{a}_e$ and $\tilde{b}_e$, respectively, which are set such that

$$\frac{1}{\tau}\sum_{\sigma=1}^{\tau}(\tilde{x}_{e-1,\sigma+1}-\tilde{p}_{e-1,\sigma})=0, \tag{3}$$

$$\frac{1}{\tau}\sum_{\sigma=1}^{\tau}(\tilde{x}_{e-1,\sigma+1}-\tilde{p}_{e-1,\sigma})(\tilde{x}_{e-1,\sigma}-\tilde{p}_{e-1,\sigma-1})=0, \tag{4}$$

where

$$\begin{aligned}\tilde{p}_{e-1,\sigma} &= \begin{cases}\tilde{p}_{e-1,\sigma-1}+\tilde{a}_e+\tilde{b}_e(\tilde{x}_{e-1,\sigma-1}-\tilde{p}_{e-1,\sigma-2}) & \text{if } \sigma>0\\ p_{e-1,\sigma} & \text{if } \sigma\le 0\end{cases}\\ \tilde{x}_{e-1,\sigma} &= \tilde{p}_{e-1,\sigma}+d_{e-1,\sigma}.\end{aligned} \tag{5}$$

Inspection of these equations should reveal that $\tilde{a}_e$ and $\tilde{b}_e$ are the parameter values such that, if investors had used trading rules in epoch $e-1$ that would have led to these values, returns (1) would be zero on average, and (2) would exhibit zero sample autocorrelation. Effectively, we are assuming that investors change their trading rules such that the price process in the new epoch (e) uses parameter values that are an average of last epoch's ($e-1$) parameter values, and the parameter values that would have made last period's prices unbiased and past returns uncorrelated in time.

To avoid obvious problems with nonstationarity, if $|\tilde{b}_e|\ge 1$, we constrain it such that $|\tilde{b}_e|<0.9$ always. Also, a_1 and b_1 are $N(0,0.01)$ draws.

The weight w_e that investors put on the new information $\tilde{a}_e$ and $\tilde{b}_e$ (the complementary weight is applied to a_{e-1} and b_{e-1}, respectively) will be left unspecified at the moment.[3] Later on, we will see how each formulation leads to different return characteristics, and, hence, inference properties. Suffice it to say that, unlike in Bayesian updating or least-squares learning, the weights will be specified to be *stationary*: Only information collected during epochs $e-1$ and $e-2$ determine w_e, and, not, for instance, information accumulated over the epochs prior to epoch $e-2$. In particular, the weights will not depend on time. This means that the problematic inference we will uncover cannot be attributed to nonstationary learning rules.

[3] A referee reminded us that, in principle, one would like to choose the weighting schemes with some kind of optimality in mind. More specifically, one would like to be able to state that the ensuing learning rules are the best reaction of agents to a multitude of possible worlds of which the actual one is just a draw. Evolutionary arguments could be appealed to in order to justify agents' knowing optimal rules. At present, however, such arguments are not yet within reach.

To determine the theoretical properties of the equilibrium return process, it is useful to know that $\tilde{b}_e$ can be computed as minus the slope coefficient in a regression of future dividends onto past returns. Because the (exogenous) dividend process is ARMA(1,2), this means that the past return is used as a proxy for the infinite series of past dividends. The evolution of $\tilde{b}_e$ over the epochs can therefore theoretically be analyzed as the outcome of a sequence of errors-in-variables regressions.

It should be emphasized that the pricing rule entertained here is by no means arbitrary. Prices are linearly related to past returns. It is easy to verify that, given the parametrization of the dividend process, returns will be close to zero on average and exhibit negligible serial correlation when $a_e = 0$ and b_e has a value between $-2/5$ and $-3/8$. In other words, for certain values of its parameters, the pricing rule is approximately rational expectations.

Let us now investigate the impact of different weightings w_e on the time-series characteristics of the equilibrium return processes and the statistics an outsider would compute with the purpose of exploiting any inefficiencies. We will focus on the sample average return and the corresponding z-statistic, denoted $\bar{x}_T$ and z_T, respectively. If the average return is found to be significantly positive or negative (as indicated by the level of the z-statistic), the outsider would infer that there are profit opportunities from purchasing the asset. That is, he/she infers that the market is inefficient, and would exploit the inefficiency with a trading strategy that depends on the sign of the average return.

$\bar{x}_T$ is computed on the basis of the complete history of returns:

$$\bar{x}_T = \frac{1}{T} \sum_{t=1}^{T} x_t. \tag{6}$$

The z-statistic is computed as follows:

$$z_T = \frac{\dfrac{1}{\sqrt{T}} \displaystyle\sum_{t=1}^{T} x_t}{\sqrt{V_T}}, \tag{7}$$

where V_T is the estimated asymptotic standard deviation of $(1/\sqrt{T})\sum_{t=1}^{T} x_t$. Because the return series clearly exhibits first-order autocorrelation (autocorrelation coefficients of higher order are almost always insignificant), V_T is computed as follows:

$$V_T = \frac{1}{T} \sum_{t=1}^{T} (x_t - \bar{x}_T)^2 + 2\,\frac{1}{T} \sum_{t=2}^{T} (x_t - \overline{x^{+}}_{T-1})(x_{t-1} - \overline{x^{-}}_{T-1}), \tag{8}$$

where

$$\overline{x^{+}}_{T-1} = \frac{1}{T-1}\sum_{t=2}^{T} x_t, \tag{9}$$

and

$$\overline{x^{-}}_{T-1} = \frac{1}{T-1}\sum_{t=2}^{T} x_{t-1}. \tag{10}$$

We take the epoch length to be 60, that is, $\tau = 60$.

We investigate two cases. First, we assume short-memory trading rules. This will produce a benchmark case: There will be no sign of inefficiency (when measured with average returns), which means that the inefficiencies have to be discovered with other means. Second, we introduce long-memory trading rules. Here, the average returns will often correctly indicate that there are inefficiencies, but, as we shall demonstrate, the inference is unreliable. In particular, the sign of the average return cannot be relied upon: If a given history produces a significantly positive return, this is no indication that the return will be positive in an independent replication; as a matter of fact, it is equally likely to be *significantly negative*. Consequently, an outsider who decides whether a market is inefficient based on the average return will correctly infer that there are inefficiencies, but cannot rely upon the sign of the average return to initiate strategies meant to exploit them.

II. Case I: Short-memory Trading Rules

Let us first look at a case where market participants are rather forgetful. In particular, we set

$$w_e = 0.5\,\frac{|s_{e-1} - s_{e-2}|}{s_{e-2}}, \tag{11}$$

where s_e denotes the sample standard deviation of the returns in epoch e (we chose $s_0 = 1$, $w_1 = 1$). Hence, the weight put on the revised pricing is one-half the percentage change in the standard deviation. This rule was chosen because it implies that agents put far more weight on the most recent epoch when there is a large change in the volatility.

The analysis will cover up to 1,000 epochs of 60 periods each. But we will discard the first 200 epochs, to be sure that our conclusions are not driven by transient nonstationarity. (Indeed, we draw initial values for the dividends, parameters and prices from rather arbitrary distributions.)

To get an idea of the nature of the resulting economy, Figure 1 plots the evolution of the cumulative dividend, the price, and the two parameters (a_e and b_e) in one draw, over epochs 201 to 300. Because each epoch lasts 60 periods, this amounts to 6,000 periods in total. Two points are worth making. First, prices generally decline in periods of mostly positive dividends

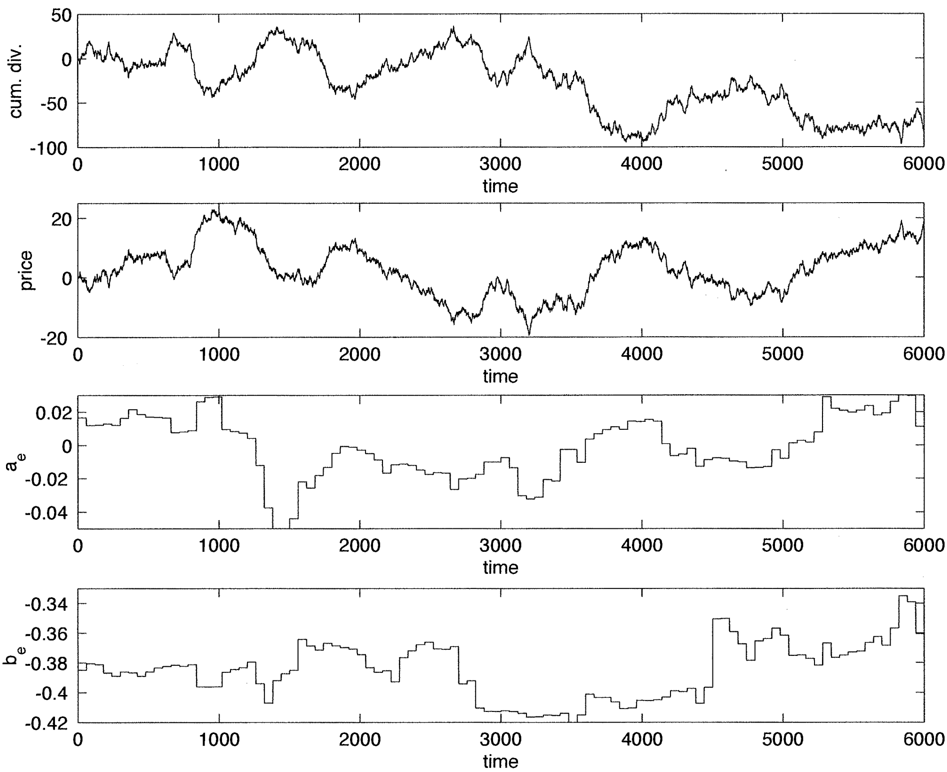

Figure 1. Dividends and prices with short-memory traders. Evolution of cumulative dividend (top plot), prices (second plot), parameters a_e (third plot), and b_e (bottom plot) in one simulation of an economy with short-memory traders (the model in equation (11)). Total time length displayed: ($T =$) 6,000 periods, or ($n =$) 100 epochs of ($\tau =$) 60 periods each. The first 200 epochs (preceding the ones shown) are deleted to eliminate the impact of the initial distribution from which the dividend process is started.

(increases in the cumulative dividend). This is a consequence of our pricing rules: parameters change such that mean returns (price changes plus dividends) remain close to zero. Second, the two parameters of the pricing equations (a_e and b_e) change regularly, but still exhibit rather long memory. The subsequent figure will reveal how (whether) this affects inference.

Figure 2 proves that the long memory in the parameters of the pricing equation does not affect asymptotic inference. It displays boxplots of the z-statistics computed from average returns (z_T in the previous section) for increasing sample sizes. For comparison, a boxplot is also shown for standard normal draws from the same random number generator that was used for the simulations themselves. Already with 100 epochs (i.e., 6,000 periods), the z-statistics are essentially standard normal. Hence, inference based on the usual asymptotic theory can be relied on.

An outsider would infer that there are no inefficiencies with a frequency equal to that if there were indeed none. There are inefficiencies; they simply

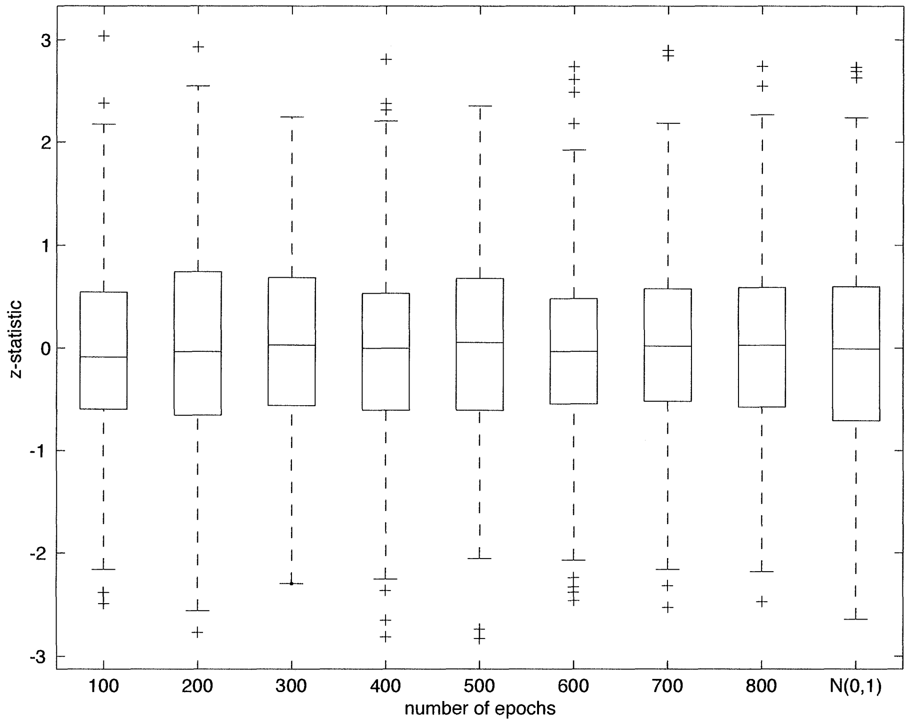

Figure 2. Returns with short-memory traders. Boxplots of distributions of the z-statistic of the return (dividend plus price change) across 400 repetitions of an economy with short-memory traders (the model in equation (11)). Shown are distributions for an increasing number of epochs ($= n$), from 100 to 800. One epoch lasts 60 periods ($= \tau$). In each simulation, the first 200 epochs (preceding the ones from which the z-statistics were computed) are deleted to eliminate the impact of the initial distribution from which the dividend process is started. The z-statistic is adjusted for mild first-order autoregression in the return. For comparison, the last boxplot shows the distribution of 400 random draws from the $N(0,1)$ distribution. These draws are based on the same random number generator on which the other simulations are based.

cannot be discerned from the z-statistic computed from average returns. This will be different in the next case, where inefficiencies will be discernible with high frequency, but the sign of the average return is an unreliable indicator of the nature of the inefficiencies.

III. Case II: Longer-memory Trading Rules

Consider an alternative weighting rule:

$$w_e = \begin{cases} \dfrac{|s_{e-1} - s_{e-2}|}{s_{e-2}}, & \text{if } \dfrac{|s_{e-1} - s_{e-2}|}{s_{e-2}} > 0.25, \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

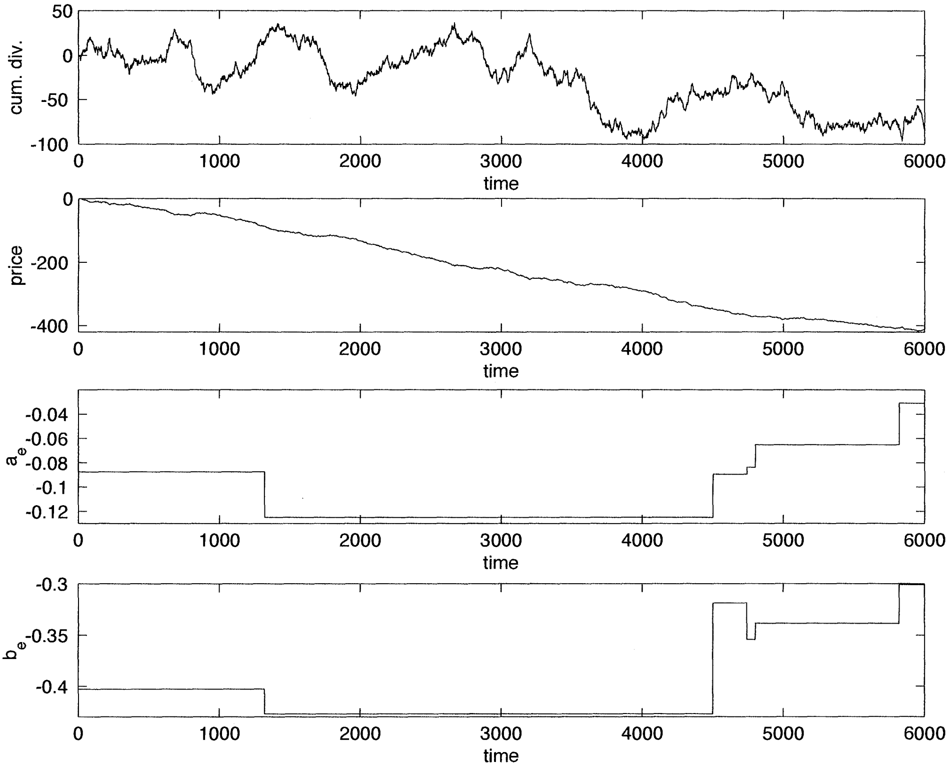

Figure 3. Dividends and prices with long-memory traders. Evolution of cumulative dividend (top plot), prices (second plot), parameters a_e (third plot), and b_e (bottom plot) in one simulation of an economy with long-memory traders (the model in equation (12)). Total time length displayed: (T =) 6,000 periods, or (n =) 100 epochs of (τ =) 60 periods each. The first 200 epochs (preceding the ones shown) are deleted to eliminate the impact of the initial distribution from which the dividend process is started.

Now the evidence from the most recent epoch is used only if the percentage change in the volatility is sufficiently large.

Figure 3 illustrates the nature of the resulting economy. Notice how the range of the prices has increased dramatically. We now discern long-run trends. This is a consequence of the pricing rule: Its parameters are revised only occasionally, as demonstrated in the bottom two panels of the figure. Indeed, the memory of the parameters is quite long.

The trend in the prices clearly implies inefficiencies. We expect to see ample rejections of the null that the average price change is zero. Yet the trend in Figure 3 is not typical: There is an equal probability of obtaining a significantly positive or negative trend. That is, discovering a significantly positive trend in one time series realization does not predict anything about the trend in another.

Figure 4 confirms this. Compared to the boxplot of the standard normal random draws, the distributions of the z-statistics are far flatter, leading the

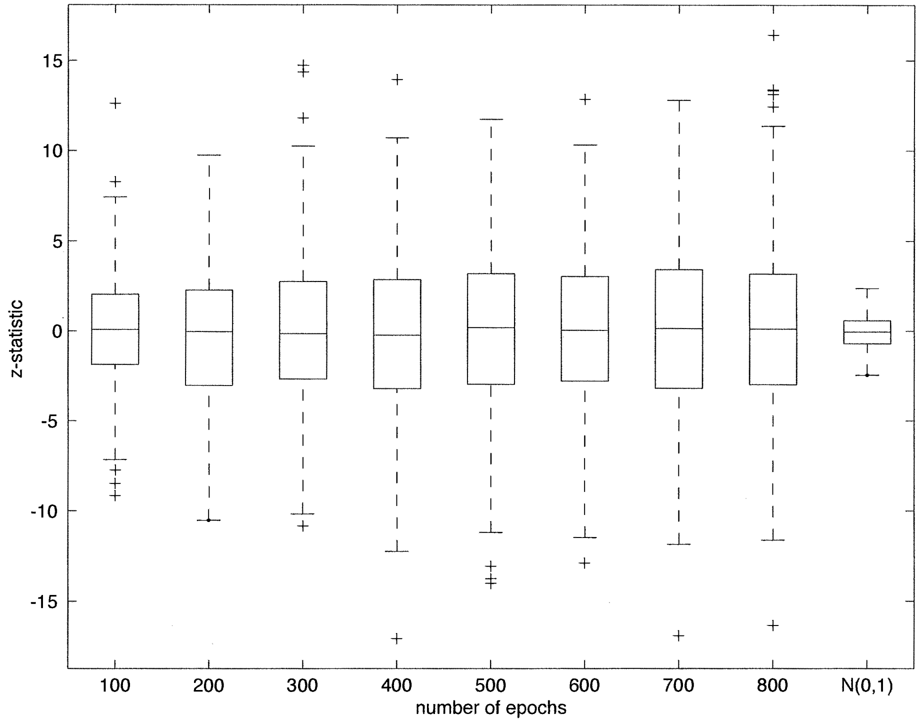

Figure 4. Returns with long-memory traders. Boxplots of distributions of the z-statistic of the return (dividend plus price change) across 400 repetitions of an economy with short-memory traders (the model in equation (12)). Shown are distributions for an increasing number of epochs ($= n$), from 100 to 800. One epoch lasts 60 periods ($= \tau$). In each simulation, the first 200 epochs (preceding the ones from which the z-statistics are computed) are deleted to eliminate the impact of the initial distribution from which the dividend process is started. The z-statistic is adjusted for mild first-order autoregression in the return. For comparison, the last boxplot shows the distribution of 400 random draws from the $N(0,1)$ distribution. These draws are based on the same random number generator on which the other simulations are based.

outsider to often reject (correctly) that this market is efficient. The distributions are more or less symmetric around zero, however, implying about an equal (and large) number of times that returns in this economy would be found to be significantly positive or negative.

This is an inefficient economy, because prices are set by means of a relatively arbitrary pricing rule. Still, the average return this generates is only slightly negative (across simulations, the average return equals -0.00012). Yet, when (correctly) rejecting efficiency, the econometrician will infer about half the time that there are significantly positive returns. When testing efficiency on an independent sample, he/she would likely come to the converse conclusion, that the market generates significantly negative returns.

So, this is a nice example of an economy where an outsider would most often correctly infer that the market is inefficient, because the z-statistic would end up in the usual rejection area. Still, he/she would not be able to

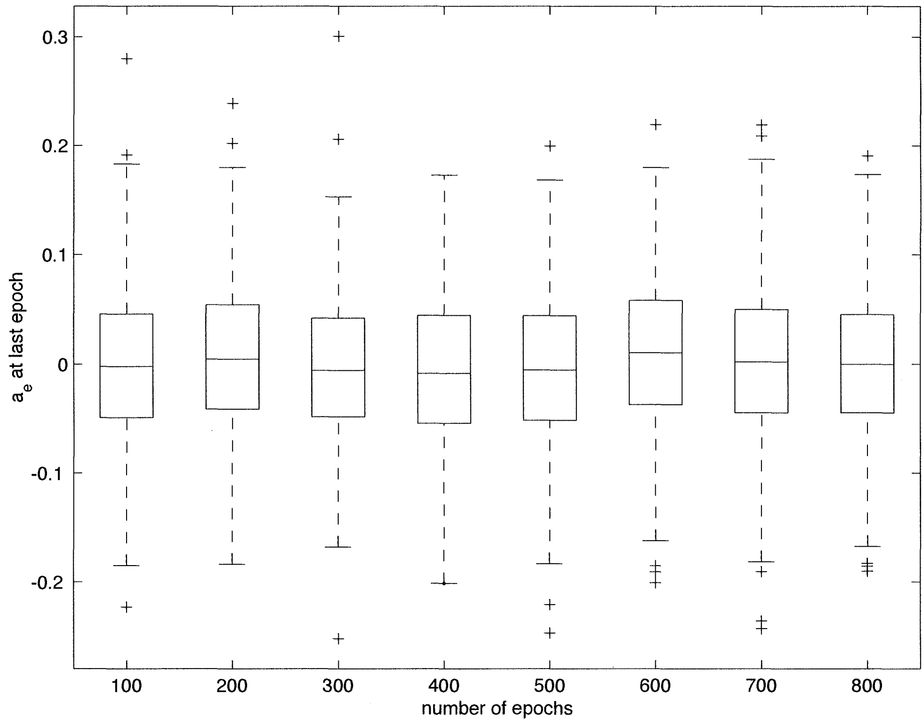

Figure 5. The distribution of a_e with long-memory traders. Boxplots of distributions of the parameter a_e across 400 repetitions of an economy with long-memory traders (the model in equation (12)). The parameter is recorded in the last epoch of histories with an increasing number of epochs (= n), from 100 to 800. One epoch lasts 60 periods (= τ). In each simulation, the first 200 epochs (preceding the ones from which the a_e is computed) are deleted to eliminate the impact of the initial distribution from which the dividend process is started.

assess what the nature of the inefficiency is. About half the time, he/she would observe a significantly positive z-statistic; at other times, the z-statistic would be significantly negative.

To prove that the anomalous behavior of the z-statistic is not a consequence of nonstationarities, Figures 5 and 6 display boxplots of the parameters of the pricing equation after 100, 200, ..., 800 epochs (where we started counting epochs from epoch 201 on). There are no discernible changes in the distributions as we increase the number of epochs, which means that the economy has indeed settled at its stationary distributions even after as few as 100 (ignoring the 200 epochs that were deleted) epochs.

The anomalous behavior of the z-statistic originates in the long memory induced in the return process by long-memory trading rules. This causes the usual estimate of the standard error in (8) to be wrong. Technically, the return process is not strong mixing anymore, so that standard asymptotic theory is invalid. It should be emphasized that violations of mixing (memory) conditions cannot be inferred from a single time series; the reason we

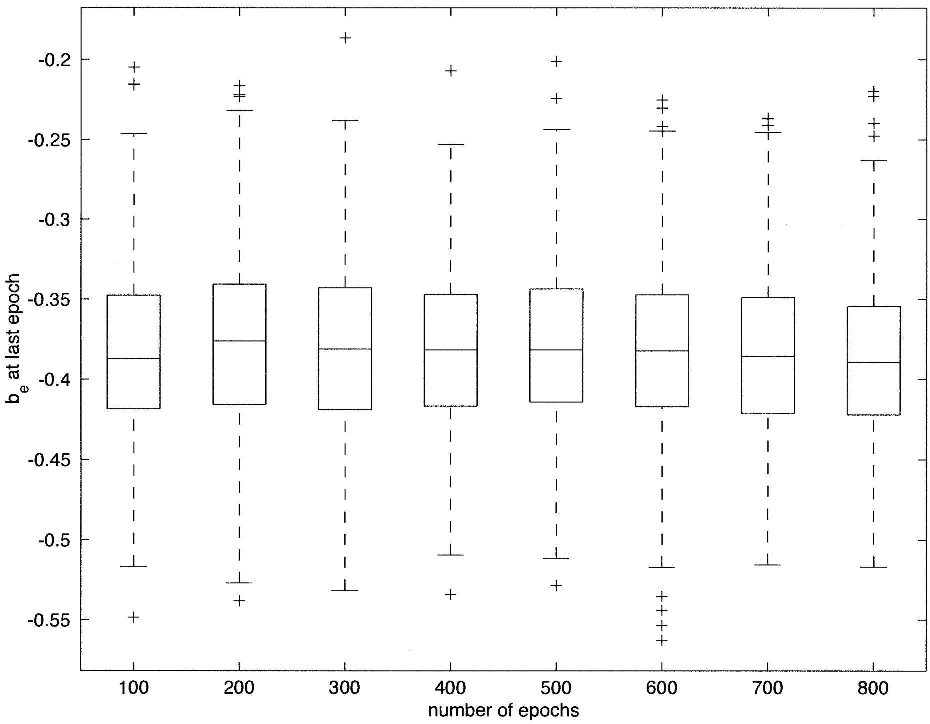

Figure 6. The distribution of b_e with long-memory traders. Boxplots of distributions of the parameter b_e across 400 repetitions of an economy with long-memory traders (the model in equation (12)). The parameter is recorded in the last epoch of histories with an increasing number of epochs (= n), from 100 to 800. One epoch lasts 60 periods (= τ). In each simulation, the first 200 epochs (preceding the ones from which the b_e is computed) are deleted to eliminate the impact of the initial distribution from which the dividend process is started.

can is because we generated a cross section of long time series. So, the problem exemplified by Figure 4 can affect analysis of any single time series.

IV. Real-life Illustrations of Neo-Austrian Ideas

The main point of this paper has been that the Neo-Austrians' rejection of market efficiency can be reconciled with their insistence on order (stationarity) only if standard large-sample inference fails. If investors cannot be confident that their test statistics have the usual asymptotic properties, then rules of thumb could work equally well. This provides the market with some tendency toward efficiency, yet definitely keeps it from ever reaching it entirely.

To illustrate that Neo-Austrian economics is relevant for actual financial markets, we must collect evidence on the failure of standard theory in predicting the large-sample behavior of some widely used statistics.

Consider the price of the three-month U.S. Treasury bill (per one dollar face value) over the period February 1, 1962, to June 25, 1999, as reported by the Federal Reserve Bank of Saint Louis. The efficient market hypothesis

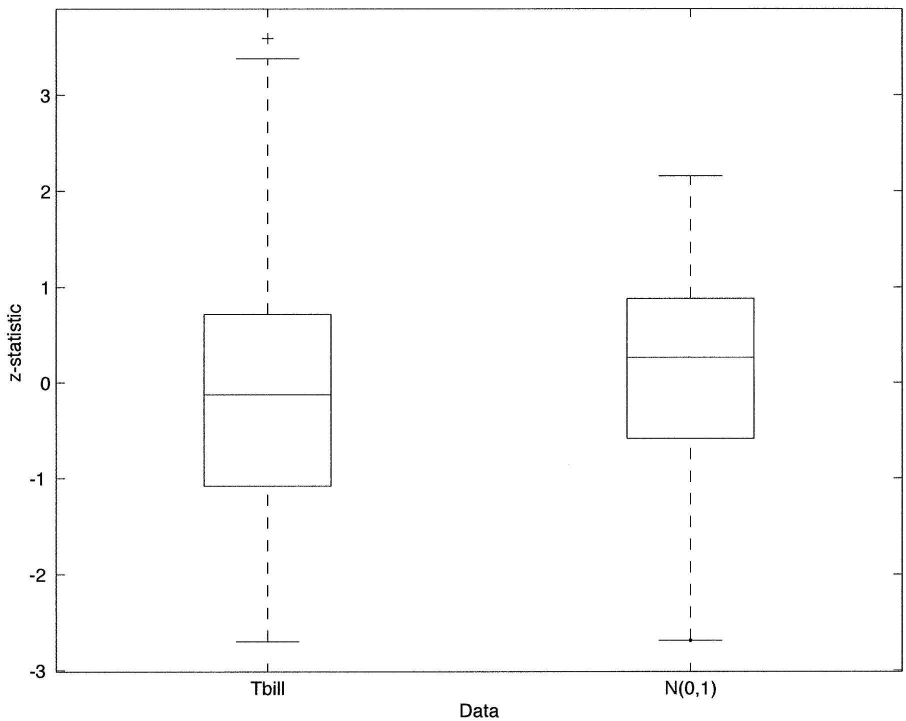

Figure 7. Daily price changes of U.S. Treasury bills. Boxplots of the distribution of the z-statistics of daily changes in the price of the three-month U.S. Treasury bill rate over subperiods of 90 business days from February 1, 1962, to June 25, 1999. For comparison, the right boxplot shows the distribution of an equal number (103) of random draws from the $N(0,1)$ distribution.

predicts that changes in this price over intervals as short as one day should not be predictable. In particular, they should be zero on average. Let us study the outcomes from standard z-tests, computed over subperiods of 90 business days (approximately four calendar months).

Figure 7 provides a boxplot of the z-statistics. For comparison, a boxplot from a simulated $N(0,1)$ random variable is also depicted. Although not as dramatic as in the figures from the previous section, there is a noticeable difference: There is a higher probability for the z-statistic from the T-bill prices to be in the tails than expected under the null of an efficient market.

One can translate the boxplots in Figure 7 into concrete numbers: 14 percent of the z-statistic is below the five percent critical value for a $N(0,1)$ random variable; 11 percent is above the 95 percent critical value. This implies, among other things, that an empiricist would have rejected the null of market efficiency in one-fourth (25 percent) of the subperiods in two-sided tests with a 10 percent p level.

The chance of finding more than 25 rejections out of 103 subperiods if one expects only 10 percent rejections is less than 0.1 percent. Hence, the data clearly reject the null of an efficient market. Yet, there are about an equal number of subperiods where the change in the price of the T bill is significantly negative or positive. That is, we are not in the position to interpret the nature of the inefficiency. Our rejecting market efficiency on the basis of standard statistical inference does not allow us to derive clear prescriptions for how to exploit the inefficiency. Rules of thumb become a necessity.

The fact that there are about an equal number of subperiods where the price change is significantly negative or positive indicates that the Treasury bill market is stationary: History repeats itself. Yet it is inefficient. The coexistence of stationarity and inefficiency is at the heart of Neo-Austrian economic thinking. It can only be explained because we are incapable of interpreting the evidence against efficiency. It also explains why Neo-Austrians insist on agents' using rules of thumb rather than fully fledged optimization policies.

V. Conclusion

The contribution of this paper can best be viewed as one where we attempt to give precise content to a major postulate of Neo-Austrian economic theory when applied to financial markets. The postulate states that the return process will be stable yet continuously feature inefficiencies, keeping the market from reaching its fully efficient equilibrium. Despite the stability (stationarity), rational, risk-averse investors are unable to exploit all inefficiencies because they cannot make reliable inferences. This would be the case if the memory of the return process is sufficiently long for statistics not to display their usual distributional properties needed to construct confidence intervals. We provide an example of an economy where this is the case.

One could object to the use of standard inference techniques in a case where the return process exhibits long memory. Adaptation of the inference procedure, however, requires knowledge of the precise level of the memory. How would agents know this? Memory is a property that they cannot infer from the history of returns. Notice that there is no ambiguity about stationarity; it concerns only memory.

This is precisely why one would conclude, as the Neo-Austrians do, that investors' knowledge, although not characterized by complete ignorance (stationarity will guarantee this), is nevertheless limited. There is not a unique prescription for how a rational investor would cope with the limitations. Some advocate that one continue to use standard probability theory, albeit dropping the additivity axiom. Others would find rules of thumb more sensible in such a situation. In fact, in our economy, this is precisely what traders do. Of course, many would restrict the rules of thumb to include only those that are successful in the long run in a given situation.

One could try to "hedge" against potentially problematic data generating processes by using the econometric procedure with the widest validity. While it is obvious that a "catch-all" procedure does not exist, it is well known that increasing generality may actually hurt. Ghysels (1998), for instance, shows how a popular econometric technique that allows for time-varying betas produces larger asset-pricing errors than a more restrictive methodology that requires constant betas.

The analysis of this paper has implications beyond Neo-Austrian economics. It exemplifies how classical inference can be unreliable when financial markets are inefficient. This should caution empiricists from concluding that they can infer the *nature* of inefficiencies once they have found that they exist. Loosely speaking, if inefficiencies are of the Neo-Austrian kind, not much can be said beyond admitting that inefficiencies exist. Market intervention that is based on potentially erroneous inference is certainly ill advised. Friedrich Hayek's distaste for government intervention in market forces now becomes understandable. Intervention may make things worse, despite the apparent inefficiencies that already exist.

REFERENCES

Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler, 1997, Asset pricing under endogenous expectations in an artificial stock market; in W. Brian Arthur, Steven N. Durlauf, and David A. Lane, eds.: *The Economy as an Evolving Complex System II* (Addison-Wesley, Reading, MA).

Bossaerts, Peter, 1995, The econometrics of learning in financial markets, *Econometric Theory* 11, 151–189.

Ghysels, Eric, 1998, On stable factor structures in the pricing of risk: Do time-varying betas help or hurt? *Journal of Finance* 53, 549–575.

Haugen, Robert, 1995, *The New Finance: The Case against Efficient Markets* (Prentice-Hall, Englewood Cliffs, NJ).

Hayek, Friedrich A., 1937, Economics and knowledge, *Economica* 4, 33–54.

Hayek, Friedrich A., 1945, The use of knowledge in society, *American Economic Review* 35, 519–530.

Hayek, Friedrich A., 1948, The meaning of competition; in Friedrich Hayek, ed.: *Individualism and Economic Order* (The University of Chicago Press, Chicago).

Hayek, Friedrich A., 1978, Competition as a discovery procedure; in Friedrich Hayek, ed.: *New Studies in Philosophy, Politics, Economics and the History of Ideas* (Routledge & Kegan Paul, London/Henley).

Kirzner, Israel M., 1997, Entrepreneurial discovery and the competitive market process, an Austrian approach, *Journal of Economic Literature* 35, 60–85.

Kurz, Mordecai, 1994, On rational beliefs equilibria, *Economic Theory* 4, 877–900.

LeBaron, Blake, 1999, Evolution and time horizons in an agent-based stock market, Working paper, Brandeis University.

Littlechild, Stephen C., 1982, Equilibrium and the market process; in Israel M. Kirzner, ed.: *Method, Process, and Austrian Economics (Essays in Honor of Ludwig von Mises)* (Lexington Books, Lexington/Toronto).

Marcet, Albert, and Juan P. Nicolini, 1997, Recurrent hyperinflations and learning, Working paper, Universitat Pompeu Fabra.

Rizzo, Mario J., 1990, Hayek's four tendencies toward equilibrium, *Cultural Dynamics* 3, 12–31.