



ELSEVIER

Journal of Banking & Finance 24 (2000) 1535–1555

Journal of
BANKING &
FINANCE

www.elsevier.com/locate/econbase

Expectations and learning in Iowa

Oleg Bondarenko^a, Peter Bossaerts^{b,*}

^a *University of Illinois, Chicago, IL 60612, USA*

^b *California Institute of Technology, Division of Humanities and Social Sciences, m/c 228-77,
Pasadena, CA 91125, USA*

Received 27 July 1997; accepted 2 September 1999

Abstract

We study the rationality of learning and the biases in expectations in the Iowa Experimental Markets. Using novel tests developed in (Bossaerts, P., 1996. Martingale restrictions on equilibrium security prices under rational expectations and consistent beliefs. Caltech working paper; Bossaerts, P., 1997. The dynamics of equity prices in fallible markets. Caltech working paper), learning in the Iowa winner-take-all markets is found to be in accordance with the rules of conditional probability (Bayes' law). Hence, participants correctly update their beliefs using the available information. There is evidence, however, that beliefs do not satisfy the restrictions of rational expectations that they reflect the factual distribution of outcomes. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: C90; D84; G14

Keywords: Arrow–Debreu securities; Experimental markets; Rational expectations equilibrium; Bayesian equilibrium; Consistent beliefs; Market efficiency; Learning; Martingales

* Corresponding author. Tel.: +1-626-395-4028; fax: +1-626-405-9841.

E-mail address: pbs@rioja.caltech.edu (P. Bossaerts).

1. Introduction

Over the last several decades, empirical analysis in virtually all of finance as well as a large part of macroeconomics has been based on the assumption of a *Rational Expectations Equilibrium* (REE). There, rational expectations are primarily understood to imply that investors' beliefs are correct.¹ This means that agents are able to weigh potential future outcomes using probabilities that reflect the frequencies with which these *factually* will occur. In other words, their priors coincide with the true probability measure. Rational expectations do not necessarily imply that investors have perfect foresight. However, to the extent that future events are uncertain, investors correctly infer the necessary information from signals (including prices in the market), applying the rules of conditional probability (Bayes' law).

The correctness of beliefs in the finance/macroeconomic definition of rational expectations is not crucial for many theoretical insights, such as the Euler equation restrictions obtained in Lucas (1978), the existence of an equivalent martingale measure shown in Harrison and Kreps (1979),² or the dealer market pricing equations of Kyle (1985). These continue to hold even if we allow for biases in investors' beliefs. But expectations must then be taken with respect to agents' subjective beliefs instead of the "true probability measure".

Instead, the attraction of the rational expectations assumption is exclusively *empirical*. It facilitates estimation and testing of asset pricing models, because agents' beliefs can readily be obtained from the empirical distribution. Parsimonious model testing ensues, as exemplified by Hansen and Singleton (1982). Rational expectations is further motivated by the fear that the possibility of biases in beliefs would require the empiricist to be much more specific about the economy, and, in particular, about learning, with the consequent danger that asset pricing models may be rejected because the evolution of beliefs was misspecified.

Still, the assumption of rational expectations has been criticized as unrealistic. Proponents defend it as the adequate long-run description of an economy by pointing out that rational expectations can be acquired by learning. Of course, this does not provide sufficient ground to ignore the initial biases in beliefs and potentially transient learning in tests of asset pricing models. For

¹ There is a micro-theoretic literature on rational expectations, which does not necessarily imply correctness of beliefs. See Radner (1967), Green (1973). Here, we refer to the concept of a *Dynamic Rational Expectations Equilibrium*, as originally defined in Lucas (1972, 1978), where beliefs are explicitly assumed to be correct. This concept seems to have gotten its inspiration from the vaguer idea of an *informationally efficient market* in finance (Fama, 1970, 1991). See also Muth (1961).

² We do rule out asymmetric information here, though. Asymmetric information would generally invalidate the existence of an equivalent martingale measure even under correct beliefs.

instance, Bossaerts (1995) proves analytically how standard rational expectations test are affected *even asymptotically* by initial biases and transient learning. That article also illustrates the fragility of the asymptotic distributional properties of test statistics as one alters the specification of beliefs and beliefs updating.

The necessity of specifying the learning environment when one deviates from REE was, however, recently challenged in Bossaerts (1996, 1997). That paper derives a set of martingale restrictions on securities prices that neither depend on the priors nor on the likelihood function. One of the securities considered in Bossaerts (1996, 1997) was of the simple Arrow–Debreu type: it pays one dollar in one state of the world, and zero in all others. Markets were assumed to be populated with risk-neutral agents endowed with unknown priors. These agents, however, used Bayes' law to infer the actual state from sequences of signals. To reflect the close relationship with the game-theoretic literature (in particular, Harsanyi, 1967), let us refer to the ensuing equilibrium as the *Bayesian Equilibrium* (BE).

The purpose of the present paper is to illustrate how the theoretical results in Bossaerts (1996, 1997) can be implemented to test for BE against the more restrictive traditional REE. In the absence of an explicit specification of priors and likelihood functions, we only test the central part of a Bayesian Equilibrium, namely investor's usage of Bayes' law to update beliefs. Since Bayesian updating is rational in the sense that it implies that the agent cannot forecast where her beliefs will go next, we effectively study the *rationality of updating*. This contrasts with tests of REE, where one requires not only rational updating, but also correct priors.

We investigate the pricing of “winner take all” contracts traded in the computerized experimental markets developed at the University of Iowa. Winner-take-all contracts are simple examples of Arrow–Debreu securities: they pay \$1 in one state of the world; \$0 in all others. We study the contracts which derive their payoff from price changes of stocks in the computer industry and the S&P500 index.

Our focusing on experimental markets should be attributed to the absence of real markets with Arrow–Debreu securities. Still, one could imply synthetic Arrow–Debreu securities prices from the prices of widely traded put and call option contracts. Our restrictions could be tested on these synthetic prices. See Bondarenko (1997) for an application to S&P500 index options. An alternative would be to directly study put and call prices: Bossaerts (1996, 1997) derived restrictions on their prices as well. A disadvantage is that these restrictions are more involved and less intuitive.

We find that prices in the Iowa markets reflect rational learning. Evidence against rational expectations surfaces, however, indicating that investors may have started with biased priors. Among other things, there is a pronounced negative correlation between the price level and the subsequent price change.

This finding is not unlike the evidence from stock markets (e.g., Keim and Stambaugh, 1986).

It deserves emphasis that the latter does not necessarily imply that there is money to be made in the future. Ex ante, and absent insider information, the priors of participants in the Iowa markets are as good as any other. To find that they were biased in the past, i.e., ex post, will not necessarily help an outsider to profit in the future. The negative correlation between the price level and the subsequent price change does indicate that market priors were biased. But we lack the necessary data to estimate the nature of this bias with sufficient precision.

The remainder of the paper is organized as follows. The next section explains the theory with the help of a simple example. (The general proof is in the Appendix.) Section 3 describes the Iowa experimental markets that are the focus of the study. Section 4 presents the results of the tests of the rationality of learning. Section 5 investigates rational expectations. Section 6 concludes the paper.

2. The theory by example

Consider a world in which a state variable, θ , can take on two values: $\bar{\theta}$ or $\underline{\theta}$. Securities are traded and markets clear at two discrete points in time, indexed $t = 1, 2$. Consider an Arrow–Debreu security which pays off at $t = 3$, \$1 if $\theta = \bar{\theta}$ and zero otherwise. Risk-neutral investors do not know θ before the end of the second period, i.e., before time $t = 3$. At time $t = 2$, however, they do receive a signal s , which they use to update their beliefs about the value of θ . The equilibrium price at time t is denoted p_t .

For the moment, we are going to assume that agents have rational expectations. In other words, they know the true probability measure P which determines the relative frequency with which signals and states are drawn. They will use this probability measure to infer from the signals whether $\theta = \bar{\theta}$ or not.

The first price, p_1 , will be set to equal the unconditional probability (prior) that $\theta = \bar{\theta}$. Mathematically

$$p_1 = P\{\theta = \bar{\theta}\}.$$

The signal at $t = 2$, s , is determined as follows. If $\theta = \bar{\theta}$ then

$$s = \begin{cases} 1 & \text{with } P = 1/2, \\ 0 & \text{with } P = 1/2. \end{cases} \quad (1)$$

In contrast, if $\theta = \underline{\theta}$,

$$s = \begin{cases} 1 & \text{with } P = 1/3, \\ 0 & \text{with } P = 2/3. \end{cases} \quad (2)$$

Setting $P\{\theta = \bar{\theta}\} = \frac{1}{4}$, we get

$$p_1 = \frac{1}{4}. \tag{3}$$

Letting $l(s|\theta)$ denote the likelihood of s given θ , Eqs. (1) and (2) imply that

$$l(1|\bar{\theta}) = \frac{1}{2},$$

$$l(1|\underline{\theta}) = \frac{1}{3},$$

$$l(0|\bar{\theta}) = \frac{1}{2},$$

$$l(0|\underline{\theta}) = \frac{2}{3}.$$

Because p_1 equals the prior that $\theta = \bar{\theta}$, we conclude: if $s = 1$,

$$p_2 = \frac{l(1|\bar{\theta})p_1}{l(1|\bar{\theta})p_1 + l(1|\underline{\theta})(1 - p_1)} = \frac{\frac{1}{2} \frac{1}{4}}{\frac{1}{2} \frac{1}{4} + \frac{1}{3} \frac{3}{4}} = \frac{1}{3}; \tag{4}$$

if $s = 0$,

$$p_2 = \frac{l(0|\bar{\theta})p_1}{l(0|\bar{\theta})p_1 + l(0|\underline{\theta})(1 - p_1)} = \frac{\frac{1}{2} \frac{1}{4}}{\frac{1}{2} \frac{1}{4} + \frac{2}{3} \frac{3}{4}} = \frac{1}{5}. \tag{5}$$

We can use this to derive several results.

Result 1. Consider the expected time-2 price:

$$E[p_2] = \frac{1}{3}P\{s = 1\} + \frac{1}{5}P\{s = 0\} = \frac{1}{3} \frac{3}{8} + \frac{1}{5} \frac{5}{8} = \frac{1}{4} = p_1.$$

Consequently,

$$E[r_2] = 0, \tag{6}$$

where

$$r_2 = \frac{p_2 - p_1}{p_1},$$

the traditional return measure. Eq. (6) states that the return is unpredictable. More generally, the return sequence will form a martingale difference sequence. This, of course, is an old result, which goes back to Samuelson (1965).

Result 2. We now take the point of view of an econometrician contemplating a single price sequence. In this sequence, θ remains fixed. We assume that the econometrician observes the price history after time $t = 3$, so that she knows the value of θ . In that case, the expectation in (6) is of little relevance. Instead,

we need to study $E[r_2|\theta]$, i.e., the expectation conditional on the value that θ takes on in the particular sequence at hand. Let us consider the case where the security matures in the money, i.e., $\theta = \bar{\theta}$. Compute

$$E[p_2|\bar{\theta}] = \frac{1}{3}P\{s = 1|\bar{\theta}\} + \frac{1}{5}P\{s = 0|\bar{\theta}\} = \frac{1}{3} \frac{1}{2} + \frac{1}{5} \frac{1}{2} = \frac{4}{15} > p_1.$$

Hence,

$$E[r_2|\bar{\theta}] > 0. \tag{7}$$

In words, there is an upward drift in prices. Of course, this is not surprising, because the econometrician picked a sequence of prices for a security that she knows matured in the money. If she had picked a security that had matured out of the money, she would have observed a negative drift.

Result 3. Is there a simple way to correct for the upward drift in (7)? Consider the conditional expectation of the *inverse* of the price at time 2:

$$E\left[\frac{1}{p_2}|\bar{\theta}\right] = 3P\{s = 1|\bar{\theta}\} + 5P\{s = 0|\bar{\theta}\} = 3 \frac{1}{2} + 5 \frac{1}{2} = 4 = \frac{1}{p_1}.$$

Hence, defining

$$x_2 = \frac{p_2 - p_1}{p_2}, \tag{8}$$

we conclude

$$E[x_2|\bar{\theta}] = 0. \tag{9}$$

In words, modifying the return by taking as basis not the past price but the *future price* makes it unpredictable!

The martingale difference result in (9) is very general. It does not depend on the likelihood function; it solely follows from investors' using Bayes' law to update their beliefs. Over time, signals can be dependent. See Bossaerts (1996, 1997) for details and a general proof. Intuitively, the result obtains from taking the future price as basis to compute a return, which cancels the drift in (7). The latter equation effectively implies that the new basis for the return measure, the future price, is higher on average than the old basis, the past price.

We are now going to relax the assumption of rational expectations. Investors no longer are required to use the "true" probability measure P . Instead, they use another measure, P^* , which is related to P , as follows. Factor P into a measure μ over the possible values of the state variable θ and a *conditional* measure P_θ which determines the likelihood of signal outcomes conditional on θ (the likelihood functions in the above example are derived from P_θ). Likewise, factor P^* into a prior λ_0 over the state variable, and a conditional measure P_θ over the signal sequence. Notice that we constrain the signal likelihoods to be

the same under both P and P^* . The priors, however, may differ: Nature uses μ to draw θ ; agents posit λ_0 .

Consider the specific case where $P^*\{\theta = \bar{\theta}\} = \frac{1}{2}$. Nature continues to draw θ according to the old law, which assigns only probability $\frac{1}{4}$ to the event $\{\theta = \bar{\theta}\}$. Apparently, investors are too optimistic. From the specification of beliefs, we conclude

$$p_1 = \frac{1}{2}. \tag{10}$$

Also, if $s = 1$, then computations that are analogous to those for the rational expectations case reveal

$$p_2 = \frac{l(1|\bar{\theta})p_1}{l(1|\bar{\theta})p_1 + l(1|\underline{\theta})(1 - p_1)} = \frac{\frac{1}{2} \frac{1}{2}}{\frac{1}{2} \frac{1}{2} + \frac{1}{3} \frac{1}{2}} = \frac{3}{5}; \tag{11}$$

if $s = 0$,

$$p_2 = \frac{l(0|\bar{\theta})p_1}{l(0|\bar{\theta})p_1 + l(0|\underline{\theta})(1 - p_1)} = \frac{\frac{1}{2} \frac{1}{2}}{\frac{1}{2} \frac{1}{2} + \frac{2}{3} \frac{1}{2}} = \frac{3}{7}. \tag{12}$$

Let us now reconsider the conditional expectations in Results 1–3. We continue to compute them under the probability measure P , which is the one that governs the data as observed by the econometrician. In other words, P is the measure that determines the behavior of statistics that the econometrician may want to compute.

Result 4. Taking expectations,

$$E[p_2] = \frac{3}{5}P\{s = 1\} + \frac{3}{7}P\{s = 0\} = \frac{3}{5} \frac{3}{8} + \frac{3}{7} \frac{5}{8} = \frac{69}{140} < p_1.$$

Consequently,

$$E[r_2] < 0. \tag{13}$$

Unlike Result 1, however, Eq. (13) is *not robust*. By changing the prior λ_0 , we can easily overturn it. In our numerical example, investors are too optimistic about the likelihood of $\bar{\theta}$; in other words, they overestimate the probability that the Arrow–Debreu security will mature in-the-money. Had we taken them to be more pessimistic, by assuming, for instance, that $P^*\{\theta = \bar{\theta}\} = \frac{1}{8}$, the inequality would have been reversed.

Result 5. Now augment the information with the knowledge that $\theta = \bar{\theta}$:

$$E[p_2|\bar{\theta}] = \frac{3}{5}P\{s = 1|\bar{\theta}\} + \frac{3}{7}P\{s = 0|\bar{\theta}\} = \frac{3}{5} \frac{1}{2} + \frac{3}{7} \frac{1}{2} = \frac{72}{140} > p_1.$$

Hence,

$$E[r_2|\bar{\theta}] > 0. \tag{14}$$

This result may not seem to be robust either, in the sense that it can be overturned by changing the prior λ_0 . This is not true; we come back to this point shortly.

Result 6. Recompute the conditional expectation of the *inverse* price

$$E\left[\frac{1}{p_2}|\bar{\theta}\right] = \frac{5}{3}P\{s = 1|\bar{\theta}\} + \frac{7}{3}P\{s = 0|\bar{\theta}\} = \frac{5}{3}\frac{1}{2} + \frac{7}{3}\frac{1}{2} = 2 = \frac{1}{p_1}.$$

Again modifying the return by taking the future price as basis

$$x_2 = \frac{p_2 - p_1}{p_2},$$

we obtain

$$E[x_2|\bar{\theta}] = 0. \tag{15}$$

The martingale property of the modified return has been unaffected by our relaxing the assumption of rational expectations!

Appendix A proves that the last result *is* robust. It does not depend on agents’ priors. It is merely a consequence of investors’ using Bayes’ law to update their beliefs. This is remarkable in view of the many instances in which the conclusion from Bayesian analysis are sensitive to the specification of the prior.

Since the modified return of “winning securities” is a martingale difference sequence, the traditional return *must have a positive expectation* when conditioning on the outcome as well. This follows readily from Jensen’s inequality (which is strict because the signal is nontrivial):

$$\begin{aligned} E[r_2|\bar{\theta}] &= E\left[\frac{1}{1-x_2} - 1|\bar{\theta}\right] \\ &> \frac{1}{1-E[x_2|\bar{\theta}]} - 1 = 0. \end{aligned}$$

Hence, Result 5 *is* robust. It holds regardless of the prior.

With respect to Arrow–Debreu securities, further developments can also be found in Bondarenko (1997). For more general securities, the reader is referred to Bossaerts (1996, 1997).

The above results have many implications. In the present paper, we focus on the analysis of a cross-section of independent price sequences. In conjunction with such a dataset, the above results can be used to distinguish between the traditional REE and the more general BE. Under the latter, priors can be arbitrary. Under REE, Result 1 can be verified in a cross-section of random

histories of prices of Arrow–Debreu securities. If the property in (6) is rejected, rational expectations must be rejected as well.³ That would not reject rationality, however. Agents could have started with biased beliefs. For rationality, they need only to apply correctly the rules of conditional probability. That is the core of the more general class of Bayesian Equilibria. In those, a *martingale result holds only for the modified return of securities that matured in the money*. This is the content of Results 4–6. A cross-section of price histories of “winning” Arrow–Debreu securities would allow us to test the proposition that investors’ learning is rational. This will be the subject of the subsequent empirical investigation.

One wonders what type of irrationality would lead to violations of Results 4–6. It can be shown (see Bossaerts and Hillion, 1997) that *over-reaction to news* leads to negative average modified returns on winning securities, and *under-reaction* (sluggishness) causes positive average modified returns. Over-reaction and under-reaction is to be defined relative to Bayes’ law. Bossaerts and Hillion (1997) also report simulations that illustrate that test statistics based on modified returns have excellent power properties. Of course, the tests will not have any power against updating rules that lead to price reactions that are similar to those generated by Bayes’ law (many frequency-analytic estimation procedures have that property).

Finally, notice that the key implication of BE, Result 2, can be tested on *biased* samples. Only price series of “winning” Arrow–Debreu securities are needed. If histories of losers are unavailable to the researcher, she can still test for BE. This is a real opportunity, because many historical financial databases do come with such a “survivorship bias”. The one we are about to present here does not suffer from survivorship bias. This allows us to demonstrate in detail the workings of our theory and tests.

3. The Iowa experimental market

The College of Business Administration at the University of Iowa runs a number of markets in winner-take-all contracts based on the returns or prices of US stock exchange listed securities and indices. Each month, a new set of contracts is offered. Contract liquidation values are determined by rates of return or changes in closing prices of the underlying stock measured from the

³ Of course, there is still the possibility that agents are risk averse, unlike what we have been assuming. The stakes in the markets we investigate in the subsequent sections, however, are small. Risk premia, which are of second order anyway, are unlikely to play an important role. Incidentally, discounting could not play a role either: the one-dollar riskfree security in the Iowa markets continuously sells for one dollar, implying that it bears a zero interest rate.

third Friday of one month to the third Friday of the next month. Trade starts on Monday following the third Friday of the month. Trading halts when the liquidation value is determined. Rates of return on the underlying stock securities are adjusted for dividend payments. Rates of return on indices are not dividend-adjusted.

Iowa manages another set of markets, where the payoffs on the traded securities depend on political events, such as the outcome of presidential or parliamentary elections. These markets pre-date the ones that we are studying here. For an in-depth analysis of the political markets, see Forsythe et al. (1992).

We investigate two sets of markets: the computer industry returns markets and the Microsoft price level markets. Four winner-take-all contracts are traded in the former. They pay one dollar depending on which underlying security had the highest return: Apple, IBM, Microsoft or the S&P500, respectively. Contracts written on losers expire worthless. Hence, these contracts are Arrow–Debreu securities written on a state space determined by the returns of the aforementioned stocks and the S&P500 index. In the Microsoft markets, two contracts are traded. One, the “High” contract, pays a dollar when Microsoft’s next month closing price is above a predetermined cut-off level. The other one, the “Low”, pays one dollar in the complementary state. The cut-off level is determined by the exercise price of the closest-at-the-money, exchange-traded option (CBOE) on Microsoft.

The trading mechanism of the Iowa Electronic Market is the computerized double auction. Participation extends beyond the University of Iowa (the market is internet-based), but is restricted to investors with academic affiliations. In addition to trading individual contracts, a “market” portfolio can be purchased and sold. This is a unit portfolio which bundles all contracts in a set of markets, and, hence, guarantees a unit payout. It can be purchased from or sold to the system for a unit price, or bought/sold in the market (the latter requires, however, that there be quotes in all contracts).

We limit our attention to daily closing prices, which are defined to be the last transaction price before midnight, or, if no transactions took place, the previous closing price. Notice that closing prices, and, hence, returns, on complementary contracts are nonsynchronous. This implies, among other things, that the closing prices do not necessarily add up to one, even if a bundle of complementary contracts guarantees a unit payoff, and that returns are not colinear. Let us look at the data to gauge the importance of nonsynchrony.

Table 1 displays some descriptive price and volume statistics. Fig. 1 plots the price paths for the two Microsoft securities over four sample months. Table 1 lists the average and standard deviation of the absolute deviation of the sum of the prices of complementary securities from one. The numbers are sizeable, indicating that nonsynchrony is a serious problem. It also means that the

Table 1
Descriptive statistics^a

		All markets	Microsoft markets	Computer industries markets
<i>N</i>		2380	795	1585
Daily return	Average	0.083	0.081	0.084
	St. dev.	1.563	1.169	1.728
	Skewness	27.7	16.8	28.2
	Kurtosis	980	325	953
Average daily volume (Units)			80.0	266.3
Measure of nonsynchrony	Average		0.006	0.019
	St. Dev.		0.049	0.062

^aThe information is based on daily closing prices from the Iowa Experimental Markets, 8/95 till 12/96 (16 months). The closing price is the price of the last transaction or, if no transaction took place during the day, the previous closing price. Sometimes, contracts continued to be traded after their payoff became known; the corresponding observations were deleted. Also, the following were discarded because they clearly represented aberrations: All computer markets closing prices on 8/28/95 and 8/29/95; Microsoft Low closing price on 5/20/96; Computer industry markets closing prices on 6/20/96. *N* denotes sample size. Nonsynchrony is measured by the average absolute deviation of the sum of the closing prices of complementary contracts from unity.

returns on complementary securities are not necessarily colinear, a fact that we will exploit in our tests. The lack of colinearity is also apparent from Fig. 1: if synchronous, the price paths of the two Microsoft securities should be perfectly symmetric with respect to a line through \$0.50. They are not.

Table 1 also documents substantial positive skewness and extreme kurtosis in the distribution of returns. Readers familiar with derivatives markets (such as the winner-take-all markets we are studying here) will not be surprised. The kurtosis is also apparent in Fig. 1. The history for the 10/95 Microsoft High contract (which pays \$1 if Microsoft's price increased from September 1995) includes a day (10/17/95) when the price jumped from \$0.024 to \$0.455, a return of 1796%! The skewness is reflected in the surprisingly high average daily return. See Table 1. The averages are less than two standard errors from zero, however, except in the pooled sample.

4. Testing rationality of learning

Results 3 and 6 of Section 2 demonstrate that one can use the modified return (where the end-of-period price is used to deflate payoffs) of winning securities to test for the rationality of learning. Those results solely depended on agents' correctly applying the rules of conditional probability (Bayes' law), and not on the initial beliefs, which could be arbitrarily biased.

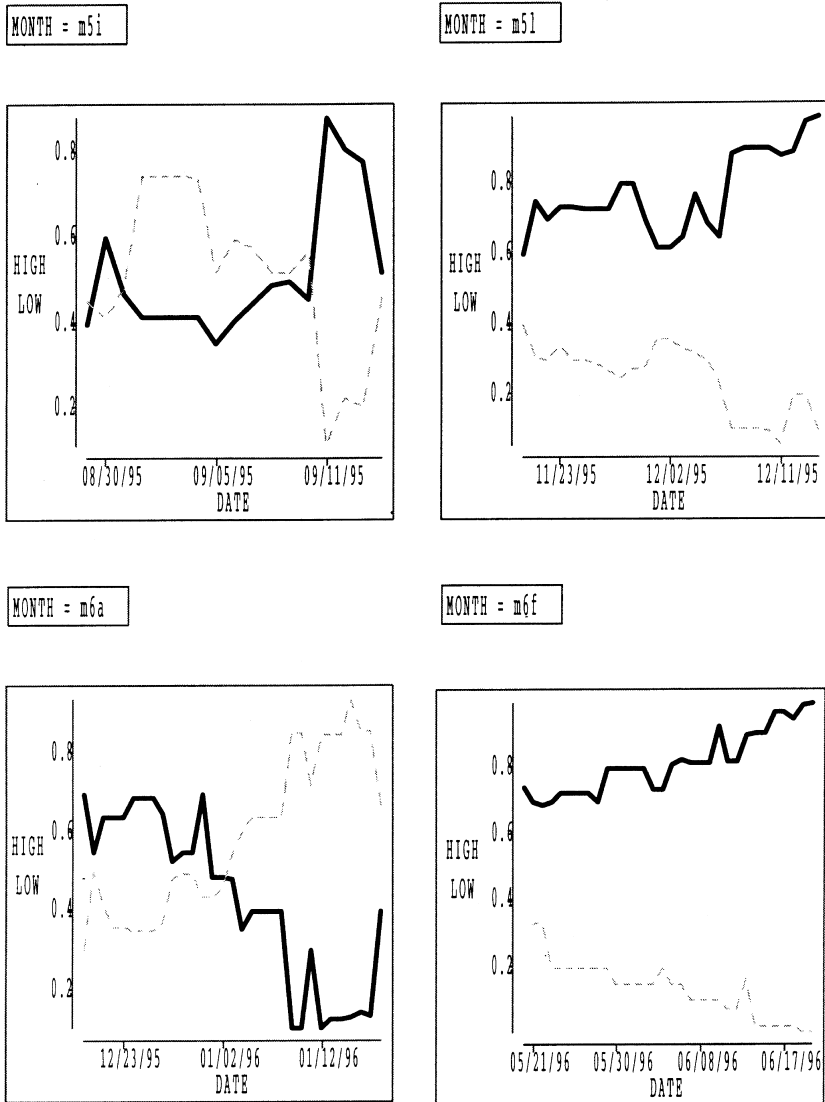


Fig. 1. Time series plots of daily closing prices for contracts microsoft high (bold line) and microsoft low (dashed line). One-month histories for four expiration months are shown: 9/95 (month m5i), 12/95 (month m5l), 1/96 (month m6a) and 6/96 (month m6f).

Let us look at the returns and modified returns *on winning securities* in the Iowa market data described in the previous section. Because of Results 2 and 5 in Section 2, we expect the traditional return to be positive on average. This contrasts with the modified return, which should be zero on average.

We computed average daily returns and modified returns for each month-long price history and then tested whether the average (of this average) was zero across the monthly histories. Since the dataset covers sixteen months, we will have a cross-section of sixteen averages for each individual set of markets (the Microsoft markets and the computer industry markets separately) and thirty-two averages when we pool both sets of markets. The outcomes in the cross-section can safely be assumed to be independent.

We can also look at the securities with payoffs complementary to those of the losing securities. Of course, the complementary securities are winners as well. Complementary securities can easily be created by purchasing the market portfolio from the system for \$1 and selling one traded security. We take the price of the complementary security to be the value of this portfolio. In the Microsoft case, each security is the complement of the other, so that results are expected to be similar to those obtained from actually traded winning securities. Discrepancies will arise, however, because securities trade nonsynchronously. In the computer industry markets, the complementary security to, say, Apple, is the one that pays \$1 if any other security than Apple wins.

In the computer industry markets, there will be correlation across monthly histories of complementary securities that traded simultaneously, i.e., during the same month. The results that we are about to report do not correct for this. When we attempted to address the problem (mainly by running Seemingly Unrelated Regressions), the results were not altered qualitatively.

Table 2 displays the average of the time series daily return average, as well as its standard error. In all cases is this average more than two standard errors above zero, corroborating Results 2 and 5. Table 2 also lists results for the average of the time series daily *modified* return average. It is always within two standard errors from zero, in accordance with Results 3 and 6.⁴

In Table 1, we documented extreme skewness and kurtosis in the distribution of returns. One wonders to what extent these affect the inference in Table 2. Therefore, Table 2 also displays the skewness and excess kurtosis of the month-long averages. The values are far more acceptable, rendering credibility to our results. This is certainly so for the modified returns.

We can conclude provisionally that learning in the Iowa markets is rational: there are no biases in the modified returns of winning securities.

In fact, Results 3 and 6 are stronger than this simple test of the average modified return may indicate. For learning to be rational, the average modified

⁴ We only report standard errors and whether the statistic is 1.8 or 2 standard errors from zero. In order to avoid unnecessary controversy, we refrain from attaching a specific p -level to the estimates. We allow the reader to use his or her favorite distributional theory to determine what the corresponding p -level is.

Table 2
 Analysis of payoffs on winning contracts ^a

		All markets	Microsoft markets	Computer industry markets
<i>Winners</i>				
	<i>N</i>	32	16	16
\bar{r}	Average	0.115** (0.030)	0.097** (0.037)	0.173** (0.047)
	Skewness	2.4	3.5	2.0
	Kurtosis	9.9	16.2	9.4
\bar{x}	Average	-0.033 (0.022)	-0.032 (0.022)	-0.035 (0.038)
	Skewness	-1.7	-1.6	-1.6
	Kurtosis	5.8	4.5	5.1
<i>Complements to losers</i>				
	<i>N</i>	64	16	48
\bar{r}	Average	0.046** (0.008)	0.086** (0.026)	0.032** (0.006)
	Skewness	4.0	3.2	2.2
	Kurtosis	25.2	14.7	8.8
\bar{x}	Average	-0.016 (0.010)	-0.028 (0.020)	-0.011 (0.011)
	Skewness	-3.9	-1.4	-4.9
	Kurtosis	21.8	3.7	30.0

^a *N* denotes sample size (number of monthly histories). \bar{r} denotes the average across all month-long histories of the time series average daily traditional return. \bar{x} denotes the average across all month-long histories of the time series average daily modified return, where the payoff is divided by the end-of-period closing price. Standard error in parentheses. * and ** indicate that the statistic is more than 1.8 and 2 standard errors from zero, respectively. The closing prices of complements to losers is obtained by subtracting the closing price of a loser from the price at which the market portfolio can be bought from the system (i.e., \$1). The price history of the complements to the losers is not identical to that of the winners, because of nonsynchrony.

return of winning securities ought to be unpredictable from *any* past information. We can test this in two ways.

First, we can correlate the modified return with past information, by computing the average of the modified return times, say, the lagged price, or the lagged price squared. Second, we can simply regress the modified return onto lagged information, e.g., the lagged price.

Let us discuss results from the first strategy. We compute monthly averages of daily observations of

$$x_{t+1}p_t$$

and

$$x_{t+1}(p_t)^2,$$

and test whether they are zero across the months. Of course, we continue to limit our attention to winners. For comparison, we also report results for monthly averages based on the traditional return, i.e., from

$$r_{t+1}p_t$$

and

$$r_{t+1}(p_t)^2.$$

Table 3 reports the results. In no case is the monthly average of $x_{t+1}p_t$ or $x_{t+1}(p_t)^2$ significantly different from zero, supporting Bayesian Equilibrium. This contrasts with the averages of $r_{t+1}p_t$ and $r_{t+1}(p_t)^2$, which are all significantly positive. Taking the results of Tables 2 and 3 together, the modified return of winning contracts is not predictable from past information embedded in the price level.

The second strategy, whereby we project the (modified) daily return onto the previous day's closing price using OLS, deserves more discussion. Many researchers, starting with Keim and Stambaugh (1986), have found that stock returns are negatively related to the lagged price levels. For winners, however, the projection of the *modified* return onto lagged prices should produce insignificant results, at least if prices are set in a Bayesian Equilibrium. We verify this for the Iowa markets.

When plainly projecting the modified return on the lagged price, however, the error exhibits substantial heteroscedasticity: its variance is far higher for low price levels. This obviously affects the inference. It turns out that there is a simple way to correct for the heteroscedasticity in the error term, namely, by regressing the modified return times the lagged price ($x_{t+1}p_t$) onto the lagged price (p_t).

Table 4 reports results of this alternative projection. For comparison purposes, we also display results for OLS projections of the traditional daily return multiplied by the lagged price (i.e., the price change) onto the lagged price. The contrast between the highly significant coefficients for the traditional return against the small and insignificant ones for the modified return is striking. As before, Bayesian Equilibrium (or rational learning) is supported.⁵

⁵ Notice that the sample sizes in Table 4 are far higher than those reported in Tables 2 and 3. However, the latter display averages of *mean* daily returns *across months*, whereas the former displays statistics based on all daily observations, pooled from the monthly time series. This difference is without consequence. For instance, when we computed averages in Tables 2 and 3 by pooling the monthly histories of daily data, as in Table 4, the results did not change significantly (obviously, the sample means did not change at all).

Table 3
 Analysis of payoffs on winning contracts: correlation with lagged information ^a

		All markets	Microsoft markets	Computer industry markets
<i>Lagged price</i>				
	N	32	16	16
\overline{rp}	Average	0.086** (0.004)	0.097** (0.004)	0.035** (0.007)
	Skewness	1.1	-0.4	0.9
	Kurtosis	6.2	2.4	5.4
	Average	-0.011 (0.008)	-0.019 (0.013)	-0.003 (0.009)
\overline{xp}	Skewness	-1.8	-2.3	-0.2
	Kurtosis	8.3	8.4	3.4
	<i>Lagged price squared</i>			
	N	32	16	16
$\overline{rp^2}$	Average	0.009** (0.002)	0.006** (0.003)	0.013** (0.004)
	Skewness	0.6	-0.8	1.5
	Kurtosis	5.3	2.4	5.9
	Average	-0.005 (0.005)	-0.011 (0.008)	0.000 (0.005)
$\overline{xp^2}$	Skewness	-2.4	-2.5	-0.0
	Kurtosis	12.0	10.3	3.6

^a N denotes sample size (number of monthly histories). \overline{rp} denotes the average across all month-long histories of the time series average daily traditional return multiplied by the lagged price level. \overline{xp} denotes the average across all month-long histories of the time series average daily modified return multiplied by the lagged price level. $\overline{rp^2}$ denotes the average across all month-long histories of the time series average daily traditional return multiplied by the square of the lagged price level. $\overline{xp^2}$ denotes the average across all month-long histories of the time series average daily modified return multiplied by the square of the lagged price level. Standard error in parentheses. * and ** indicate that the statistic is more than 1.8 and 2 standard errors from zero, respectively.

For space considerations, we do not report results for the securities that are complementary to the losers. They do not lead to different inference. The interested reader can obtain the results for these securities from the authors.

5. Testing for rational expectations

As is well-known, REE can readily be tested by verifying (in-sample) predictability of returns using past information. We must not, however, condition on future information, such as, e.g., the eventual payoff on a security. In the previous section, we did so, because the tests were inspired by Results 3 and 6.

Table 4
 Analysis of daily payoffs on winning contracts: OLS projections of returns times lagged prices onto lagged prices ^a

	All markets		Microsoft markets		Computer industry markets	
	Return	Modified return	Return	Modified return	Return	Modified return
<i>N</i>	601	601	299	299	301	301
Intercept	0.044** (0.010)	-0.014 (0.014)	0.056** (0.018)	-0.032 (0.030)	0.034** (0.013)	-0.011 (0.014)
Slope	-0.044** (0.017)	0.003 (0.024)	-0.067** (0.026)	0.021 (0.045)	-0.015 (0.028)	0.011 (0.031)
<i>R</i> ²	0.011	0.000	0.021	0.001	0.001	0.000

^aTo compute the modified return, the payoff is divided by the end-of-period price instead of the beginning-of-period price. *N* denotes sample size (number of daily returns). * and ** indicate that the statistic is more than 1.8 and 2 standard errors from zero, respectively. The standard return times the lagged price is identical to the price change (which is analyzed further in Table 5). Returns are multiplied by the lagged price in order to mitigate heteroscedasticity.

In contrast, tests of rational expectations are based on Result 1, where one conditions out (i.e., averages across) eventual outcomes.

We initially tested REE by projecting returns onto lagged prices. As with the projections leading to Table 4, however, we discovered pronounced heteroscedasticity in the error term: its variance decreased with the price level. Again, a simple transformation purged the error of heteroscedasticity, namely, multiplication of the return by the lagged price. Of course, this means that the regressand became simply the change in the price over the subsequent day.

Table 5 reports the results of the OLS projections of price changes onto lagged prices. A positive intercept and a negative slope coefficient appear, both highly significant. It is interesting to note that this is precisely what has been observed for US stock returns as well (see, e.g., Keim and Stambaugh, 1986). Of course, this link is only suggestive: it is quite a leap of faith to claim that

Table 5
 OLS projections of daily changes in prices onto lagged price levels ^a

	All markets	Microsoft markets	Computer industry markets
<i>N</i>	1792	591	1200
Intercept	0.050** (0.006)	0.058** (0.013)	0.048** (0.007)
Slope	-0.053** (0.009)	-0.069** (0.019)	-0.048** (0.010)
<i>R</i> ²	0.020	0.022	0.019

^a *N* denotes sample size (number of daily returns). * and ** indicate that the statistic is more than 1.8 and 2 standard errors from zero, respectively.

findings for the Iowa experimental markets have any implications for official US stock exchanges.

The evidence in Table 5 leads us to refute REE at high confidence levels. Because we concluded in the previous section that closing prices in the Iowa markets reflect rational learning, however, we cannot attribute this rejection to over-reaction (the explanation usually advanced for this phenomenon in the stock markets). It must be agents' beliefs. They are biased, in violation of one of the assumptions behind rational expectations.

In what sense are beliefs biased? A closer investigation of the data revealed the following. Take the Microsoft markets, for instance. The Microsoft High contract matured in the money twelve times out of sixteen (months). This seems to indicate that chances for Microsoft High to pay off \$1 are 75%, which would lead us to assign a value of \$0.75 at the beginning of trading. We noticed, however, that the first closing price for Microsoft High was often substantially different from \$0.75. The deviation for the 9/95 contract, for instance, was more than \$0.30 (see Fig. 1).

By themselves, these deviations from what we consider to be a fair value for the Microsoft High contract may merely be evidence of the superior information that markets possess over us, or they may reflect that the Microsoft High contract does not always start trading at the money. However, together with the negative relationship between price levels and subsequent returns, the deviations in the initial Microsoft High prices from fair value seem to indicate that Iowans started out with more extreme priors than can be justified *ex post*. Beliefs subsequently reverted back to some more acceptable level.

One would of course want more precise estimates of the nature of the belief biases over the particular period under study. Unfortunately, we observe only sixteen (one-month) histories of the evolution of the beliefs. This means that we implicitly only have sixteen sample points with which to compare the final outcome with the initial price (belief). Such a sample size is simply too small to make any confident statements about the direction of the bias in the market's priors.

In this respect, it is worth emphasizing that a comparison between the negative correlation between price levels and subsequent price changes (documented in Table 5), on the one hand, and the absence of correlation between the modified return and lagged prices for winning securities (documented in Table 4), on the other hand, did allow us to conclude that initial beliefs must have been biased. It underscores the power of our methodology. However, we cannot go further and estimate precisely what the bias in initial beliefs was like.

As a final note, we should caution the reader against overinterpreting our evidence against rational expectations. Return measures do not always accurately reflect potential profitability. Our returns are computed on the basis of *transaction prices*. Transactions could take place either at the bid or at the ask. If today's closing transactions took place at the bid, and tomorrow's at the ask,

the return one computes from the closing prices measures incorrectly the payoff on a strategy of purchasing the security at today's close and selling at tomorrow's close.⁶

Among other things, the regression results reported in Table 5 could be affected by this bid–ask bounce, if closes with a low price predominantly take place at the bid. The subsequent return would be biased upward. This bias would translate into a negative correlation between the error in the regression of returns onto previous closing prices and the regressor. The OLS slope coefficient would be negatively biased. Absent quote data, however, we have no way to tell how severe this problem is.

Still, it is important to appreciate that if bid–ask bounce biased the regressions with the traditional return as regressor, it ought to have affected the ones with the modified return as regressand as well. Tables 3 and 4 documented, however, that the results for the latter conform nicely to the theory. This makes our evidence against rational expectations less ambivalent.

6. Conclusion

We cannot find any predictability in return sequences of winning securities in the Iowa experimental markets where the payoffs are defined with respect to US exchange-listed stock. The returns were suitably modified to adjust for the selection bias caused by our focusing on securities that matured in the money. Since these modified returns are unpredictable if agents correctly apply the rules of conditional probability, we conclude that we cannot reject that market prices in Iowa reflect rational learning.

In contrast, we find evidence of in-sample predictability of returns, in the form of a negative correlation between the daily price change and the level of the previous closing price. Since in-sample predictability violates rational expectations but agents appear to have correctly implemented Bayes' law, we conclude that their priors must have been biased. We do not have enough data, however, to estimate the precise nature of the biases in initial beliefs that participants in the Iowa markets appear to have had over the period under investigation.

What does this mean in terms of the traditional notion of market efficiency (Fama, 1970, 1991), i.e., for a newcomer's ability to "make money"? Based on the evidence reported here, a newcomer should refrain from betting against the Iowa markets. There is no money to be made from mistakes against the rules of conditional probability. We did find that priors in the Iowa markets were

⁶ Several papers in the finance literature point out the potential biases from using transaction prices, starting with Blume and Stambaugh (1983).

biased. But these biases may have been period-specific. In other words, there is no guarantee that they would still be there in the future. Moreover, we did not have enough data to estimate their nature within acceptable limits. The bottom line is: the marginal bet in the Iowa market is pure speculation.

Acknowledgements

We thank Tom Rietz for sending us price and volume series from the Iowa Experimental Market. We are grateful for the comments from Stephen D. Smith as well as an anonymous referee.

Appendix A. A general proof of Result 6 (Section 2)

We use the notation of Section 2, but allow the signals to be general. In particular, let s denote the signal, with (continuous) likelihoods $l(s|\bar{\theta})$ and $l(s|\underline{\theta})$. We prove here that

$$E[x_2|\bar{\theta}] = 0,$$

which is Eq. (15). Writing the expectation as an integral, remembering that $x_2 = 1 - p_1/p_2$, and applying Bayes' law, we obtain

$$\begin{aligned} E[x_2|\bar{\theta}] &= \int \left[1 - \frac{p_1}{p_2} \right] l(s|\bar{\theta}) \, ds \\ &= 1 - \int \frac{p_1 [p_1 l(s|\bar{\theta}) + (1 - p_1) l(s|\underline{\theta})]}{p_1 l(s|\bar{\theta})} l(s|\bar{\theta}) \, ds \\ &= 1 - \int [p_1 l(s|\bar{\theta}) + (1 - p_1) l(s|\underline{\theta})] \, ds = 0. \end{aligned}$$

Notice that the result does not depend on the prior (p_1), which cancels out. For more details and generalizations, see Bossaerts (1996, 1997).

References

- Blume, L., Stambaugh, R.F., 1983. Biases in computed returns: An application to the size effect. *Journal of Financial Economics* 12, 387–404.
- Bondarenko, 1997. Testing rationality of financial markets: An application to S&P 500 index options. Caltech working paper.
- Bossaerts, P., 1995. The econometrics of learning in financial markets. *Econometric Theory* 11, 151–189.

- Bossaerts, P., 1996. Martingale restrictions on equilibrium security prices under rational expectations and consistent beliefs. Caltech working paper.
- Bossaerts, P., 1997. The dynamics of equity prices in fallible markets. Caltech working paper.
- Bossaerts, P., Hillion, P., 1997. IPO post-issue markets: Questionable predictions but diligent learners? Caltech working paper.
- Fama, E., 1970. Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25, 383–417.
- Fama, E.F., 1991. Efficient capital markets: II. *Journal of Finance* 46, 1575–1618.
- Forsythe, R., Nelson, F., Neumann, G.R., Wright, J., 1992. Anatomy of an experimental political stock market. *American Economic Review* 82, 1142–1661.
- Green, J., 1973. Information, efficiency and equilibrium. Discussion Paper 284. Harvard Institute of Economic Research, Cambridge, MA.
- Hansen, L.P., Singleton, K., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50, 1269–1286.
- Harrison, J.M., Kreps, D.M., 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20, 381–408.
- Harsanyi, J.C., 1967. Games with incomplete information played by Bayesian players, I: The basic model. *Management Science* 14, 159–182.
- Keim, D., Stambaugh, R., 1986. Predicting returns in the stock and bond markets. *Journal of Financial Economics* 17, 357–391.
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1336.
- Lucas, R., 1972. Expectations and the neutrality of money. *Journal of Economic Theory* 4, 103–124.
- Lucas, R., 1978. Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.
- Muth, J.F., 1961. Rational expectations and the theory of price movements. *Econometrica* 29, 315–335.
- Radner, R., 1967. Equilibre des marchés à terme et au comptant en cas d'incertitude, *Cahiers du Séminaire d'Économétrie*, pp. 35–52.
- Samuelson, P.A., 1965. Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* 6, 41–50.