DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

RATIONAL PRICE DISCOVERY
IN EXPERIMENTAL AND FIELD DATA

Peter Bossaerts

SOCIAL SCIENCE WORKING PAPER 952

July 1995

Rational Price Discovery In Experimental And Field Data

Peter Bossaerts

Abstract

The methodology of tests for martingale properties in return series is analyzed. Martingale results obtain frequently in finance. One case is focused on here, namely, *rational price discovery.* Price discovery is the process by which a market moves towards a new equilibrium after a major event. It is rational if price changes cannot be predicted from commonly available information. The price discovery process, however, cannot be assumed stationary. Hence, to avoid false inference in the presence of nonstationarities, event studies of field data have been advocating the use of cross-sectional information in the computation of test statistics. Under the martingale hypothesis, however, this inference strategy is shown to add little except if higher moments of the return series do not exist. On the contrary, the cross-sectional approach may even be invalid if there is cross-sectional heterogeneity in the price discovery process. The time series statistic of Patell [1976], originally suggested in the context of *i.i.d.* time series but cross-sectional heteroscedasticity, may be preferable. It will not provide valid inference either, if higher serial correlation coincides with higher volatility. Unfortunately, this appears to be the case in the dataset which is used in the paper to illustrate the methodological issues, namely, transaction price changes from experiments on continuous double auctions with stochastic private valuations.

Keywords: Price Discovery, Rational Expectations, Martingales, Nonstationarity, Nonergodic Central Limit Theorems, Event Studies, Experimental Economics.

Rational Price Discovery In Experimental And Field Data*

Peter Bossaerts

1 Introduction

Let r_t denote the change of the price of an asset over the period $t-1, t$. One often conjectures that return series are a martingale relative to publicly available information, i.e.,

$$E[r_t|\mathcal{F}_{t-1}] = 0, \tag{1}$$

where $\mathcal{F}_{t-1}$ denotes the public information set at time $t-1$. The aim of this paper is to discuss the methodology of tests of such martingale hypotheses.

This will be done in a the following context. In a market mechanism where participants meet continuously, a price adjustment process is set in motion each time a shock hits the market. During this process, the market searches for the new equilibrium price level. Indeed, price adjustment is rarely instantaneous, and the transient *price discovery* has become the focus of attention among many researchers, especially in experimental economics (see, e.g., Plott and Sunder [1988], Smith, Suchanek and Williams [1988], Camerer and Weigelt [1993], Cason and Friedman [1993] and [1994], Gjerstad [1994], Kagel [1994]).

Price discovery is essentially a *learning process*: market participants observe signals in the trading process, from which they attempt to infer the likely continuation of the adjustment. It is also a *bargaining process*: participants' quotes and transactions reveal information, which generate subsequent trades, and, hence, transaction prices. Explicit

*California Institute of Technology; e-mail: pbs@rioja.caltech.edu. The content of this paper was the basis of an invited talk at the 1994 European Summer Meetings of the Econometric Society. Charles Plott and David Porter graciously provided the author with their experimental data. They and Tom Palfrey commented on an earlier draft. The research was partly supported through a grant from First Quadrant to the California Institute of Technology.

models of this learning-*cum*-bargaining process can be found in Wilson [1987], Easley and Ledyard [1993], Friedman [1984], Gjerstad [1994].

Here is a hypothesis that could be the basis of a general empirical analysis: *price discovery ought to be rational.* As mentioned before, price discovery involves learning. Learning can be called rational whenever the learner cannot predict where her beliefs will go next. If the learner starts with certain probabilistic assessments (priors) about the uncertainty around her, rationality requires that she apply correctly the rules of conditional probability (Bayes' law). If transaction prices reflect the learning of one specific, risk-neutral participant in the market (the marginal or representative agent), rationality will be reflected in absence of predictability from that participant's point of view. This person will refrain from speculating, because she beliefs that prices follow a martingale. If, in addition, the marginal agent has correct priors, it follows that transaction prices must not be predictable even when weighting all likely outcomes with their actual frequencies. In other words, if her beliefs are unbiased across repetitions of the market adjustment process, prices will be a martingale for an outside observer as well. Formally, one obtains the hypothesis in Eqn. (1). The assumption that priors are unbiased is usually referred to as *rational expectations.*

A test of the rationality of price discovery, Eqn. (1), seems straightforward. One could collect time series of returns following a particular event, and check whether they are serially uncorrelated. This could be done by means of the average of future times past returns, which should be zero under the null. To avoid parametric assumptions, the sample average is computed from a long time series, and central limit results are appealed to. These state that the sample average divided by its asymptotic standard error is standard normally distributed in large samples. The asymptotic standard normality of this *z-statistic* provides the basis for inference.

The implementation of this methodology, however, is not without problems. In particular, learning drives price discovery. Learning is nonstationary, so that price discovery will be as well. In particular, the conditional volatility can be expected to change systematically over time as the market accumulates information about the equilibrium price level. One would suspect that this poses problems for the estimation of the asymptotic standard error of the rescaled sample average. An estimator that requires homoscedasticity will almost certainly cause incorrect inference. Unfortunately, such an estimator has been widely used in experimental studies of price discovery (Cason and Friedman [1993] and [1994], Gjerstad [1994], Kagel [1994]). In contrast, many event studies of field data have been advocating the use of *cross-sectional* information in the estimation of the standard error. There are many examples: e.g., the studies of return patterns following initial public offerings (e.g., Ritter [1991]), or, studies of the adjustment process following dividend and earnings announcements (e.g., Patell and Wolfson [1984]).

This paper evaluates the theoretical validity of cross-sectional tests of the martingale hypothesis in Eqn. (1). To clarify rather abstract results from the mathematical statistics literature, data from a set of experiments are used. The experiments were recently held at Caltech, and concern continuous double auctions. Traders had private valuations that sometimes changed randomly from one trading round to another. Only when the private valuations changed, a new price discovery process was set in motion. Otherwise, price adjustment is absent. Hence, the experiments include a set of 'normal' trading rounds against which the influence of the price discovery process could be evaluated. They correspond to the *control samples* in the aforementioned event studies.

The paper argues that there is little theoretical reason for a cross-sectional estimate of the standard error. The martingale assumption (Eqn. (1)) imposes enough structure for a simple time-series heteroscedasticity-adjusted estimator to be sufficient. Adjustment for autocorrelation is obviously not called for. Most importantly, it is robust to cross-sectional heterogeneity in the price discovery process. In particular, the conditional volatility patterns may be sample-dependent, i.e., volatility is *nonergodic.* The preferred statistic is the one that was originally proposed in Patell [1976] in the context of *i.i.d.* return series with cross-sectional heteroscedasticity.

Path dependencies are likely to arise in the context of price discovery. The evolution of the precision of agents' beliefs will depend on the accumulated evidence. Consequently, volatility is likely to be history-dependent. Likewise, price discovery in a continuous market is the outcome of a complex multilateral bargaining game with asymmetric information. These games are known to have multiple equilibria. The actual equilibrium which is played will be context-dependent (parameters, agents). Therefore, the volatility of transaction price changes is likely to vary from one equilibrium to another, and, hence, across experiments.

One of the main advantages of the experimental data (or, of event study data, for that matter) is that z-statistics can be computed in several ways. Evidence for time series anomalies can easily be discovered when comparing the inference across the various z-statistics. To understand this, remember that *all* z-statistic would generate the same conclusions, for instance, if the data were *i.i.d.* over time and in cross-section (provided, of course, that the samples are large enough). The experimental data clearly indicate that the inference differs dramatically across z-statistics. These discrepancies form the basis for the theoretical-statistical analysis of tests of the rationality of price discovery.

One question obviously emerges: which way of computing the z-statistic is correct? A comprehensive analysis of the data reveals that *even Patell's approach is invalid.* The data appear not to satisfy the conditions for asymptotic normality of Patell's statistic, outlined in Hall [1977]. In particular, there is strong evidence that serial correlation increases whenenever volatility is higher, in violation of one of Hall's most crucial assumptions.

This paper complements Bossaerts [1995], where the actual large-sample distribution of simple statistics is derived in the context of two models of price adjustment.[1] The asymptotic distributions are nonstandard. They are those of functionals of Brownian motions. The present paper illustrates the concerns of Bossaerts [1995], not by means of theoretical examples, but using experimental data. For the purpose of this paper, Bossaerts [1995] should therefore be viewed as providing a theoretical basis for the lack of stationarity and ergodicity in price discovery, which will be documented here to be actually present in experimental data.

Bossaerts [1995] and the present paper complement each other in another respect. In the former, the statistics are fixed beforehand, and the actual large-sample behavior is derived. Here, we attempt to alter the statistics in order to obtain asymptotic standard normality. In other words, we attempt to find the statistic that is robust to nonstationarities and even nonergodicities in the data.

The next section presents the experiments with which the paper illustrates the theoretical-statistical arguments. Section 3 isolates a control sample, where no price adjustment takes place. The striking differences in serial correlation between the control sample and the remaining sample motivates the theoretical-statistical analysis. Section 4 discusses a general martingale central limit theorem that is to form a robust basis for further inference. Section 5 elaborates on and illustrates two implementations of this martingale invariance result. Section 6 explores the causes behind the discrepancies between the results from different implementation strategies. Section 7 concludes.

2 The Experiments

We study five experiments that Charles Plott recently ran at the California Institute of Technology. The aim was to study rational expectations equilibria in continuous double auctions with stochastic private valuations. The experiments are a large-scale replication of the experiments reported in Plott and Agha [1983]. They consist of pairs of trading rounds. In the first trading round, trade takes place under certainty, i.e., the distribution of private valuations is known. In the second trading round, some uncertainty is introduced, in the form of stochastic private valuations for agents on the demand side. Traders are given the opportunity to carry over inventory between two periods in a pair.

The experiments generally lasted twenty periods (ten pairs of two trading rounds). We shall ignore the first couple, as these were practice runs designed to familiarize subjects with the computer software. Participants were not paid during the practice rounds. They

[1] Bossaerts [1995] focuses only on the learning aspect of price discovery, ingnoring the strategic issues, which form the core of the analyses of Wilson [1987] and Easley and Ledyard [1993].

were paid, however, during the remainder of the experiment. The pay depended on their trading performance.

We shall look at transaction prices only, and, in particular, their serial correlation patterns. Table 1 lists simple first-order autocorrelation coefficients, as well as the corresponding z-statistics computed by rescaling (multiplying) the sample autocorrelation coefficients with the square root of the time series length, and denoted z_{tsh}.[2] Only intra-trading-round transaction price changes are used.

The amount of *negative serial correlation* is striking. It confirms the results of other experiments on continuous double auctions with stochastic private valuations (Cason and Friedman [1993] and [1994], Gjerstad [1994], Kagel [1994]). One should be careful, however, not to draw too quickly the conclusion that price adjustment in experimental markets is irrational. From studies of field data, we know that negative serial correlation will emerge if the minimum tick size puts a lower bound on the bid ask spread and if the volatility of transaction price changes is low (see Roll [1984]; for evidence, see, e.g., Patell and Wolfson [1984]).[3] In the experiments, the tick size is very low (less than 0.5% of the equilibrium price level), but the volatility decreases dramatically as prices reach the equilibrium level or in the first of each pair of trading round, where there is no uncertainty about the distribution of private valuations. Figure 1 illustrates this: it plots transaction price changes in transaction time (as opposed to calendar time) for one of the experiments. The intra-trading-round transaction prices were concatenated to obtain one long time series.

There may be an additional cause for spurious negative serial correlation in transaction price changes. In the experiments, private valuations lie on a discrete grid. Once infra-marginal units have changed hands, an artificially high bid-ask spread would emerge merely as a result of this discrete grid. Of course, there is still the issue that no transactions should take place anymore once infra-marginal units have traded. In other words, the artificially high bid-ask spread ought not to show up in transaction price changes.

In any event, to remedy the problem of potentially spurious negative serial correlation, one could look at only one type of transactions, namely accepted asks (one could also investigate the other side, namely, of accepted bids; the conclusions do not differ). These cannot possibly be affected by spurious autocorrelation because of bid-ask bouncing. Table 2 lists serial correlation properties of accepted asks in the five experiments. The

[2]Actually, the z-statistic is computed in a slightly different way. It is based on a sample autocorrelation coefficient for which neither the covariance (in the numerator) nor the variances (in the denominator) are adjusted for the mean. Under the null of Eqn. (1), this is allowed. Adjustment for the mean has no qualitative effect on the results.

[3]An explicit relation between serial correlation and the bid-ask spread is given in Roll [1984]. Harris [1994] provides an in-depth analysis of the economic effect (volume, quotation size, etc.) of minimum tick size.

picture is different. Highly significant negative autocorrelation continues to appear in the majority of the experiments, but a significantly positive serial correlation shows up for the fourth experiment. The uniformity in the pattern of autocorrelations of transaction price changes is broken.

3 Introducing Control Samples

As mentioned before, the experiments consist of pairs of trading rounds. In the first one, there is no uncertainty about private valuations. Price discovery ought to be absent. In the second trading round, private valuations on the demand side are stochastic. Hence, price discovery should be apparent. A comparison of transaction price patterns across the two types of trading rounds should reveal something about the effect of price discovery. To verify this, let us split the samples in two parts: one were we concatenate price changes from the first trading rounds in each pair, and a second part, where price changes from the second trading rounds are linked together to form one long series. We shall refer to the first series as the "control sample". The second series will be called the "adjustment sample."

Table 2 reports autocorrelation coefficients and corresponding z-statistics for the two samples. There is a pronounced difference. In only one experiment can one find significant (negative) serial correlation in the control sample. In contrast, significance emerges in all but one adjustment sample. The difference in the serial correlation patterns between the control and adjustment samples is striking. The significant autocorrelation coefficients in the adjustment sample indicate that price discovery is far from rational.

But is this really the conclusion we should draw? The statistical properties of the adjustment sample may differ in many respects, to the point that the z-statistics are not standard normal anymore, invalidating our inference. There is one feature of price discovery which leads us to suspect that this is the case, namely nonstationarity.

4 A Martingale Central Limit Theorem

Price discovery is predominantly a learning phenomenon: traders accumulate information about other agents' private valuations and about the equilibrium price level. The cumulation of the information over time is likely to induce nonstationarity in the transaction price process. One can illustrate this in a theoretical model, as in Bossaerts [1995]. In the experiments, however, there may be an opportunity to see nonstationarity directly in action.

One likely type of nonstationarity that price adjustment could induce in transaction price changes is *heteroscedasticity.* A look at Figure 1 confirms that there are periods of sudden increases in volatility (corresponding to the beginning of trading rounds with stochastic valuations), followed by a steady decay. This upsets the inference from the z-statistics in Tables 1 and 2: they are standard normally distributed *only* under homoscedasticity.

What would be an appropriate z-statistic? The only restriction the null hypothesis imposes is Eqn. (1). In all other respects, generality should be aimed for. Hall [1977] derived a result that could form the basis for alternative z-statistics. It states the following. Define:

$$S_T = \sqrt{T}\left(\frac{1}{T}\sum_{t=1}^{T} r_t r_{t-1}\right).$$

Theorem 4.1 *Assume: (i) Eqn. (1), (ii)*

$$\lim_{T\to\infty} E[\max_{1\leq t\leq T} \frac{(r_t r_{t-1})^2}{T}] = 0$$

(a Lindeberg condition), (iii)

$$\frac{1}{T}\sum_{t=1}^{T} E[(r_t r_{t-1})^2 | \mathcal{F}_{t-1}] \to \sigma \tag{2}$$

in probability, with $\sigma > 0$ almost surely, and (iv), σ is measurable in $\mathcal{F}_0$. Then:

$$\frac{S_T}{\sigma} \rightsquigarrow N(0,1).$$

This is a typical martingale central limit result in that it restricts the behavior of conditional volatilities instead of the memory (mixing) of the process in order to generate asymptotic standard normality. Notice how general it is: σ is not constrained to be constant across time series. The reader may be unfamiliar with martingale processes where the average conditional volatility is sample-dependent, so, let us assume for the time being that σ is a constant. Only if the data force us to reconsider this assumption will we do so.

σ is unknown, and, therefore, has to be estimated somehow from the data. In other words, the implementability of Theorem 4.1 depends crucially on whether one can formulate a consistent estimator for σ. Two strategies will be discussed next.

5 Implementation

Assuming that one has a cross-section of independent time series, there is a very simple estimator for σ, namely the cross-sectional standard deviation of S_T. Let $S_T(i)$ denote the value of S_T in the ith time series ($i = 1, ..., N$). Then:

$$(\hat{\sigma}_c)^2 = \frac{1}{N}\sum_{i=1}^{N}(S_T(i))^2. \tag{3}$$

σ is really the asymptotic *cross-sectional* standard deviation of S_T. Provided the realizations of S_T are independent in cross-section, the cross-sectional sample standard deviation ($\hat{\sigma}_c$) provides a consistent estimator.

The simplicity of the estimator in (3) has lead many to use it in event studies, where one has a cross-section of time series, namely the return pattern after the event being studied. If the time series are clustered in time, cross-sectional dependence could be reduced by a matched-sample technique, whereby appropriately matched non-event time series are subtracted from the return series at hand. For examples, see the event analyses of Patell and Wolfson [1984] and Ritter [1991].

But there is an alternative, equally simple estimator. Consider

$$(\hat{\sigma}_{ts})^2 = \frac{1}{T}\sum_{t=1}^{T}(r_t r_{t-1})^2. \tag{4}$$

$\hat{\sigma}_{ts}$ is a consistent estimator of σ under very weak conditions. To see this, write $(\hat{\sigma}_{ts})^2$ as the sum of a term that converges to σ^2 and the sample average of a martingale difference sequence:

$$\begin{aligned}
&(\hat{\sigma}_{ts})^2 \\
&= \frac{1}{T}\sum_{t=1}^{T} E[(r_t r_{t-1})^2|\mathcal{F}_{t-1}] + \left(\frac{1}{T}\sum_{t=1}^{T}(r_t r_{t-1})^2 - \frac{1}{T}\sum_{t=1}^{T} E[(r_t r_{t-1})^2|\mathcal{F}_{t-1}]\right) \\
&= \frac{1}{T}\sum_{t=1}^{T} E[(r_t r_{t-1})^2|\mathcal{F}_{t-1}] + \frac{1}{T}\sum_{t=1}^{T}\left((r_t r_{t-1})^2 - E[(r_t r_{t-1})^2|\mathcal{F}_{t-1}]\right).
\end{aligned}$$

By assumption, the first term converges to σ. The second term is the sample average of a martingale difference sequence. For the latter to converge to zero, it suffices that higher moments of the elements of the martingale difference sequence exist (a martingale law of large numbers; see, e.g., White [1983], p. 58). No further requirements are needed.

As a matter of fact, $\hat{\sigma}_{ts}$ is a very robust estimator of σ. In particular, σ could vary in cross-section. $\hat{\sigma}_c$ cannot possibly capture cross-sectional variation in σ, and, hence, would be an inconsistent estimator in that case. We shall come back to this point.

Let us first look at the experimental data and compute z-statistics on the basis of $\hat{\sigma}_c$ and $\hat{\sigma}_{ts}$. Denote:

$$z_c = \frac{S_T}{\hat{\sigma}_c}. \tag{5}$$

Likewise:

$$z_{ts} = \frac{S_T}{\hat{\sigma}_{ts}}. \tag{6}$$

Table 3 reports the z-statistics for the experiments. Only results from the disaggregated samples (control sample, adjustment sample) are displayed. The z_{tsh}s (the z-statistics that requires homoscedasticity) are taken from Table 2.

Several observations can immediately be made. First, notice the striking contrast between z_{tsh}, on the one hand, and z_c or z_{ts}, on the other hand, in the adjustment sample. No such discrepancy emerges for the control sample, indicating that non-event return series can even be taken to be homoscedastic. In the adjustment sample, both z_c and z_{ts} lead to different inference compared to z_{tsh}. This illustrates how nonstationarity affects inference when there is price discovery. In other words, one cannot simply assume that the post-event adjustment series is stationary.

Second, there are discrepancies even between z_c and z_{ts}. Not only across experiments, but also in the overall χ^2-statistic. The latter is based on the following simple argument. Both z_c and z_{ts} are standard normal (asymptotically). The sum of uncorrelated squared standard normal random variables is χ^2 distributed, with degrees of freedom equal to the number of terms in the sum. The χ^2-statistics in Table 3 are computed by taking the squares of the z-statistics and summing across experiments. Under the null, the result should be χ^2 distributed with five degrees of freedom (one can safely assume that there is independence across experiments).

The overall χ^2 rejects the null that returns form a martingale difference sequence (Eqn. (1)) when based on z_{ts}, but not when based on z_c. The discrepancy is not marginal. Moreover, z_c and z_{ts} do not attribute rejections to the same experiments. z_{ts}, for instance, finds substantial positive serial correlation in experiment cp112493, but z_c does not agree: zero autocorrelation cannot be rejected even at low levels of significance.

6 Explaining The Discrepancies

It is unlikely that z_{ts} fails because $\hat{\sigma}_{ts}$ is an inconsistent estimator of σ: this would occur only if higher moments of the return series do not exist. z_c, however, leads to correct inference only if σ does not change in cross-section. z_{ts} does not require this. It is the appropriate estimator with which to implement Theorem 4.1 when σ is suspected of varying across time series.

It was already mentioned that Theorem 4.1 is very general, because σ, the asymptotic time series average conditional variance, may be a random variable. This contrasts with traditional martingale central limit theorems, where σ remains constant across time series (see Brown [1971], McLeish [1974]). Actually, the Theorem can also be described as an asymptotic mixtures-of-normals result, i.e., it is a *nonergodic* central limit theorem. In particular, one can write:

$$S_T \rightsquigarrow \sigma . z,$$

where z is a standard normal random variable, mean-independent of σ. The mean-independence of z relative to σ comes from the restriction in the Theorem that σ be measurable in $\mathcal{F}_0$. It is crucial, as Example 3.3 in Hall [1974] illustrates. Loosely speaking, it means that there must not be correlation between the level of the conditional volatilities (actually, the time series average conditional volatility) and the (signed) return series (more precisely, the rescaled average of returns times past returns).

Before citing evidence that σ varies across experiments in the data, we should briefly discuss whether there are *a priori* reasons to believe that such would be the case in the context of price discovery. There are at least two reasons. As mentioned before, the process of price adjustment in a continuous market is essentially a learning process. In such a process, the evolution of the precision of the beliefs depend critically on the accumulated evidence. Therefore, it is path-dependent, or *nonergodic*. The lack of ergodicity of the precision translates into path-dependence (nonergodicity) of the price adjustment process in general, and the conditional volatility in particular. A simple analytical example can be found in Bossaerts [1995] (his "Case II"). Hence, the asymptotic average conditional volatility (σ) will reflect this: it will be path-dependent.

Another reason for the sample dependence of σ comes from a game-theoretic analysis of the price discovery process. One way to model the price adjustment is to consider it to be a concatenation of bargaining games under asymmetric information (as in Wilson [1987]). We know that multiplicity of equilibria plague the analysis of such games, and, hence, that the price adjustment path will be case-dependent. In other words, one expects to see heterogeneity across experiments. In particular, the asymptotic average conditional volatility may vary. Theorem 4.1 allows this.

At this point, I should comment on the way I organize the data, which is to concatenate transaction price changes from different trading rounds in an experiment into a single time series. Why not consider each trading round as a separate price adjustment process, i.e., as a separate time series? The resulting time series would be of insufficient length to make reliable inference. The concatenation technique has the added advantage that it effectively considers the entire experiment to be a single, long price adjustment, consisting of shorter price discovery subperiods. Recent statistical analysis of game-theoretic experiments has found systematic learning effects over entire experiments, even if this consisted of a sequence of independent subgames. Agents appear to extract information about the likely behavior of opponents, to be used in future, independent replications of the same game. See, e.g., McKelvey and Palfrey [1993] and El-Gamal, McKelvey and Palfrey [1994]. In other words, whereas the parameters may be independent across trading rounds of an experiment, beliefs, and, hence, the price discovery processes, are not.[4]

There is evidence in the experimental data that σ varies in cross-section. Table 3 reports $\hat{\sigma}_{ts}$s. They vary substantially across experiments. Hence, sample dependence of σ seems to be the reason why there is a discrepancy between the inference from z_c and z_{ts}.

Because of the cross-sectional variation of σ, z_c is not a valid implementation of Theorem 4.1. In contrast, z_{ts} would be, provided the rescaled average return times lagged return (S_T) is mean-independent of σ. To ascertain whether this condition is satisfied in the experimental data, one could regress S_T onto $\hat{\sigma}_{ts}$. An ordinary least squares regression generated the following results:

$$\begin{array}{ccccccc} S_T & = & 86.4 & - & 1.3 & \hat{\sigma}_{ts} & + \ \hat{\epsilon} \\ & & (0.150) & & (-4.525) & & \end{array} \tag{7}$$

(t-statistics in parentheses).[5] There appears to be a significant negative relationship between S_T and $\hat{\sigma}_{ts}$: strong negative serial correlation coincides with high volatility. Because of errors-in-variables, the true dependence between S_T and σ is likely to be even more pronounced.

Consequently, we cannot trust z_{ts} either. It is not a correct implementation of Theorem 4.1. For valid inference, conditional volatility and serial correlation must not be correlated.

7 Conclusion

When testing the simple martingale difference hypothesis in Eqn. (1) in the context of price discovery, one should be careful about nonstationarity and nonergodicity in transaction price changes. The experimental data from continuous double auctions with stochastic private valuations amply illustrate this. Lack of stationarity and even ergodicity affect these data to the point that none of the traditional test strategies are valid. In

[4]In a previous draft, I reported results from tests on the alternative organization of the data, whereby each trading round in an experiment is considered to be a separate time series. Computing the corresponding z_cs for each experiment reveals absolutely no evidence of serial correlation in transaction price changes, not even in the adjustment sample.

[5]The estimates are based on a pooling of the adjustment and control samples; the results do not alter qualitatively when considering only the adjustment sample

other words, *the traditional z-statistics do not generate conclusive evidence as to whether the fairly substantial serial correlation in the transaction prices is significant.*

Is there a way to make valid inference in nonstationary, nonergodic series? In general, the answer is negative. It appears that test statistics and their properties have to be developed in the context of a particular model of price adjustment. The asymptotics on which the inference in based are likely to be nonstandard. An example can be found in Bossaerts [1995], where the asymptotic distribution of simple z-statistics is derived for two particular price adjustment models. Both cases generate asymptotic laws that involve functionals of Brownian motions.

There is a negative relationship between serial correlation and average conditional volatility in the second model of Bossaerts [1995]. The same relationship is present in the experimental data and lead us to conclude that Patell's statistic cannot be used as a basis for inference, despite its obvious advantages.

References

Bossaerts, P., 1995, "The Econometrics of Learning in Financial Markets," *Econometric Theory* 11, 151-189.

Brown, B.M., 1971, "Martingale Central Limit Theorems," *Annals of Mathematical Statistics* 42, 59-66.

Camerer, C. and K. Weigelt, 1993, "Convergence in Experimental Double Auctions for Stochastically Lived Assets," in: *The Double Auction Market*, Eds. D. Friedman and J. Rust. Reading, MA, Addison-Wesley.

Cason, T. and D. Friedman, 1993, "An Empirical Analysis of Price Formation in Double Auction Markets," in: *The Double Auction Market*, Eds. D. Friedman and J. Rust. Reading, MA, Addison-Wesley.

Cason, T. and D. Friedman, 1994, "Price Formation in Double Auction Markets," University of Southern California, working paper.

Easley, D. and J. Ledyard, 1993, "Theories of Price Formation and Exchange in Double Oral Auctions," in: *The Double Auction Market*, Eds. D. Friedman and J. Rust. Reading, MA, Addison-Wesley.

El-Gamal, M., R. McKelvey and T. Palfrey, "A Bayesian Sequential Experimental Study of Learning in Games," *Journal of the American Statistical Association* 80, 428-435.

Gjerstad, S., 1994, "Empirical Features of Experimental Markets," University of Minnesota, working paper.

Hall, P., 1977, "Martingale Invariance Principles," *Annals of Probability* 5, 875-887.

Harris, L., 1994, "Minimum Price Variations, Discrete Bid-Ask Spreads, and Quotation Sizes," *Review of Financial Studies* 7, 149-178.

Kagel, J., 1994, "Double Auction Markets with Stochastic Supply and Demand Schedules: Call Markets and Continuous Auction Trading Mechanisms," University of Pittsburgh, working paper.

McKelvey, R. and T. Palfrey, 1993, "An Experimental Study of the Centipede Game," *Econometrica* 60, 803-836.

McLeish, D.L., 1974, "Dependent Central Limit Theorems and Invariance Principles," *Annals of Probability* 2, 620-628.

Patell, J.M., 1976, "Corporate Forecasts of Earnings per Share and Stock Price Behavior: Empirical Tests," *Journal of Accounting Research* 14, 246-276.

Patell, J.M. and M.A. Wolfson, 1984, "The Intraday Speed of Adjustment of Stock Prices to Earnings and Dividend Announcements," *Journal of Financial Economics* 13, 223-252.

Plott, C. and G. Agha, 1983, "Intertemporal Speculation with a Random Demand in an Experimental Market," in: *Aspiration Levels in Bargaining and Economic Decision Making*, ed. R. Rietz. New York, Springer Verlag.

Plott, C. and S. Sunder, 1988, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Securities Markets," *Econometrica* 56, 1085-1118.

Ritter, J.R., 1991, "The Long Run Performance of Initial Public Offerings," *Journal of Finance* 46, 3-27.

Roll, R., 1984, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance* 39, 1127-1139.

Smith, V.L., G.L. Suchanek, A.W. Williams, 1988, "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica* 56, 1119-1152.

White, H., 1984, *Asymptotic Theory for Econometricians*. Orlando: Academic Press.

Wilson, R., 1987, "On Equilibria of Bid-Ask Markets," in: *Arrow and the Ascent of Modern Economic Theory*, ed. G. Feiwel. London, MacMillan.

Table 1
Serial correlation properties of intra-trading-round transaction prices in five experiments

Experiment	ρ	z_{tsh}	T
cp032294	-.397	-6.861^{**}	298
cp032394	-.419	-7.972^{**}	362
cp110193	-.522	-8.638^{**}	274
cp112493	-.253	-4.534^{**}	322
cp120893	-.483	-9.471^{**}	384

Remarks: ρ denotes the first-order autocorrelation coefficient of transaction price changes; z_{tsh} is the corresponding z-statistic, obtained by multiplying ρ by the square root of the sample size; T is the sample size. ** denotes significance at the 1% level.

Table 2
Serial correlation properties of intra-trading-round accepted asks in five experiments

Experiment	Total Sample			Adjustment Sample				Control Sample			
	ρ	z_{tsh}		ρ	z_{tsh}		T	ρ	z_{tsh}		T
cp032294	-.294	-3.672	**	-.297	-2.987	**	101	-.117	-0.867		55
cp032394	-.165	-1.140		-.200	-1.058		28	.055	0.247		20
cp110193	-.623	-8.363	**	-.624	-7.169	**	132	-.114	-0.789		48
cp112493	.260	2.883	**	.289	2.539	**	77	.053	0.359		46
cp120893	-.429	-6.892	**	-.430	-5.585	**	169	-.443	-4.181	**	89

Remarks: ρ denotes the first-order autocorrelation coefficient of transaction price changes; z_{tsh} is the corresponding z-statistic, obtained by multiplying ρ by the square root of the sample size; T is the sample size. ** denotes significance at the 1% level.

Table 3
Serial correlation properties of intra-trading-round accepted asks in five experiments: evidence adjusted for nonstationarity

Experiment	Adjustment Sample z_{tsh}		z_c		z_{ts}		$\hat{\sigma}_{ts}$	Control Sample z_{tsh}		z_c		z_{ts}	$\hat{\sigma}_{ts}$
cp032294	-2.987	**	-1.048		-0.842		4538.1	-0.867		-0.148		-0.628	58.1
cp032394	-1.058		-0.265		-0.957		1008.2	0.247		0.206		0.381	133.3
cp110193	-7.169	**	-2.038	*	-2.340	**	3176.7	-0.789		-0.012		-0.756	4.0
cp112493	2.539	**	0.172		2.538	**	246.9	0.359		0.082		0.316	63.9
cp120893	-5.585	**	-0.864		-1.210		2604.0	-4.181	**	-2.229	*	-1.574	349.2
χ^2	99.2	**	6.1		15.0	**		19.1	**	5.0		3.6	

Remarks: z_{tsh} tests for serial correlation of a homoscedastic series. z_c tests for serial correlation of a nonstationary series. z_{ts} tests for serial correlation of a nonstationary, nonergodic series. χ^2 is obtained by squaring z-statistics and summing across rows. * and ** denote significance at the 5% and 1% level, respectively.

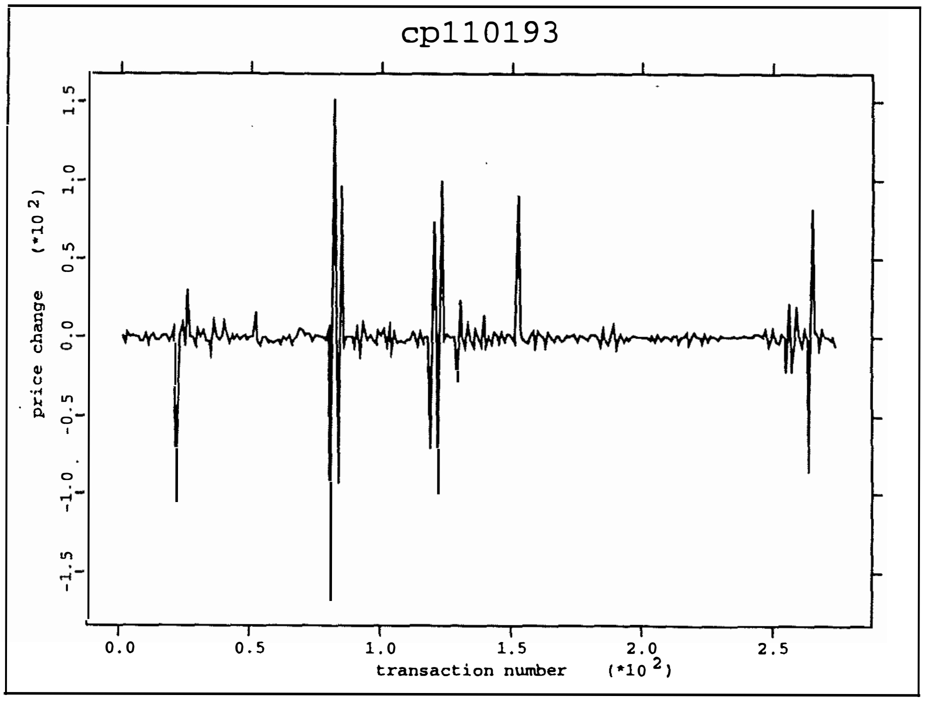

Figure 1: Transaction price changes against transaction time for experiment cp110193.