The Dynamics Of Equity Prices In Fallible Markets*

Peter Bossaerts

January 1998 (first version: April 1995)

*California Institute of Technology. Mailing address: m/c 228-77, California Institute of Technology, Pasadena, CA 91125; Phone (626) 395-4028; e-mail: pbs@rioja.caltech.edu. This paper is a revision of Bossaerts [1996]. The first part came out of a sketchy paper that the author presented at the Université des Sciences Sociales, Toulouse, in April 1995, while he was at CentER, Tilburg University. He is grateful for the many comments from the seminar participants. In addition, the paper benefited through discussions at seminars at Carnegie Mellon University, U.C. Riverside (Economic Theory and Econometrics), University of Minnesota, Washington University at Saint Louis, the 1996 Asset Pricing summer meeting at NBER, the 1996 European Meetings of the Econometric Society, and through comments from and discussions with Oleg Bondarenko, David Easley, Alan Kraus, Guy Laroque, P.C.B. Phillips, Richard Roll and Philippe Weil. Oleg Bondarenko provided able research assistance. Ravi Jagannathan, the Executive Editor, and a referee provided invaluable comments and suggestions. The usual disclaimer applies.

Abstract

In an efficient securities market, prices correctly reflect news about future payoffs. This paper argues that there are two aspects to correctness: (i) correct updating of beliefs from news, (ii) correct prior beliefs. Traditionally, empirical research has implicitly insisted on both. Lucas' rational expectations equilibrium theory also assumes both, explicitly. Nevertheless, rationality requires only the former, but not the latter. This paper develops restrictions on the random behavior of prices of equity-like contracts when (i) is maintained, but the market may have mistaken priors about the likelihood of the bankruptcy state, in violation of (ii). The restrictions are cast in the form of familiar martingale difference results. They do not necessarily restrict returns as traditionally computed, however. Most importantly, the restrictions appear only when the empiricist deliberately imposes a selection bias. In particular, the price histories of securities that are in the money at the terminal date are to be separated from those of securities that end out of the money (i.e., in the bankruptcy state). As a result, this paper also demonstrates that something can be learned about market efficiency from samples subject to survivorship bias or the Peso problem.

JEL Classification : C22, D84, G14.

Keywords: Market Efficiency, Incomplete Information, Biased Priors, Rational Learning, Martingales, Default, Survivorship Bias, Peso Problem.

1 Introduction

Financial markets are populated with human beings. According to the *Chambers 20th Century Thesaurus*, there are many synonyms for the adjective *human*. One of them is: reasonable. This should capture man's affinity to logical reasoning. Over the last thirty years, rationality has been made the cornerstone of the theory of finance. The hypothesis is that securities prices would reflect the rationality of the market participants. The label *efficiency* is used to characterize the outcome. In particular, a market is deemed efficient if *its prices correctly incorporate the available information*. This is the Efficient Markets Hypothesis (EMH; see Fama [1970], [1991]).

The empirical implementation of market efficiency and the subsequent theoretical refinements (in particular, Lucas' dynamic rational expectations equilibrium - Lucas [1972], [1978]) have given an extreme interpretation of the qualifier "correctly" in the above definition of an efficient market: not only is the market supposed to update its beliefs rationally using the available information, its beliefs are assumed to be unbiased at all times. Therefore, the market is endowed with far more rationality than is typically associated with even the most sophisticated of its human participants. Indeed, the *Chambers 20th Century Thesaurus* lists another synonym for the adjective human: fallible. Investors do make mistakes; they are not omniscient. Why, then, wouldn't the market sometimes have mistaking expectations?

In its defense, one must note that EMH leads to a dramatic simplification of empirical research. Indeed, as long as the world is stationary and ergodic, actual outcomes provide good estimates of the market's ex-ante expectations. As a consequence, a large body of evidence on the validity of the theory has emerged. Unfortunately, in violation of one of the main implications of EMH, returns in excess of the riskfree rate have often been found to be predictable beyond acceptable compensation for risk,[1] or in a way that was opposite to risk.[2]

Humans are fallible. In particular, they do not always know the true frequency of events that determine securities prices. This certainly occurs in new or rare events, such as initial public offerings in new industries, new political arrangements (Hong Kong's return to China, Russia's change to a democracy), or economic and monetary arrangements (the European common currency), exceptional historical developments (Gulf war), etc.

The purpose of this paper is to study securities price behavior in a market that is modeled in the image of its participants, as a reasonable but fallible human. The market puts a (potentially biased) prior on the aspects (parameters) of the environment it is uncertain about, and learns using the rules of conditional probability (Bayes' law).

[1] E.g., Mehra and Prescott [1985], De Bondt and Thaler [1985], Jegadeesh [1990].

[2] E.g., Dunn and Singleton [1986], Fama and French [1992].

The novelty is that the market's prior is allowed to be arbitrary. At first, this would seem to deprive securities prices of any meaningful restriction, even if attention is restricted to risk neutrality. Unless, of course, one makes the implausible assumption that the biases in market beliefs are known. In other words, it would be reasonable to conjecture that any price history could be explained in terms of some bias in initial beliefs. The examples in Bossaerts [1995] certainly seem to lend support to this conjecture.

The conjecture is wrong. This paper demonstrates that the dynamic behavior of the prices of equity-like contracts is subject to easily verifiable and powerful restrictions. The "equity-like contracts" are finite-lived securities which have a simple final payoff structure: in one state, the contract pays a random quantity; in another one, to be referred to as the bankruptcy state, it pays zero (or some known amount). This obviously describes the payoff structure of common stock, call and put options, as well as simple digital options ("Arrow-Debreu securities"). Many other financial contracts are known to be combinations of these securities, or could be imagined as such.

Initially, the market does not know the state of the world. Most importantly, the market may have incorrect (biased) beliefs about the frequency of occurrence of the bankruptcy state. As a sophisticated, rational human, however, the market updates its beliefs as evidence ("signals") about the true state emerges, using Bayes' law. Still, it will be assumed that the market has correct beliefs about the distribution of the signals and the final payoff conditional on the state of the world. For instance, in the case of common stock, the market is supposed to correctly anticipate the distribution of terminal values of shares in the company, conditional on no default. Therefore, the analysis of this paper is limited to a small, yet important, deviation from the level of rationality assumed in the standard theory of efficient markets.

One could refer to the extension of market efficiency proposed here as a *theory of Efficiently Learning Markets* (ELM), a terminology that this paper will adopt. Borrowing from game theory (Harsanyi [1967]), it could also be called a theory of efficient markets under incomplete information. This terminology may be somewhat confusing, because it resembles market efficiency under asymmetric information.[3]

There is a relationship with recent attempts to model "overconfidence" and other supposedly

[3]The equilibrium concept underlying the theory of efficient markets is Lucas' dynamic rational expectations equilibrium. The equilibrium concept that could justify the modeling of this paper is somewhere between a temporary equilibrium (see Grandmont [1977] and references cited there) and Lucas' equilibrium. More specifically, Bray and Kreps [1987]'s rational learning equilibrium could form a starting point for developing the general equilibrium aspects of the concept of market efficiency proposed here. To reflect the importance of agreement among investors, earlier versions of this paper talked about a consistent beliefs equilibrium. See Bossaerts [1996].

irrational phenomena in financial markets (e.g., Daniel, Hirshleifer and Subrahmanyam [1997]; Cecchetti, Lam and Mark [1997]). These attempts, however, put "irrationality" both in the priors (mistaken expectations) and in the learning (ad hoc learning rules). The present paper demonstrates that a theory that attributes "irrationality" entirely to mistaken expectations generates a rich set of results that inherit the attractive analytics of the theory of efficient markets, and that allow one to readily incorporate existing asset pricing theory.

In fact, this description of the present paper has to be qualified. Indeed, it cannot be emphasized enough that *to have mistaken expectations is no sign of irrationality*, for rationality is purely a property of learning, and not of beliefs.[4] This point is also implicit in the theory of rational decision making under uncertainty. Savage [1954], for instance, shows that strong rational choice axioms merely induce a mathematical representation of preferences in terms of expected utility over subjective beliefs, and, therefore, one that does not involve correct (unbiased) beliefs.[5]

The restrictions on securities prices are cast in the familiar language of martingales. In contrast with the theory of efficient markets, however, they may involve returns measured as changes in prices divided by *end-of-period* prices, and weighted appropriately. Most importantly, they emerge only as a consequence of *deliberate conditioning on the outcome*. Indeed, the simplest restriction (which this paper focuses on) holds only for securities that expired "in the money." This obviously has implications far beyond the attempt to relax EMH. Financial data often come with the very selection bias that has to be deliberately applied in order to test for ELM. Indeed, many datasets suffer from *survivorship bias*, whereby one observes only price histories of securities that were "in the money" at the final date (see, e.g., Brown, Goetzmann and Ross [1995]).[6]

One can state this result differently: it is shown that a weak form of market efficiency can be tested even on a sample subject to selection bias (survivorship bias). One cannot test that initial beliefs are unbiased, however. So, EMH as it is traditionally understood cannot be tested.

[4]This bold statement may be surprising, especially to those who study the rational expectations equilibrium originally defined in Muth [1961]. In it, the market correctly forecasts prices because it knows that they emerge from equilibration of demand and supply. Forecast mistakes would reflect lack of understanding of economic theory. Nevertheless, correctness of the rational expectations forecasts depends crucially on correctness of the *belief* that every agent in the economy is rational. If enough agents behave suboptimally, the rational expectations forecasts will be wrong. Hence, forecasts that deviate from the rational expectations belief are *not* irrational. They merely express the belief that there are non-optimizing agents, and this could be right.

[5]A referee reminded me of Savage's views on this matter.

[6]A related problem is the *Peso problem*, which occurs when some (random) event did not happen in a history, yet was anticipated in the price setting (see, e.g., Bekaert, Hodrick and Marshall [1995]).

But one can test that the market is efficient otherwise. In particular, one can verify whether its learning is rational, i.e., that it does not under-react or over-react relative to its own beliefs. Hence, ELM can be tested.

Some results may lack intuition. Their generality certainly challenges imagination. Therefore, a simple case will be analyzed first. Also, ideas will be developed by means of empirical and numerical examples. This way, the theory will be built up, and the reader will be able to gain perspective as to its content, scope, and limitations.[7]

Here is how the paper proceeds. Digital options will be studied first. These are options that pay one dollar in one future state, and zero dollars otherwise. The analysis will start with an empirical example, using data from the internet-based Iowa Experimental Market (IEM). Biases in the market's beliefs will not (yet) be the main concern. Rather, the example will be used to illustrate survivorship bias. In particular, we will study the impact on tests of market efficiency when only price histories of "winning" options (options that expired in-the-money) are used in the inference. The empirical exercise will be complemented with a numerical example. At that point, a correction to the traditional return measure will be suggested. It will eliminate the effect of survivorship bias. This will be demonstrated with the numerical example, and the resulting test will be implemented on the IEM data.

We then turn to general equity-like payoffs. We follow the same chronology. Survivorship bias is illustrated, this time using prices of S&P500 index call options that expired in-the-money. A numerical example accompanies the analysis. A correction to the standard return measure is advanced. It is verified on the numerical data, and implemented on the index call option dataset.

We do have the full dataset of S&P500 index call options, including those that expired out-of-the-money. Therefore, standard market efficiency (EMH) can be tested. We will do this and document violations. This may seem puzzling, because our correction for survivorship bias on the returns for winning index call options appeared to work. With the numerical example, however, we will demonstrate that our correction to the return measure eliminates survivorship bias no matter what the market's prior is about the chances that the call expires in-the-money. The results rely only on correctness of learning in the marketplace, i.e., on the usage of Bayes' law. Therefore, the survivorship-bias adjusted return can be used to test the weaker ELM, and not the stronger EMH.

We then proceed to stating the main formal result (Theorem 2). Still, this is not presented in its full generality. Risk neutrality and zero discounting will be assumed, as in the numerical

[7] The plan follows the suggestions of the Executive Editor.

examples. Extensions of the theory that accommodate risk aversion and/or random discounting are delegated to the Appendix, where all the formal proofs are collected as well.

2 Digital Options

Digital options are securities that pay one dollar in one future state, and zero in all others. While they are the building blocks of modern asset pricing theory (where they are referred to as Arrow-Debreu securities), they are not widely traded. One exception is the internet-based Iowa Experimental Market (IEM), to which we now turn to illustrate the problem of survivorship bias.

The University of Iowa organizes an electronic market in "winner-take-all" contracts (digital options). A few of those contracts derive their payoff from stock price changes of firms in the computer industry. Each month, a new set of contracts is offered. Contract liquidation values are determined by changes in closing prices of the underlying stock measured from the third Friday of one month to the third Friday of the next month. Trade starts on Monday following the third Friday of the month. In our illustration, we will focus on the digital options that are written on common stock of Microsoft. A comprehensive study of all IEM financial winner-take-all contracts can be found in Bondarenko and Bossaerts [1997].

Two Microsoft digital options 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 option written on Microsoft and traded at the CBOE. We will focus on daily closing prices, which are defined to be the last transaction price before midnight, or, if no transaction took place, the previous closing price. Figure 1 plots price paths for the two Microsoft securities over four sample months. The prices of the two (complementary) digital options do not add up to one because of some serious nonsynchrony. The "Low" contract expired out-of-the-money in each of the four months.

2.1 Illustrating Survivorship Bias Using Data From The IEM

Survivorship bias is the effect on sample average returns caused by absence of price data for securities that did not perform well. In the case of digital options, survivorship bias could be induced by (deliberate or unintentional) deletion of price histories of options that expired out-of-the-money. With only data on successful digital options, it is clear that the potential average return is over-estimated. If the market is risk-neutral, applies zero discount rates, and has unbiased beliefs (in accordance with the EMH), then ex-ante expected returns would be zero. When computed on the basis of price histories of winning securities only, however, ex-post returns will be positive on average.

We can illustrate survivorship bias using the IEM data. Across the sixteen months in the dataset, the average of the (time series) mean daily return on Microsoft High is 8.1%. With a standard error of 4.2%, this is marginally significant at the 5% level.[8] This does indicate some evidence against EMH, which would require the average return to be equal to zero (assuming risk neutrality and zero discounting). We are primarily interested, however, in illustrating survivorship bias. So, let us compare this to the average return on winning IEM digitals. Each month, one of the two Microsoft digitals expires in-the-money, because they are complementary. So, we have sixteen monthly series of daily returns on winning IEM digital options. Across those sixteen months, the average of their mean daily return is 9.7%. With a standard error of 3.7%, this is highly significant. And it is much higher that the average mean daily return on Microsoft High, because of the selection bias.

It is possible to illustrate survivorship bias with a numerical example.

2.2 A Numerical Example

Consider a world in which a state variable, θ, can take two values: $\overline{\theta}$ or $\underline{\theta}$. A security is traded and its market clears at two discrete points in time, indexed $t = 1, 2$. The security is a digital option. It expires at $t = 3$, when it pays \$1 if $\theta = \overline{\theta}$ and zero otherwise. The market is risk-neutral and the interest rate is zero. It does not know θ before the end of the second period, i.e., before time $t = 3$. At time $t = 2$, however, it receives a signal s, which it uses to update its beliefs about the value of θ. The equilibrium price at time t is denoted p_t.

The market is efficient, i.e., EMH holds. This means that prices are set according to the true probability measure P with which signals and states are actually drawn. In particular, the market uses this probability measure to infer from the signals whether $\theta = \overline{\theta}$ or not.

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

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

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

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

[8] The significance levels in all tests of the paper are based on the t-distribution. There is no reason to suspect violations to the cross-sectional independence that would be required to validate the t-distribution, since cross-sections of non-overlapping time series are studied throughout.

The signal at $t = 2$, s, is determined as follows. Let $l(s|\theta)$ denote the likelihood of s given θ. Assume:

$$\begin{aligned} l(1|\overline{\theta}) &= \frac{1}{2} \\ l(1|\underline{\theta}) &= \frac{1}{3} \\ l(0|\overline{\theta}) &= \frac{1}{2} \\ l(0|\underline{\theta}) &= \frac{2}{3} \end{aligned}$$

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

$$\begin{aligned} p_2 &= P\{\theta = \overline{\theta}|s = 1\} && (2) \\ &= \frac{l(1|\overline{\theta})p_1}{l(1|\overline{\theta})p_1 + l(1|\underline{\theta})(1 - p_1)} \\ &= \frac{\frac{1}{2}\frac{1}{4}}{\frac{1}{2}\frac{1}{4} + \frac{1}{3}\frac{3}{4}} \\ &= \frac{1}{3}; && (3) \end{aligned}$$

if $s = 0$:

$$\begin{aligned} p_2 &= P\{\theta = \overline{\theta}|s = 0\} && (4) \\ &= \frac{l(0|\overline{\theta})p_1}{l(0|\overline{\theta})p_1 + l(0|\underline{\theta})(1 - p_1)} \\ &= \frac{\frac{1}{2}\frac{1}{4}}{\frac{1}{2}\frac{1}{4} + \frac{2}{3}\frac{3}{4}} \\ &= \frac{1}{5}. && (5) \end{aligned}$$

In the above, Bayes' law is used to compute conditional probabilities such as $P\{\theta = \overline{\theta}|s = 0\}$. In other words, the market is assumed to update its priors rationally. While this assumption is standard in finance, and an integral part of EMH, it is crucial to obtain the results that we are about to discuss.

The above data can be used to study the impact of survivorship bias. First, consider the expected time-2 price:

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

Consequently:

Result 1:

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

where

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

the traditional return measure. Eqn. (6) states that the return is zero on average. More generally, the return sequence will form a martingale difference sequence. This, of course, is a result that goes back to Samuelson [1965]. It is the mathematical representation of EMH (under risk neutrality and zero discounting).

Now introduce a selection bias and condition on a favorable final payoff, i.e., on $\theta = \overline{\theta}$. In other words, re-compute the average return, but only for cases where the digital option expired in-the-money. The difference between the conditional and the unconditional mean returns will give an indication of survivorship bias, i.e., the bias in average return caused by deleting price histories of losing securities.

Compute:

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

Hence:

Result 2:

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

In words: there is an upward bias in the mean return.

2.3 A Correction For Survivorship Bias

There happens to be a very simple correction for the above survivorship bias. Consider the following. Compute the conditional expectation of the *inverse* of the price at time 2:

$$\begin{aligned} E[\frac{1}{p_2}|\overline{\theta}] &= 3P\{s=1|\overline{\theta}\} + 5P\{s=0|\overline{\theta}\} \\ &= 3\frac{1}{2} + 5\frac{1}{2} \\ &= 4 \\ &= \frac{1}{p_1}. \end{aligned}$$

Hence, defining

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

we conclude the following.

Result 3:

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

In words: modifying the return by taking as basis not the past price but the *future price* makes it again zero on average.

The return measure that one obtains using the future price as basis will be referred to as the *modified return* and x_t will be used as symbol.

To a certain extent, Result 3 is intuitive. Consider price histories of digitals that eventually expire in-the-money. One expects these to reveal a distinct positive trend, reflecting on-balance positive news about the final payoff. Result 2 confirms this. To bring the average return down to zero, one ought to make positive returns smaller and negative returns larger (in absolute value), by multiplying the return with an appropriate factor. The most straightforward candidate is the ratio of today's price over tomorrow's price: it is smaller than one when the return is positive, and larger than one otherwise. It turns out that this factor does the trick. The performance measure that results from multiplying the traditional return with this factor is precisely our modified return. Figure 2 illustrates this idea.

The result in (9) is very general. It really only relies on the assumption that the market uses Bayes' law to update its beliefs about the chances that the digital option expires in-the-money. It is too early to sharply delineate the scope of this result. Things will become more precise as we go along, however. At this point, it is more useful to observe how the result works when implemented on the IEM data.

2.4 Re-Visiting The IEM Data

Table 1 also displays the average of the mean modified daily return across the sixteen months in the IEM dataset. In each month, the digital option that expired in-the-money is determined, modified returns are computed from its price series, and their (time series) mean is computed. Subsequently, the mean modified returns across the sixteen months are averaged, and the standard error is computed in the usual fashion.

Table 1 documents that the average mean modified daily return on winning IEM digitals is -3.2%. With a standard error of 2.2%, this number is insignificant. The insignificant -3.2% should be contrasted with the highly significant 9.7% average mean traditional daily return on winning

IEM digitals. Apparently, adjusting the return along the lines of the previous subsection does indeed eliminate the selection bias that was caused by our considering only winning securities.

Because the modified return filters selection biases from price series of digital options that expire in-the-money, the reader may be inclined to conclude that EMH can now be tested on samples subject to survivorship bias. That is, EMH is verifiable *even if* a biased sample of only winning securities is all one has. Such a conclusion will have to be qualified. We will do so after introducing the survivorship bias correction for returns on general equity-like securities. At this point, we also need to clarify something about the impact of perfectly revealing signals.

2.5 Back To The Numerical Example: The Second Period

In the first period of our numerical example, signals were not perfectly revealing. That is, they did not unambiguously signal whether $\theta = \overline{\theta}$. The data for the second period ($t = 2$ to $t = 3$) allow us to discuss the effect of perfect revelation.

At $t = 3$, the true state is revealed and the security pays off. There is no real trading, but, if there were, prices would obviously equal the announced payoffs. That is, $p_3 = 1$ if $\theta = \overline{\theta}$, and $p_3 = 0$ if $\theta = \underline{\theta}$.

Let us assume that $s = 1$ in the first period. Hence, the time-2 price equals 1/3, and the return over the second period, r_3, is either

$$r_3 = \frac{p_3 - p_2}{p_2} = \frac{1 - \frac{1}{3}}{\frac{1}{3}} = 2$$

if $\theta = \overline{\theta}$, or

$$r_3 = \frac{p_3 - p_2}{p_2} = \frac{0 - \frac{1}{3}}{\frac{1}{3}} = -1$$

otherwise. On average:

$$E[r_3|s = 1] = 2P\{\theta = \overline{\theta}|s = 1\} - 1P\{\theta = \underline{\theta}|s = 1\} = 2\frac{1}{3} - \frac{2}{3} = 0,$$

in accordance with EMH and Result 1 above.

Let us introduce a selection bias, and investigate whether the *modified* return eliminates it. The modified return equals

$$x_3 = \frac{p_3 - p_2}{p_3}.$$

So,

$$x_3 = \frac{1 - \frac{1}{3}}{1} = \frac{2}{3}$$

if $\theta = \overline{\theta}$, and

$$x_3 = \frac{0 - \frac{1}{3}}{0} = -\infty$$

otherwise. Computing the expectation conditional on the option's expiring in-the-money, i.e., $\theta = \overline{\theta}$, one obtains:

$$
\begin{aligned}
& E[x_3 | s = 1, \theta = \overline{\theta}] \\
= \; & \frac{2}{3} P\{\theta = \overline{\theta} | s = 1, \theta = \overline{\theta}\} - \infty P\{\theta = \underline{\theta} | s = 1, \theta = \overline{\theta}\} \\
= \; & \frac{2}{3} - \infty 0, \qquad (10)
\end{aligned}
$$

which is indeterminate.

Using limit arguments, one can resolve the indeterminacy in (10) and restore Result 3. See Bossaerts [1997]. Such arguments are purely mathematical, however. From an empirical point of view, it is only of relevance to ask what the empiricist will actually see. Conditional on $\theta = \overline{\theta}$ (and $s = 1$), the empiricist will *only* observe: $p_2 = 1/3$ and $p_3 = 1$. So, the average modified return across replications of histories where the digital option expired in-the-money (and $s = 1$) is:

$$x_3 = \frac{1 - \frac{1}{3}}{1} = \frac{2}{3},$$

in violation of Result 3.

Earlier, this paper proposed a simple modification to the traditional return measure that eliminates biases which result when only price histories of winning digital options are investigated. From the example we have just discussed, it is clear that this adjustment will not work if there is a chance that signals fully reveal whether θ equals $\overline{\theta}$. Bossaerts [1997] discusses the impact if this chance is only tiny.

3 General Equity-Like Payoffs

The limited liability of equity gives it a peculiar payoff structure: either a positive, random amount is received, or zero. The latter outcome is usually referred to as "default" or "bankruptcy." So, equity is really an extension of the digital option discussed in the previous section. When equity is "in the money," however, the holder receives a random, positive pay, as opposed to a fixed $1.

Many other securities have equity-like payoff patterns. Among those are the quintessential options, the put and the call. Call options pay the maximum of the difference between the value of the underlying asset and its strike price, or zero. As mentioned in the Introduction, there are many other securities that have equity-like payoff patterns. Because of their simplicity, however, call options will be used to illustrate the theoretical developments of this paper.[9]

[9]The reader may ask whether common stock (equity) really belongs in the category of securities that is being

In particular, we will investigate European S&P500 index call options, traded on the CBOE (symbol: SPX). Four-week series of daily prices will be studied, covering the years 1991-1995. At the start of each series, the option was at the money and five weeks from expiration. The prices were estimated from closing prices of actually traded put and call options, using the smoothing technique developed in Bondarenko [1997] (constrained convex least squares regression). The latter filters the data for bid-ask bounce and violations of simple arbitrage bounds.

In order to get an idea of the nature of the time series, Figure 3 plots 4 four-week histories, covering the contracts expiring in 1/93, 2/93, 3/93 and 4/93. The index call option matured in-the-money in the months 1/93 and 3/93. It expired worthless in the other two months.

Let us proceed as in the previous subsection, and first document the effect of imposing a deliberate survivorship bias.

3.1 Illustrating Survivorship Bias Using S&P500 Index Call Option Prices

Table 2 displays the average of the daily mean return across the 58 four-week periods in the sample. Returns were computed after inflating each beginning-of-period price with the price of a one-day, riskfree pure-discount bond; the three-month Treasury bill rate was used to proxy for the riskfree rate. The average equals -0.8%, which, with a standard error of 0.7%, is insignificant.

So, the evidence in Table 2 is consistent with EMH: one cannot reject that average returns (adjusted for the riskfree rate) are zero on average. This evidence is based on an unbiased sample: both "winning" and "losing" call options are included, in the proportion that was dictated by Nature.

To illustrate the impact of survivorship bias, Table 2 also reports the average daily mean return across the 40 four-week histories when the index call option expired in-the-money. In other words, it reports an estimate of the daily mean return that is affected by the usual survivorship

studied here, pointing to its lacking a definite maturity date. Notice, however, that the theory continues to hold if we specify the maturity date of common stock to be a specific date in the future (e.g., one year from now), setting V equal to the stock price that would obtain at that point in time. Of course, there is the issue of prior revelation of bankruptcy. As illustrated in Section 2, securities prices will not be observed to behave as the theory requires if there is a chance that bankruptcy is announced before maturity. Hence, when implementing the tests in the context of equity pricing, there is an effective limit on how far the horizon (maturity date) can be extended into the future. Since the theory takes bankruptcy to be liquidation (with zero payoff), one has to set the maturity date such that definite news about liquidation is unlikely to emerge prior to the chosen date. Nevertheless, a small, positive probability of full revelation of bankruptcy can be tolerated, as the simulations in Bossaerts [1996] illustrate.

bias. As expected, it is positive (1.3%), and, with a standard error of 0.6%, significantly so. Hence, computing average returns only on the basis of winning histories does induce a positive bias in the estimate of the average return.

Before introducing an adjustment of the return measure that corrects for survivorship bias, let us first discuss a numerical example. It will be used to demonstrate how the return adjustment works.

3.2 A Numerical Example

We will build on the numerical example from the previous section. We keep the same state and signal structure, but we now add the complication that our option pays a random amount V^+ whenever $\theta = \overline{\theta}$. In particular, we specify:

$$V^+ = \begin{cases} 2 & \text{if } s = 1; \\ 1 & \text{if } s = 0. \end{cases}$$

By linking the signal at time $t = 2$ and the payoff at $t = 3$, we introduce realism in the example. Indeed, signals about the likelihood that a security expires in-the-money usually also reveal information about what the final payoff will be in the case the security indeed expires in-the-money. The index call option from the previous subsection is a good example: the main signal for the likelihood of this option's expiring in-the-money is the level of the S&P500 index; at the same time, the S&P500 index also reveals a lot about how far the option will be in the money if it will be at all.

The first price, p_1, will now be:

$$\begin{aligned} p_1 &= E[V^+ 1_{\{\theta=\overline{\theta}\}}] \\ &= P\{\theta = \overline{\theta}\}\left(2P\{s = 1|\overline{\theta}\} + 1P\{s = 0|\overline{\theta}\}\right) \\ &= \frac{1}{4}\left(2\frac{1}{2} + 1\frac{1}{2}\right) \\ &= \frac{3}{8}. \end{aligned}$$

Similarly, prices at $t = 2$ change. If $s = 1$,

$$\begin{aligned} p_2 &= 2P\{\theta = \overline{\theta}|s = 1\} \\ &= \frac{2}{3}; \end{aligned}$$

and, if $s = 0$,

$$\begin{aligned} p_2 &= 1P\{\theta = \overline{\theta}|s = 0\} \\ &= \frac{1}{5}. \end{aligned}$$

This allows us to re-derive Result 1.

$$\begin{aligned} E[p_2] &= \frac{2}{3}P\{s=1\} + \frac{1}{5}P\{s=0\} \\ &= \frac{2}{3}\frac{3}{8} + \frac{1}{5}\frac{5}{8} \\ &= \frac{3}{8}. \end{aligned}$$

Hence,

<u>Result 4:</u>

$$E[r_2] = E[\frac{p_2 - p_1}{p_1}] = 0. \tag{11}$$

Again, this is the translation into mathematics of what EMH is about.

To determine the effect of survivorship bias, condition on the state $\theta = \overline{\theta}$:

$$\begin{aligned} E[p_2|\overline{\theta}] &= \frac{2}{3}P\{s=1|\overline{\theta}\} + \frac{1}{5}P\{s=0|\overline{\theta}\} \\ &= \frac{2}{3}\frac{1}{2} + \frac{1}{5}\frac{1}{2} \\ &= \frac{104}{240}. \end{aligned}$$

This is higher than p_1, which equals 3/8, i.e., 90/240. Consequently,

<u>Result 5:</u>

$$E[r_2|\overline{\theta}] > 0.$$

Survivorship of winning securities induces a clear upward bias in average returns.

3.3 Again, A Correction For Survivorship Bias

As with the digital option, there is a simple correction of the traditional return measure that will purge it of its bias. The adjustment is slightly different (in fact, the adjustment for digital options is nested).

Consider the average inverse p_2. We will need a *weighted average*, where the weights are determined by the final payoff of the security, V^+:

$$\begin{aligned} E[\frac{1}{p_2}V^+|\overline{\theta}] &= (\frac{3}{2})2P\{s=1|\overline{\theta}\} + (5)1P\{s=0|\overline{\theta}\} \\ &= \frac{3}{2}2\frac{1}{2} + 5\frac{1}{2} \\ &= 4. \end{aligned} \tag{12}$$

Likewise, compute the conditional average final payoff:

$$E[V^+|\overline{\theta}] = 2\frac{1}{2} + 1\frac{1}{2} = \frac{3}{2}.$$

This, together with (12), implies the following.

Result 6:

$$\begin{aligned} E[x_2 V^+|\overline{\theta}] &= E[\frac{p_2 - p_1}{p_2} V^+|\overline{\theta}] \\ &= E[V^+ - p_1 \frac{1}{p_2} V^+|\overline{\theta}] \\ &= \frac{3}{2} - \frac{3}{8} 4 \\ &= 0. \end{aligned} \tag{13}$$

Our standard adjustment to the return measure, together with a weighting using the final payoff, eliminates the survivorship bias entirely. Notice that Result 3 (for digital options) is nested in Result 6. Indeed, for digital options, $V^+ = 1$ always, the weighting becomes trivial, and the average reduces to (9).

Again, this result is very general. Let us look at how it works in the context of CBOE S&P500 index call options.

3.4 Back To The CBOE Options Data

Table 2 also displays the weighted average of the mean daily modified return across the 40 four-week periods that the S&P500 index call option expired in-the-money. For the reader who is unsure about how this weighted modified return was computed, let C_t^i denote the closing price on day t of the ith four-week history ($i = 1, \ldots, 40$, $t = 1, \ldots, T^i$). Let r_t^i be the corresponding riskfree rate (expressed as percentage per day; as mentioned before, the three-month Treasury bill rate was used as proxy). Modified return (i, t) was computed as:

$$x_t^i = \frac{C_t^i - C_{t-1}^i (1 + r_t^i)}{C_t^i}.$$

The weight $V^{+,i}$ was computed from the last closing price of the option:

$$V^{+,i} = C_{T^i+5}^i.$$

(At the beginning of each four-week period, options have a maturity of five weeks; therefore, expiration occurs five trading days after the end of each history, i.e., at $T^i + 5$, and the option's final payoff can be computed from the closing price for that day.) Then, the weighted average modified return was calculated using the following formula:

$$\frac{1}{40} \sum_{i=1}^{40} V^{+,i} \left(\frac{1}{T^i} \sum_{t=1}^{T^i} x_t^i \right).$$

The cross-section of the weighted mean daily modified return formed the basis for calculating standard errors.

Table 2 reports that the weighted average of the mean daily modified return on winning S&P500 index call options equals 0.4%. With a standard error of 11.4%, this is insignificant at any reasonable level. Our adjustment to the standard return measure, together with the peculiar weighting scheme, appear to have eliminated the bias that we introduced by looking only at 40 four-week price histories for options that expired in-the-money.

To gain perspective, it should be emphasized that the lack of significance of the weighted average modified return in Table 2 is in no way caused by our weighting with a variable with substantial range, namely V^+. To illustrate this, Table 2 also displays the weighted average of the daily mean *traditional* return, which, like its unweighted counterpart, should also be affected by our deliberately imposed selection bias. Table 2 reports that it equals a sizeable 35.3%, which is highly significant in view of the standard error of 12.7%.

Still, the issue of power is interesting. Because of lack of space, we cannot address it here. The interested reader is referred to Bossaerts [1997], as well as Bossaerts and Hillion [1997]. The latter includes a Monte Carlo exercise and reports excellent power properties against alternative hypotheses with an over-reacting market.

3.5 Projections

Just like Results 1 and 4 (Equations (6) and (11)) are known to be extendible to become *martingale difference restrictions*, Results 3 and 4 (Equations (9) and (13)) generalize. This would mean that modified returns or weighted modified returns must not be predictable from past information. Any projection of modified or weighted modified returns onto past information must lead to insignificant coefficients (except for the usual type I error).

Let us investigate how projections would work on the index options data. Figure 4 plots weighted daily modified returns of winning options against lagged information. The lagged price level was chosen as predictor, because it was one of the few variables that did reveal some anomaly.[10] The solid line through the scatter plot is a kernel estimate of the regression function. For options that are priced below $5, there seems to be some indication that the lagged price level could be used to predict weighted daily modified returns, with lower prices indicating lower weighted modified returns. This would violate our restriction. For prices beyond $5, however, there is clearly no discernible pattern, confirming the restriction. One could suspect that bid-ask

[10] Also, the lagged price is a popular predictor for asset returns, ever since Keim and Stambaugh [1986] discovered that it could be used to forecast stock returns.

bounce is behind the rejection for low-priced options, but this is not the appropriate place to further investigate such a possibility. We will interpret the evidence from Figure 4 to be roughly in line with the martingale-difference extension of Result 6.

Instead, we come back to a question that was raised at the end of the previous section. We can adjust returns to offset survivorship biases. Does that imply that EMH can be tested on biased samples? It is time to address this question.

4 Biases In Beliefs

In EMH, the market has unbiased beliefs. Among other things, it is supposed to know *ex ante* the true frequency with which options expire in-the-money. As far as the CBOE data is concerned, we have found no clear violations of EMH. A closer look, however, does reveal some striking violations.

4.1 Re-Visiting S&P500 Index Option Prices

Figure 5 plots daily changes in S&P500 index call option prices against lagged call price levels across *all* histories. That is, all daily price changes were included, no matter whether the option eventually expired in-the-money (winners) or out-of-the-money (losers). Price changes were preferred over returns: the plot with the latter on the Y axis was harder to interpret because of some obvious heteroscedasticity.

Figure 5 clearly displays strong evidence against EMH. There is a pronounced, positive, linear relationship between price level and subsequent price changes. This relationship is uniform across all price levels. The intercept of the OLS projection depicted in Figure 5 is -0.894. Its standard error is only 0.114. Likewise, the slope coefficient is 0.097, with a standard error of only 0.010.

Figure 5 should be contrasted with Figure 4. The latter revealed some weak evidence against the proposition that weighted modified returns should not be predictable, but only for low-priced options. In contrast, the evidence from the former is sharp and convincing: option price changes are predictable across the board, and, hence, EMH has to be rejected.

This empirical example suggests that tests based on weighted modified returns of securities that expire in-the-money must be validating a weaker theory than EMH. So, the impression that one could use weighted modified returns to test EMH in biased samples seems to be wrong. Of course, there may be other explanations for the discrepancies between Figures 4 and 5, such as power, market microstructure effects, skewness, etc. But one should at least entertain the possibility that weighted modified returns of winners can be used only to test a weaker restriction on beliefs than EMH imposes.

We will demonstrate this by means of our numerical example. We will impose a bias on the market's beliefs, and prove that Result 6 continues to hold. In other words, weighted average modified returns remain zero even with biases in the market's prior about the chances of the security's expiring worthless (default).

4.2 Changing The Market's Beliefs In The Numerical Example

Let us change the market's prior about the event $\{\theta = \overline{\theta}\}$, from $P\{\theta = \overline{\theta}\} = 1/4$, to $P^*\{\theta = \overline{\theta}\} = 1/2$. From now on, P^* will indicate the market's beliefs, in order to clearly distinguish them from the "true" probabilities, which remain the same as before, and which will still be indicated with P. To distinguish whether expectations are based on P^* or on P, we should use the notation E^* and E, respectively. Notice that the market is too optimistic, in violation of EMH.

It should be emphasized that we only change the prior about the event $\{\theta = \overline{\theta}\}$. We do not change the market's beliefs otherwise. In particular, the market continues to hold *correct conditional beliefs*. For instance,

$$P^*\{s = 1|\overline{\theta}\} = P\{s = 1|\overline{\theta}\} = \frac{1}{2}.$$

The first price, p_1, will change, to:

$$\begin{aligned} p_1 &= E^*[V^+ 1_{\{\theta=\overline{\theta}\}}] \\ &= P^*\{\theta = \overline{\theta}\} \left(2P^*\{s = 1|\overline{\theta}\} + 1P^*\{s = 0|\overline{\theta}\}\right) \\ &= \frac{1}{2}\left(2\frac{1}{2} + 1\frac{1}{2}\right) \\ &= \frac{3}{4}. \end{aligned}$$

Similarly, prices at $t = 2$ change. If $s = 1$,

$$\begin{aligned} p_2 &= 2P^*\{\theta = \overline{\theta}|s = 1\} \\ &= \frac{l(1|\overline{\theta})P^*\{\theta = \overline{\theta}\}}{l(1|\overline{\theta})P^*\{\theta = \overline{\theta}\} + l(1|\underline{\theta})P^*\{\theta = \underline{\theta}\}} \\ &= \frac{\frac{1}{2}\frac{1}{2}}{\frac{1}{2}\frac{1}{2} + \frac{1}{3}\frac{1}{2}} \\ &= \frac{6}{5}; \end{aligned}$$

and, if $s = 0$,

$$p_2 = 1P^*\{\theta = \overline{\theta}|s = 0\}$$

$$
\begin{aligned}
&= \frac{l(0|\overline{\theta})P^*\{\theta = \overline{\theta}\}}{l(0|\overline{\theta})P^*\{\theta = \overline{\theta}\} + l(0|\underline{\theta})P^*\{\theta = \underline{\theta}\}} \\
&= \frac{\frac{1}{2}\frac{1}{2}}{\frac{1}{2}\frac{1}{2} + \frac{2}{3}\frac{1}{2}} \\
&= \frac{3}{7}.
\end{aligned}
$$

Because EMH is violated, Result 1 must be violated as well. Indeed:

$$
\begin{aligned}
E[p_2] &= \frac{6}{5}P\{s = 1\} + \frac{3}{7}P\{s = 0\} \\
&= \frac{6}{5}\frac{3}{8} + \frac{3}{7}\frac{5}{8} \\
&= \frac{201}{280}.
\end{aligned}
$$

Since $p_1 = 3/4 = 210/280$,

$$
E[r_2] = E[\frac{p_2 - p_1}{p_1}] = \frac{201}{210} - 1 < 0.
$$

The impact of survivorship bias remains the same, however.

$$
\begin{aligned}
E[p_2|\overline{\theta}] &= \frac{6}{5}P\{s = 1|\overline{\theta}\} + \frac{3}{7}P\{s = 0|\overline{\theta}\} \\
&= \frac{6}{5}\frac{1}{2} + \frac{3}{7}\frac{1}{2} \\
&= \frac{114}{140}.
\end{aligned}
$$

This is higher than p_1, which equals 3/4, i.e., 105/140. Consequently, Result 5 is upheld:

$$
E[r_2|\overline{\theta}] > 0.
$$

What about Result 6? First compute the weighted expected inverse price.

$$
\begin{aligned}
E[\frac{1}{p_2}V^+|\overline{\theta}] &= \frac{5}{6}2P\{s = 1|\overline{\theta}\} + \frac{7}{3}1P\{s = 0|\overline{\theta}\} \\
&= \frac{5}{6}2\frac{1}{2} + \frac{7}{3}\frac{1}{2} \\
&= 2.
\end{aligned}
$$

The conditional average final payoff remains the same as before.

$$
E[V^+|\overline{\theta}] = 2\frac{1}{2} + 1\frac{1}{2} = \frac{3}{2}.
$$

All this implies the following.

$$
\begin{aligned}
E[x_2V^+|\overline{\theta}] &= E[\frac{p_2 - p_1}{p_2}V^+|\overline{\theta}] \\
&= E[V^+ - p_1\frac{1}{p_2}V^+|\overline{\theta}] \\
&= \frac{3}{2} - \frac{3}{4}2 \\
&= 0.
\end{aligned}
$$

This confirms Result 6!

So, the validity of Result 6 does not depend on correctness of the market's priors (EMH). Result 6 derives from the other assumptions: correctness of conditional beliefs, and, most importantly, whether the market uses Bayes' law to update its beliefs. The importance of Bayes' law was already alluded to in the numerical example of a digital option. A market with biased priors but which otherwise reacts correctly to new information (uses Bayes' law) has been referred to as an *Efficiently Learning Market* (ELM). Therefore, the weighted average modified return on winning securities can be used to test ELM, and not EMH.

4.3 Discussion Of The Evidence From The CBOE Options Data

How, then, do we interpret the evidence from the CBOE index options data? Because we found predictability of the traditional return in an unbiased sample, but little evidence of predictability of the weighted modified return in a biased sample (only those options that matured in-the-money), we must conclude: (i) EMH is rejected, (ii) the evidence against ELM is weak. This implies that the options market is one which learns correctly, albeit from biased priors. It may not always know correctly what the probability is that an S&P500 index call option expires in-the-money, but it does react correctly to the arrival of new information.

You will remember, however, that we did uncover some slight evidence against Result 6, and, hence, against ELM, in low-priced options. There may be various explanations for this, including bid-ask spreads, or risk premia (remember that we have been assuming risk neutrality throughout)[11]. It is interesting, however, to take the evidence at face value and wonder what type of irrational market would generate a (conditionally) negative weighted average modified return. The paper would become inordinately burdened if we were to try to answer this question here. The analysis can be found, however, in Bossaerts and Hillion [1997]. In particular, it appears that a market that over-reacts to new information (relative to Bayes' law) would generate significantly negative weighted average modified returns.

5 General Result

It is now time to state the main result formally. Here is the framework. Let time be indexed $t = 1, 2, 3, \ldots, T+1$. The possible states of the world are listed in an outcome space $\Omega = \Theta \times \tilde{\Omega}$. Θ is binary, with two elements, $\overline{\theta}$ and $\underline{\theta}$. Generically, the elements of Θ will be referred to as θ.

[11] The evidence against the theory comes mainly from low-priced options. Generally, these are out-of-the-money options, with a high beta. Hence, models such as the CAPM would predict that they carry a large risk premium.

Θ is the basic state space. It determines the payoff on the one security that is of interest to us. This security's payoff is denoted V, and will be realized at time $T+1$. V is a random variable defined on Ω. It relates to Θ as follows:

$$V = \begin{cases} V^{+} & \text{if } \theta = \overline{\theta}; \\ 0 & \text{if } \theta = \underline{\theta}, \end{cases} \tag{14}$$

This is meant to reflect the typical payoff structure on equity and equity-related contracts (such as options). $\underline{\theta}$ can be referred to as the "default state," because our security pays \$0 when it occurs. When $\theta = \overline{\theta}$, the security will mature "in the money."

The flow of information to the marketplace will be represented by a filtration $\{\mathcal{F}_t\}_{t=1}^{T+1}$. $V \in \mathcal{F}_{T+1}$, but, importantly, θ is not in $\mathcal{F}_t$ when $t \leq T$. In other words, the information never fully reveals θ. The importance of this assumption was discussed at the end of the second section.

Nature's drawings generate a probability measure P. In accordance with the factorization of the outcome space, we split P into an unconditional measure that determines how θ is picked, λ, and a conditional measure, P_θ, that determines how $\tilde{\omega}$ (and, hence, V, as well as the information flow $\{\mathcal{F}_t\}_{t=1}^{T+1}$) is drawn, conditional on θ. That is,

$$P = \lambda P_\theta.$$

Not only does the market not know θ before $T+1$, it may also have biased beliefs about how θ is drawn. The market's beliefs will be represented by a probability measure P^*, to be factored into a (subjective) λ^* over Θ, and an objective conditional measure P_θ over $\tilde{\Omega}$. Since we take P_θ instead of any subjective belief, the market is assumed to have *correct conditional beliefs*. Hence:

$$P^* = \lambda^* P_\theta.$$

Because λ and λ^* may differ, EMH may be violated. To distinguish whether expectations are computed over the objective P or over the subjective P^*, the notation E (when P is involved) and E^* (expectations over P^*) will be used.

The market in our security clears at times $t = 1, \ldots, T$. Let p_t denote the time-t equilibrium price. $p_t \in \mathcal{F}_t$. Subject to risk neutrality and zero discounting, p_t is obtained as:

$$p_t = E^*[V|\mathcal{F}_t], \tag{15}$$

where the conditional expectation E^* is computed under P^*, i.e., under the market's subjective beliefs.

If the market's subjective beliefs are correct, i.e., if EMH holds, we obtain the following well-known result, originally derived in Samuelson [1965]. As before, let r_{t+1} denote the return over the period $t, t+1$:

$$r_{t+1} = \frac{p_{t+1} - p_t}{p_t}. \tag{16}$$

Lemma 1 *Under risk neutrality and zero discounting, and if the market is efficient (i.e., EMH holds),*

$$E[r_{t+1}|\mathcal{F}_t] = 0,$$

$t = 1, \ldots, T-1.$

(All proofs are collected in the Appendix.) In words: returns cannot be predicted from past information.

Lemma 1 will obviously not obtain if subjective beliefs (P^*) and actual frequencies (given by P) differ. If the market is too optimistic about chances that the security expires in-the-money, as in the numerical example of Section 4, we expect:

$$E[r_{t+1}] < 0.$$

Absent knowledge about the nature of the bias in the market's priors, we cannot put a definite sign on the expected return (conditional or unconditional). In other words, under ELM, we have no unambiguous results. Of course, we could frankly assume that we know whether the market is always too optimistic or too pessimistic, but that would be implausible. Also, we do want the theory to work when the market's priors change randomly from one history to another. Right now, they are fixed (and equal λ^*), but the proofs in the Appendix allow them to be random.

This discussion reveals that little can be inferred from an unbiased sample of price histories. The restrictions that ELM imposes only become transparent when a deliberate selection bias is imposed. In particular, we have to investigate the price behavior of winning securities separately from those of losing securities. This was also the message from the examples in Sections 2 to 4.

We now introduce such a selection bias: we condition on knowledge that the security expired in-the-money, i.e., that $\theta = \overline{\theta}$. For digital options, the resulting survivorship bias can be signed. Digital options are securities for which $V^+ = 1$. We have the following result.

Theorem 1 *Consider the traditional return r_{t+1} on a digital option. Assume risk neutrality and zero discounting. If the market learns efficiently (i.e., under ELM),*

$$E[r_{t+1}|\mathcal{F}_t, \overline{\theta}] \geq 0,$$

$t = 1, \ldots, T-1.$

This survivorship bias was illustrated with empirical and numerical examples in the previous sections.

It is somewhat surprising that the survivorship bias can be signed under ELM, because we have not fixed the market's priors: λ^* can differ arbitrarily from λ. In fact, in the proof of Theorem 1 in the Appendix, we allow λ^*, and, hence, the market's beliefs bias, to be arbitrarily random. One could have expected that Theorem 1 would not obtain if it is known that the market always tends to be too optimistic, inducing a negative drift on securities prices. Theorem 1 states that this conjecture is wrong.

As the empirical and numerical examples of the previous section illustrated, one can adjust the traditional return measure and correct for survivorship bias under ELM. In particular, consider the *modified return* x_{t+1}:

$$x_{t+1} = \frac{p_{t+1} - p_t}{p_{t+1}}.$$

When weighting x_{t+1} with the final payoff on the security, we eliminate survivorship bias. This is stated formally in the following theorem.

Theorem 2 *Assume risk neutrality and zero discounting. If the market learns efficiently (i.e., under ELM),*

$$E[x_{t+1}V^+|\mathcal{F}_t, \overline{\theta}] = 0,$$

$t = 1, \ldots, T-1$.

The IEM Microsoft "High" data illustrated the working of this Theorem in the case of digital options (for which $V^+ = 1$). The CBOE S&P500 index call options data did this for general equity-like payoffs (V^+ is the difference between the S&P500 index level and the option's strike price). The numerical examples in the previous sections complemented the empirical illustrations.

The Appendix provides the proof of Theorem 2. It may seem demanding (especially because random priors are allowed for), but its core is simple, at least for digital options. The core relies on arguments that have already been employed elsewhere in the mathematical-statistics literature, to show that sequences of ratios of likelihood functions are martingale processes (see, e.g., Doob [1953], II.7 and VII.9). Likelihood functions appear here because, through Bayes' law, they determine how prices change as the result of information accumulation.

To understand how the proof works for digital options, let us assume that the information at $t+1$ can be represented with a continuous signal s_{t+1} taking values in some set S_{t+1}. Its likelihood given θ (derived from P_θ) is continuous, and denoted $l_{t+1}(s_{t+1}|\theta)$. Also, let $\lambda_t^*(\theta)$ denote the market's posterior about θ, given all the information up to time t ($\mathcal{F}_t$). In particular,

$\lambda_t^*(\overline{\theta})$ denotes the market's posterior about the event $\{\theta = \overline{\theta}\}$. Since we are investigating a digital option, $\lambda_t^*(\overline{\theta})$ will also equal the equilibrium market price, p_t. The market's posterior at $t+1$, $\lambda_{t+1}^*(\overline{\theta})$, is obtained from the signal and $\lambda_t^*(\overline{\theta})$ using Bayes' law. It will also equal the price at time $t+1$, p_{t+1}. Hence,

$$\begin{aligned}\frac{p_t}{p_{t+1}} &= \lambda_t^*(\overline{\theta})\frac{l_{t+1}(s_{t+1}|\overline{\theta})\lambda_t^*(\overline{\theta}) + l_{t+1}(s_{t+1}|\underline{\theta})\lambda_t^*(\underline{\theta})}{l_{t+1}(s_{t+1}|\overline{\theta})\lambda_t^*(\overline{\theta})} \\ &= \frac{l_{t+1}(s_{t+1}|\overline{\theta})\lambda_t^*(\overline{\theta}) + l_{t+1}(s_{t+1}|\underline{\theta})\lambda_t^*(\underline{\theta})}{l_{t+1}(s_{t+1}|\overline{\theta})}.\end{aligned}$$

(Notice that the prior $\lambda_t^*(\overline{\theta})$ in front of the first fraction canceled; this is why prior beliefs are irrelevant in Theorem 2.) Taking conditional expectations using $l_{t+1}(s_{t+1}|\overline{\theta})$ generates:

$$\begin{aligned}&E[\frac{p_t}{p_{t+1}}|\mathcal{F}_t,\overline{\theta}] \\ &= \int_{S_{t+1}} \frac{l_{t+1}(s_{t+1}|\overline{\theta})\lambda_t^*(\overline{\theta}) + l_{t+1}(s_{t+1}|\underline{\theta})\lambda_t^*(\underline{\theta})}{l_{t+1}(s_{t+1}|\overline{\theta})} l_{t+1}(s_{t+1}|\overline{\theta})ds_{t+1} \\ &= \int_{S_{t+1}} l_{t+1}(s_{t+1}|\overline{\theta})\lambda_t^*(\overline{\theta}) + l_{t+1}(s_{t+1}|\underline{\theta})\lambda_t^*(\underline{\theta})ds_{t+1} \\ &= \left(\int_{S_{t+1}} l_{t+1}(s_{t+1}|\overline{\theta})ds_{t+1}\right)\lambda_t^*(\overline{\theta}) + \left(\int_{S_{t+1}} l_{t+1}(s_{t+1}|\underline{\theta})ds_{t+1}\right)(1-\lambda_t^*(\overline{\theta})) \\ &= 1.\end{aligned}$$

(Notice the cancellation of likelihood functions, possible because we assume that the market has correct conditional beliefs.) Consequently,[12]

$$E[x_{t+1}|\mathcal{F}_t,\overline{\theta}] = E[1 - \frac{p_t}{p_{t+1}}|\mathcal{F}_t,\overline{\theta}] = 0.$$

6 Extensions

The theory from the previous section can be extended in various directions.

First, risk aversion and/or nonzero riskfree rates can easily be built into the analysis, provided one has a candidate state-price deflator (stochastic discount factor). The Appendix discusses how, and restates Theorem 2 in this more general framework.

Second, under mild conditions, the *unweighted* modified return is restricted as well. In particular, Bossaerts [1997] proves when the following obtains:

$$E[x_{t+1}|\mathcal{F}_t,\overline{\theta}] \leq 0. \tag{17}$$

[12]Lones Smith independently proved a weaker version of this result, under correct beliefs and independent signals. His motivation was entirely different. The question whether the modified return satisfied martingale difference restrictions showed up as part of the determination of optimal investment strategies. See Smith [1996].

Finally, until now, we have only investigated the price pattern of winning securities. One obviously wonders whether prices of losing securities (i.e., those that expired out-of-the-money) reveal anything about ELM. Subject to certain conditions, the answer is affirmative. The restrictions require, however, that one investigate the return series in reverse time. Because the ensuing complexity distracts from the main objective of this paper, the analysis of returns on losing securities is not included here. The interested reader may consult Bossaerts [1997].

7 Conclusion

The concept of market efficiency is central to asset pricing. Intuitively, it is understood to mean that the market does not over-react or under-react to new information. In other words, the market updates its beliefs rationally, using the rules of conditional probability (Bayes' law). Empirical tests of asset pricing models, however, implicitly assume also that the market's initial beliefs are unbiased. Likewise, Lucas' theory of rational expectations equilibria explicitly posits that the representative consumer correctly assesses the frequency of future random events.

This paper is a first attempt to relax the assumption of correct priors. Its aim is to demonstrate that (i) securities prices remain restricted even if one allows the market to have biased priors, (ii) the resulting retrictions are simple and powerful. The paper does so in the context of equity-like contracts: securities that pay nothing in one state (the bankruptcy state), and a random amount in the other state. The market is allowed to have wrong priors about the likelihood of the bankruptcy state, yet is assumed to update its beliefs rationally as information arrives.

In such a world, restrictions emerge only after deliberately imposing a selection bias on price histories. The paper thereby demonstrates that market efficiency can be tested on samples that are subject to survivorship bias. The concept of market efficiency that is being tested is, however, weaker than the traditional EMH: only rational updating is assumed; priors can be biased. We referred to this concept as ELM (Efficiently Learning Markets).

One aspect that this paper does not dwell on is the nature of the financial market where the restrictions would fail. Let us assume that prices were correctly adjusted for risk and/or nonzero discounting. (The body of the paper actually assumes risk neutrality and zero discounting; the Appendix, however, explains how to accommodate risk aversion and/or nonzero discounting.) Such a market must then be making mistakes in updating its own beliefs. It either under-reacts or over-reacts to new information. Probability theorists would describe the situation as one where a Dutch book could be used to make money in gambles against the market, with no risk.[13] Asset pricing theorists would refer to it as an (asymptotic) arbitrage opportunity. It can be shown, for instance, that over-reaction to new information leads to *negative* weighted average modified returns on winning contracts; conversely, under-reaction leads to *positive* weighted average modified returns. Over-reaction and under-reaction are measured relative to the Bayesian update. The interested reader is referred to Bossaerts and Hillion [1997] for further analysis.

[13] See, e.g., Schervish [1995], p. 655-6.

References

Bekaert, G., R.J. Hodrick and D.A. Marshall, 1995, "'Peso Problem' Explanations For Term Structure Anomalies," Stanford University working paper.

Bondarenko, O., 1997, "Testing Rationality of Financial Markets: An Application to S&P 500 Index Options," Caltech working paper.

Bondarenko, O. and P. Bossaerts, 1996, "Expectations and Learning in Iowa," Caltech working paper.

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

Bossaerts, P., 1996, "Martingale Restrictions On Securities 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. and P. Hillion, 1997, "IPO Post-Issue Markets: Questionable Predilections But Diligent Learners?" Caltech working paper.

Bray, M. and D. Kreps, 1987, "Rational Learning and Rational Expectations," in: *Arrow and the Ascent of Modern Economics Theory*, ed. G.R. Feiwel. New York: NYU Press.

Brown, S., W. Goetzmann and S. Ross, 1995, "Survival," *Journal of Finance* 50, 853-874.

Cecchetti, S., P. Lam and N. Mark, 1997, "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to be True?" Ohio State University working paper.

Daniel, K., D. Hirshleifer and A. Subrahmanyam, 1997, "A Theory of Overconfidence, Self-Attribution, and Securities Returns Anomalies," UCLA Anderson School working paper.

De Bondt, W. and R.H. Thaler, 1985, "Does the Stock Market Overreact?" *Journal of Finance* 40, 793-805.

Dunn, K. and K. Singleton, 1986, "Modeling the Term Structure of Interest Rates under Non-Separable Utility and Durability of Goods," *Journal of Financial Economics* 17, 27-56.

Doob, J.L. (1952): *Stochastic Processes.* New York: Wiley.

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.

Grandmont, J.M., 1977, "Temporary General Equilibrium Theory," *Econometrica* 45, 535-572.

Hansen, L.P. and R. Jagannathan, 1991, "Implications of Security Market Data For Models Of Dynamic Economies," *Journal of Political Economy* 99, 225-262.

Harsanyi, J.C., 1967, "Games with Incomplete Information Played by 'Bayesian' Players, I: The Basic Model," *Management Science* 14, 159-82.

Jegadeesh, N., 1990, "Evidence of Predictable Behavior of Security Returns," *Journal of Finance* 45, 881-898.

Keim, D.B. and R.F. Stambaugh (1986): "Predicting Returns in Stock and Bond Markets ," *Journal of Financial Economics* 17, 357-390.

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.

Mehra, and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145-161.

Muth, J.F., 1961, "Rational Expectations and the Theory of Price Movements," *Econometrica* 29, 315-35.

Samuelson, P. A., 1965, "Proof that properly anticipated prices fluctuate randomly," *Industrial Management Review* 6, 41–50.

Savage, L., 1954, *The Foundations Of Statistics.* New York: Wiley.

Schervish, M., 1995, *Theory of Statistics.* New York: Springer Verlag.

Smith, L., 1996, "Private Information and Trade Timing," MIT working paper.

Appendix

There are two parts to this Appendix. First, proofs of the Lemma and Theorems in the paper are given, albeit in the more general context of random market priors (we will discuss what that means). Second, the main result, Theorem 2, is extended to allow for risk aversion and/or nonzero discounting.

Part I: Proofs

I.A. Preliminaries

To capture potential randomness in the market's prior over θ, λ^*, extend Ω ($= \Theta \times \tilde{\Omega}$) by the Cartesian product with the set of probability measures on Θ, denoted $\Pi(\Theta)$. $\lambda^* \in \Pi(\Theta)$. We obtain the following outcome space:

$$\Theta \times \Pi(\Theta) \times \tilde{\Omega}.$$

Extend the probability measure P accordingly, obtaining Q, to be factored as follows:

$$Q = \lambda \pi_\theta P_{\theta\lambda^*}. \tag{18}$$

Under traditional market efficiency, $\pi_\theta = \delta_\lambda$, i.e., the probability measure that puts unit mass on the correct belief about θ, namely λ. Nothing fundamental is added by allowing the outcomes in $\Pi(\Theta)$ (i.e., market beliefs) and in $\tilde{\Omega}$ (the latter determine signals and payoffs) to be independent conditional on θ. In fact, assuming the converse would mean that knowledge of the market's priors reveals information about the likely signals and eventual payoff, which is rather awkward as an empirical proposition. Hence, assume that $P_{\theta\lambda^*}$ does not change with λ^*. The earlier notation P_θ wil be used when referring to $P_{\theta\lambda^*}$.

In the Appendix, the expectations E will be computed with respect to the more general measure, Q, which allows market beliefs to be random.

The following notation will be used throughout the remainder of the Appendix.

Let $s_1, s_2, \ldots, s_T$ denote the sequence of signals that investors receive. Each s_t lives in a space S_t. One does not have to be specific about what the S_ts are. Just let $l_t(s_t|\theta)$ denote the density (if s_t is continuous) or probability mass (otherwise) of s_t corresponding to P_θ. Following the discussion of Section 5, P^*_θ equals P_θ. Hence, $l_t(s_t|\theta)$ also denotes the density or probability mass corresponding to P^*_θ.

Agents' beliefs about θ will be written as a sequence of elements in $\Pi(\Theta)$, $\{\lambda^*_t\}_{t=1}^T$, which are recursively defined by:

$$\lambda_t^*(\overline{\theta}) = \frac{l_t(s_t|\overline{\theta})\lambda_{t-1}^*(\overline{\theta})}{\sum_{\theta\in\Theta} l_t(s_t|\theta)\lambda_{t-1}^*(\theta)} \tag{19}$$

when $t = 2, \ldots, T$;

$$\lambda_1^*(\overline{\theta}) = \lambda^*(\overline{\theta});$$

and:

$$\lambda_t^*(\underline{\theta}) = 1 - \lambda_t^*(\overline{\theta}),$$

all t.

I.B. Proof Of Lemma 1

Under EMH, market prices are determined as in (15), with E substituted for E^*. Using the law of iterated expectations,

$$\begin{aligned}
&E[r_{t+1}|\mathcal{F}_t] \\
&= E[\frac{p_{t+1} - p_t}{p_t}|\mathcal{F}_t] \\
&= \frac{1}{p_t}E[p_{t+1}|\mathcal{F}_t] - 1 \\
&= \frac{1}{p_t}E[E[V|\mathcal{F}_{t+1}]|\mathcal{F}_t] - 1 \\
&= \frac{1}{p_t}E[V|\mathcal{F}_t] - 1 \\
&= 0.
\end{aligned}$$

□

A.III. Proof Of Theorem 1

Let $1_{\{\theta=\overline{\theta}\}}$ denote the indicator function, which takes the value 1 if $\theta = \overline{\theta}$, and 0 otherwise. From (15) and (19), the time-$(t+1)$ price of the digital option equals:

$$\begin{aligned}
p_{t+1} &= E^*[V|\mathcal{F}_{t+1}] \\
&= E^*[1_{\{\theta=\overline{\theta}\}}|\mathcal{F}_{t+1}] \\
&= \lambda_{t+1}^*(\overline{\theta}) \\
&= \frac{l_{t+1}(s_{t+1}|\overline{\theta})\lambda_t^*(\overline{\theta})}{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)}.
\end{aligned} \tag{20}$$

Using $p_t = \lambda_t^*(\overline{\theta})$,

$$\frac{p_{t+1}}{p_t} = \frac{\lambda_t^*(\overline{\theta})}{\lambda_t^*(\overline{\theta})} \frac{l_{t+1}(s_{t+1}|\overline{\theta})}{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)}.$$

Now consider:

$$\begin{aligned} & E[\frac{p_{t+1}}{p_t}|\mathcal{F}_t,\overline{\theta}] \\ = \quad & E[E[\frac{p_{t+1}}{p_t}|\mathcal{F}_t,\overline{\theta},\lambda^*]|\mathcal{F}_t,\overline{\theta}]. \end{aligned}$$

Compute the *inner* conditional expectation, using the assumption that $P_{\theta\lambda^*} = P_\theta$:

$$\begin{aligned} & E[\frac{p_{t+1}}{p_t}|\mathcal{F}_t,\overline{\theta},\lambda^*] \\ = \quad & \int_{S_{t+1}} \frac{l_{t+1}(s_{t+1}|\overline{\theta})}{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)} P_{\overline{\theta}}(ds_{t+1}) \\ = \quad & \int_{S_{t+1}} \frac{1}{\frac{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)}{l_{t+1}(s_{t+1}|\overline{\theta})}} l_{t+1}(s_{t+1}|\overline{\theta}) ds_{t+1} \\ \geq \quad & \frac{1}{\int_{S_{t+1}} \frac{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)}{l_{t+1}(s_{t+1}|\overline{\theta})} l_{t+1}(s_{t+1}|\overline{\theta}) ds_{t+1}} \\ = \quad & \frac{1}{\sum_{\theta\in\Theta}\int_{S_{t+1}} l_{t+1}(s_{t+1}|\theta) ds_{t+1}\lambda_t^*(\theta)} \\ = \quad & 1. \end{aligned}$$

The result immediately obtains:

$$\begin{aligned} E[r_{t+1}|\mathcal{F}_t,\overline{\theta}] \quad &= \quad E[\frac{p_{t+1}}{p_t}|\mathcal{F}_t,\overline{\theta}] - 1 \\ &= \quad E[E[\frac{p_{t+1}}{p_t}|\mathcal{F}_t,\overline{\theta},\lambda^*]|\mathcal{F}_t,\overline{\theta}] - 1 \\ &\geq \quad 0. \end{aligned}$$

□

A.IV. Proof Of Theorem 2

Let us first prove the following Lemma.

Lemma 2

$$E[\frac{\lambda_t^*(\overline{\theta})}{\lambda_{t+1}^*(\overline{\theta})}|\mathcal{F}_t,\overline{\theta}] = 1.$$

Proof:

$$\begin{aligned} & E[\frac{\lambda_t^*(\overline{\theta})}{\lambda_{t+1}^*(\overline{\theta})}|\mathcal{F}_t,\overline{\theta}] \\ = \quad & E[E[\frac{\lambda_t^*(\overline{\theta})}{\lambda_{t+1}^*(\overline{\theta})}|\mathcal{F}_t,\overline{\theta},\lambda^*]|\mathcal{F}_t,\overline{\theta}] \end{aligned}$$

Computing the inner conditional expectation, using arguments from the proof of Theorem 1:

$$
\begin{aligned}
& E[\frac{\lambda_t^*(\overline{\theta})}{\lambda_{t+1}^*(\overline{\theta})}|\mathcal{F}_t,\overline{\theta},\lambda^*] \\
&= \int_{S_{t+1}} \frac{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)}{l_{t+1}(s_{t+1}|\overline{\theta})} P_{\overline{\theta}}(ds_{t+1}) \\
&= \int_{S_{t+1}} \frac{\sum_{\theta\in\Theta} l_{t+1}(s_{t+1}|\theta)\lambda_t^*(\theta)}{l_{t+1}(s_{t+1}|\overline{\theta})} l_{t+1}(s_{t+1}|\overline{\theta}) ds_{t+1} \\
&= \sum_{\theta\in\Theta} \int_{S_{t+1}} l_{t+1}(s_{t+1}|\theta) ds_{t+1} \lambda_t^*(\theta) \\
&= 1.
\end{aligned}
$$

This leads to the desired result.

(End of Proof of Lemma 2.)

Now the proof of Theorem 2. First, notice from (15), the assumption of correct conditional beliefs and of independence between V^+ and λ^* (remember: $P_{\theta\lambda^*} = P_\theta$):

$$
\begin{aligned}
p_t &= E^*[V|\mathcal{F}_t] \\
&= E^*[V^+ 1_{\{\theta=\overline{\theta}\}}|\mathcal{F}_t] \\
&= E^*[V^+|\mathcal{F}_t,\overline{\theta}]\lambda_t^*(\overline{\theta}) \\
&= E[V^+|\mathcal{F}_t,\overline{\theta},\lambda^*]\lambda_t^*(\overline{\theta}) \\
&= E[V^+|\mathcal{F}_t,\overline{\theta}]\lambda_t^*(\overline{\theta}).
\end{aligned}
$$

Hence,

$$
\begin{aligned}
& E[V^+ x_{t+1}|\mathcal{F}_t,\overline{\theta}] \\
&= E[x_{t+1} E[V^+|\mathcal{F}_{t+1},\overline{\theta}]|\mathcal{F}_t,\overline{\theta}] \\
&= E[V^+|\mathcal{F}_t,\overline{\theta}] \\
&\quad - E\left[\frac{E[V^+|\mathcal{F}_t,\overline{\theta}]\lambda_t^*(\overline{\theta})}{E[V^+|\mathcal{F}_{t+1},\overline{\theta}]\lambda_{t+1}^*(\overline{\theta})} E[V^+|\mathcal{F}_{t+1},\overline{\theta}]|\mathcal{F}_t,\overline{\theta}\right] \\
&= E[V^+|\mathcal{F}_t,\overline{\theta}] \\
&\quad - E[V^+|\mathcal{F}_t,\overline{\theta}] E\left[\frac{\lambda_t^*(\overline{\theta})}{\lambda_{t+1}^*(\overline{\theta})}|\mathcal{F}_t,\overline{\theta}\right] \\
&= 0,
\end{aligned}
$$

where the last equality follows from Lemma 2.

□

Part II: Risk Aversion And/Or Nonzero Discounting

Until now, risk neutrality and zero interest rates have been assumed. To allow for risk aversion and/or nonzero interest rates, i.e., to accommodate stochastic discount rates, one should chose a state-price deflator and use it to normalize prices.

Here is how it works. Let M_t denote the chosen (stochastic) state-price deflator, for $t = 1, \ldots, T+1$. If $\tilde{p}_t$ is the raw price level, then the following restriction characterizes equilibrium prices:[14]

$$M_t \tilde{p}_t = E^*[M_{T+1} V | \mathcal{F}_t], \tag{21}$$

The deflated price p_t is obtained as $M_t \tilde{p}_t$. Defining:

$$M = M_{T+1},$$

we obtain the following pricing formula:

$$p_t = E^*[MV | \mathcal{F}_t], \tag{22}$$

For our purposes, the representation in terms of deflated prices (22) is more useful than the original representation using raw prices (21). This should be clear once we get to the proof of the main result. Define:

$$x_{t+1} = \frac{p_{t+1} - p_t}{p_{t+1}}.$$

(Remark: this modified return is based on *deflated prices*!)

Theorem 3 *If the market learns efficiently (i.e., under ELM),*

$$E[x_{t+1} M V^+ | \mathcal{F}_t, \overline{\theta}] = 0,$$

$t = 1, \ldots, T-1$.

Proof: Notice the similarity between the pricing equations (22) and (15): the former can be obtained by substituting MV for V in the latter. The Theorem follows from implementing the same substitution in the proof of Theorem 2, and remembering that, conditional on $\theta = \overline{\theta}$, $V = V^+$.

□

[14] The (strictly positive) ratio M_{T+1}/M_t is often referred to as pricing kernel. See, e.g., Hansen and Jagannathan [1991].

Table 1

Securities Prices In The Iowa Experimental Markets:

Digital Options Written On Microsoft (MS) Common Stock

	N	Mean
All MS High Contracts		
Return	16	0.081*
		(0.042)
All MS Winning Contracts		
Return	16	0.097**
		(0.037)
Modified Return	16	-0.032
		(0.022)

Remarks: To compute the modified return, the *end-of-period* price is used as basis; N is the number of time series (each one month long); The averages are computed as the cross-sectional average of the time series mean daily return or modified return; Standard errors in parentheses; *: significant at the 5% level; **: significant at the 1% level.

Table 2

Prices Of CBOE S&P500 Index Call Options

	N	Average
All Contracts		
Return	58	-0.008
		(0.007)
Winning Contracts		
Return	40	0.013*
		(0.006)
Weighted Return	40	0.353**
		(0.127)
Weighted Modified	40	0.004
Return		(0.114)

Remarks: Modified returns are computed based on the *end-of-period* price; Weighted returns are returns multiplied by the payoff on the security at maturity; N is the number of time series (each four weeks long); The averages are computed as the cross-sectional average of the time series mean daily return or modified return; Standard errors in parentheses; *: significant at the 5% level; **: significant at the 1% level.

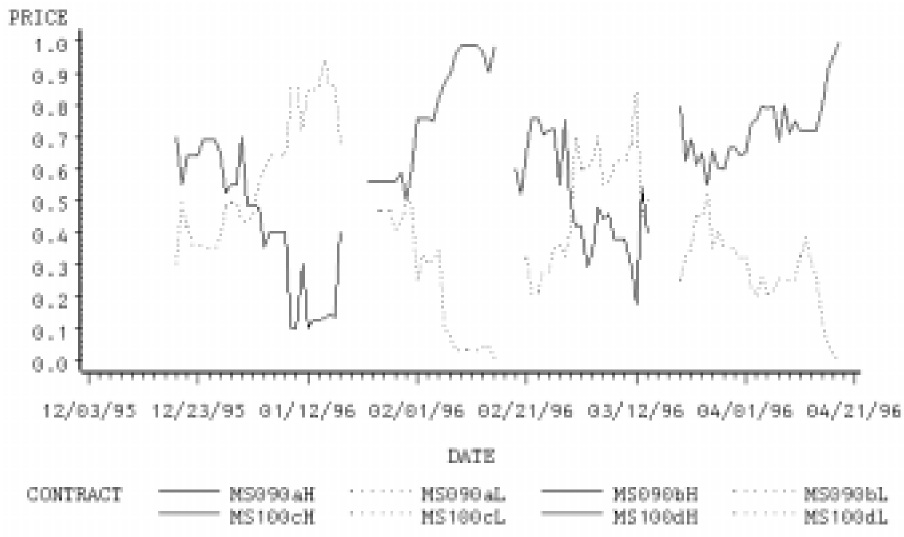

Figure 1: Time series plots of daily closing prices for contracts Microsoft High (bold lines) and Microsoft Low (dotted lines) in the Iowa experimental markets. One-month histories for expiration months 1/96 (contracts MS090a*) till 4/96 (contracts MS100d*) are shown.

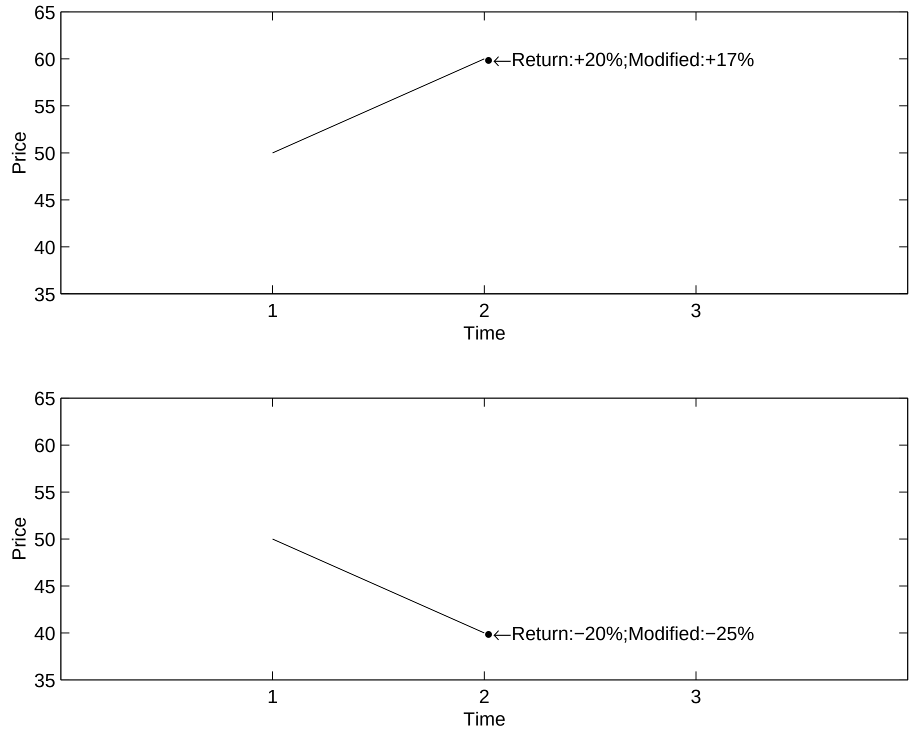

Figure 2: Left panel: when prices increase, the modified return (based on the end-of-period price level) is lower than the traditional return (based on the beginning-of-period price level). Right panel: when prices decrease, the modified return is larger (in absolute value) than the traditional return. For digital options that matured in-the-money, the traditional return is positive on average. The modified return can be proven to be zero on average, provided that the market reacts rationally to news about the eventual payoff. The result obtains even if the market has biased expectations *ex ante*.

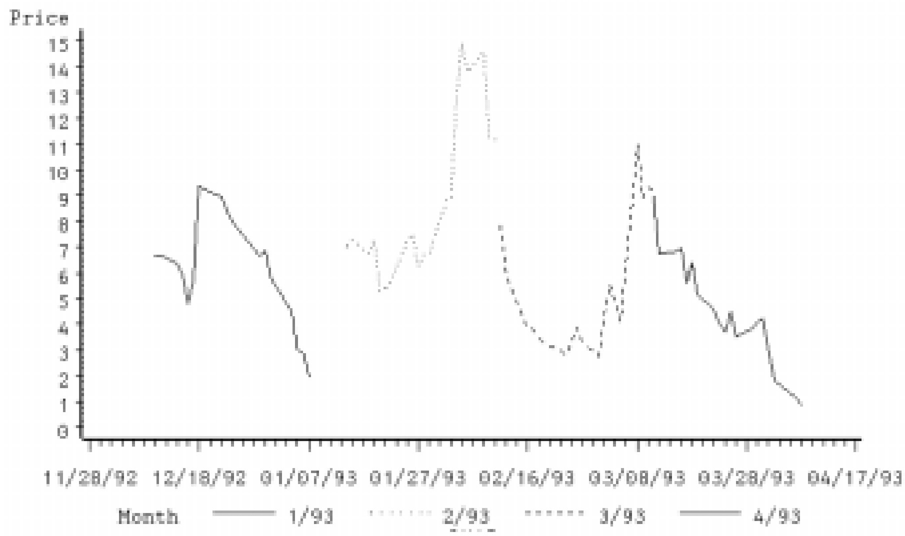

Figure 3: Time series plots of daily closing prices for CBOE S&P500 index call options. The options were at the money at the beginning of each time series, and expired five weeks later. Four-week histories for four expiration months are shown.

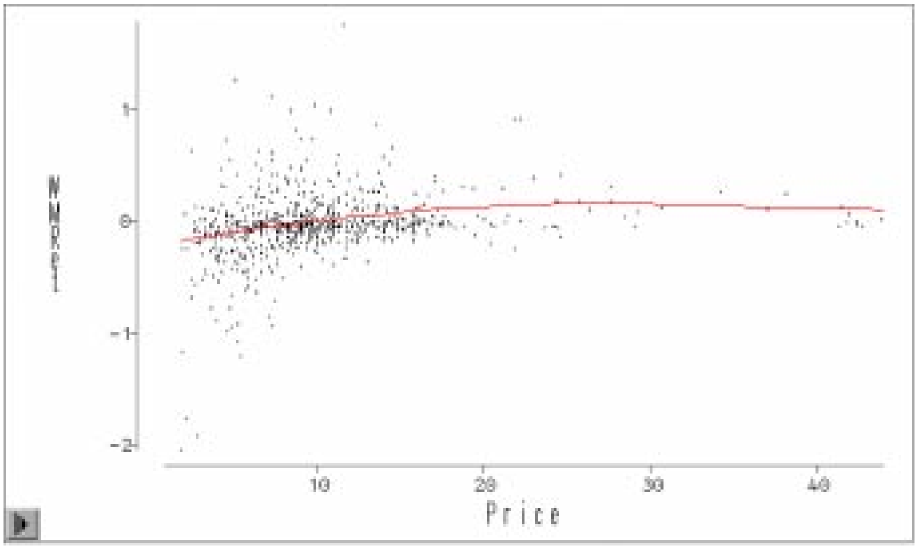

Figure 4: Scatter plot of weighted daily modified returns divided by lagged prices (WMRet) against lagged prices for CBOE S&P500 index call options. The line is an estimate of the regression function (local polynomial estimation). The conditional expectation is negative and increasing for prices below \$10; it is zero above \$10.

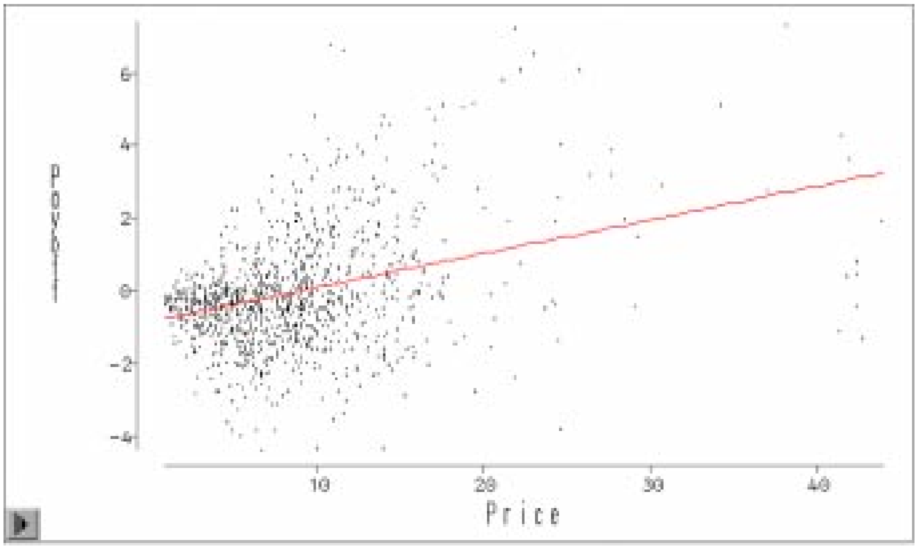

Figure 5: Scatter plot of daily payoffs (returns times lagged prices) against lagged prices for CBOE S&P500 index call options. Straight line is the OLS projection.