Equilibrium Asset Pricing Under Heterogeneous Information

By

Bruno Biais

Toulouse University and CEPR

Peter Bossaerts

California Institute of Technology and CEPR

and

Chester Spatt

Carnegie Mellon University

Revised January, 2003

Equilibrium Asset Pricing Under Heterogeneous Information

Abstract

We analyze theoretically and empirically the implications of heterogeneous information for equilibrium asset pricing and portfolio choice. Our theoretical framework, directly inspired by Admati (1985), implies that with partial information aggregation, portfolio separation fails, buy-and-hold strategies are not optimal, and investors should structure their portfolios using the information contained in prices in order to cope with winner's curse problems. We implement empirically such a price-contingent portfolio allocation strategy and show that it outperforms economically and statistically the passive/indexing buy-and-hold strategy. We thus demonstrate that prices *reveal* information, in contrast with the homogeneous information CAPM, but only *partially*, consistent with Noisy Rational Expectations Equilibrium. The success of our price-contingent strategy does not proxy for the success of trading strategies based purely on historical performance, such as momentum investment.

1 Introduction

The theory of financial markets under homogeneous information has generated a rich body of predictions, extensively used in the financial industry, such as the optimality of indexing, the nature of arbitrage, and equilibrium–based pricing relations, as illustrated by the CAPM. In contrast, the theory of capital markets under heterogeneous information has not been used much to guide asset pricing and portfolio allocation decisions. The goal of the present paper is to derive some of the implications of partially revealing (noisy) rational expectations equilibria for asset pricing and asset allocation, and to test their empirical relevance.

Our theoretical framework is inspired directly by Admati (1985). Her analysis, which involves multiple risky assets and mean–variance preferences, encompasses the CAPM, and extends it to the case where investors observe private signals. In this context, when the supply of risky assets is known by all agents, prices are fully revealing, the CAPM holds, and consequently all investors hold the market portfolio. In contrast, with random asset supplies, the rational expectations equilibrium in Admati (1985) is noisy, i.e., prices only partially reveal private information.[1] Note that the randomness of supply shocks implies that investors do not know the exact structure of the market portfolio. Thus, the key assumption that assets are in random supply is in line with the Roll (1977) Critique, emphasizing the unobservability of the market portfolio.

We focus on a special case of this model, where all investors have identical precisions, but different signals. In this context, we show that equilibrium prices are equal to the prices that would arise in a representative-agent economy. The fictitious representative agent would observe the structure of the market portfolio and have beliefs equal to the average beliefs of the investors. Holding the market portfolio is optimal from the perspective of this representative agent. While the corresponding equilibrium pricing relation is similar to the CAPM, portfolio choices of the actual investors in this economy differ markedly from their CAPM counterparts. Each investor holds a portfolio that deviates from the market portfolio in response to his or her private signal. While this choice is rational, given the information set of the agents, it entails a winner's curse: the investor holds more (resp. less) of an asset than the market portfolio when his or her signal is above (resp. below) the average investor's opinion of the value of the asset, i.e., they hold the market portfolio plus a "tilt" portfolio that reflects their private signal. Investors cope with the winner's curse problem by complementing their private information with the information reflected in prices. Consequently, buy-and-hold strategies are not optimal. This is true even for

---

[1]Admati (1985) extends the noisy rational expectations model in which investors receive diverse signals with one risky asset in Hellwig (1980) to many risky assets. The presence of many risky assets is crucial to a CAPM analysis.

investors who do not receive a private signal as they also will rebalance their portfolios to reflect the informational content of prices. That is, optimal portfolio allocations are price–contingent, even for uninformed investors.

To understand why buy-and-hold is suboptimal, one can draw an analogy with auction theory. In first-price auctions bidders adjust for the winner's curse by shading the price they bid, conditional on winning. In our setup markets are competitive, and hence, price adjustments are ruled out. Instead, agents can and will make quantity adjustments, conditional on prices.

The suboptimality of buy-and-hold contrasts with the case where information is homogeneous or prices are fully revealing. In that context, price changes cause portfolio weights to change exactly as necessary for the benchmark portfolios to remain optimal. This differs from the Noisy Rational Expectation Equilibrium, where prices are only partially revealing.

To evaluate the empirical relevance of this theoretical analysis, we study the performance of the price–contingent portfolio allocation strategy it suggests, using monthly U.S. stock data over the period 1927-2000. Taking the perspective of an uninformed but rational agent, we extract the information contained in prices by projecting returns onto (relative) prices. We use the corresponding expected returns and variance–covariance matrix to construct the mean–variance optimal portfolio. We then compare the performance of this portfolio, as measured by its Sharpe ratio, to that of passively buying and holding the value–weighted CRSP index. We find that the optimal price–contingent portfolio outperforms the buy-and-hold strategy both economically and statistically. Our findings thus demonstrate empirically how prices reveal *some* information, in contrast with the homogeneous information CAPM, while *not fully* aggregating information, consistent with Noisy Rational Expectations Equilibrium.[2]

It should be emphasized that the optimal price–contingent portfolio allocation strategy we analyze is entirely based on ex–ante information. Portfolio decisions made at the beginning of month t rely on price and return data prior to month t. Thus, we only use information available to market participants when they chose their portfolios. Hence, our result that the optimal price– contingent allocation strategy outperforms the buy-and-hold indexing strategy differs from the

---

[2]There have been a number of interesting applications of the noisy rational expectations framework. For example, Cho and Krishnan (2000) evaluate the importance of prices in aggregating private information in the S&P 500 futures market by estimating the primitive parameters of the Hellwig (1980) (single risky asset) model. Brennan and Cao (1997) study the implications for international investment flows. Grundy and Kim (2002) study the implications for volatility of partially revealing rational expectations equilibrium. An application of this type of equilibrium to the study of market microstructure is Kalay and Wohl (2001). Our analysis differs from these studies because we focus on the implications for the Capital Asset Pricing Model and implementing the portfolio choices implied by the theory.

Fama and French (1996) findings. Fama and French show that, based on return means, variances and covariances estimated as empirical moments over a period including month t as well as later months, an optimal combination of their "factor portfolios" outperforms the index. However, Cooper, Gutierrez, and Marcum (2002) show that, if one estimates these empirical moments using only information prior to month t, the factor portfolios fail to outperform the buy-and-hold strategy. To put our results into perspective we replicate their results, translating them into the mean–variance framework of our theoretical model.

The price-contingent portfolio allocation strategy that we implement implicitly uses past return information. Hence, our finding that it outperforms buy-and-hold might be related to the success of strategies based on historical performance, in particular, the popular momentum strategy.[3] Still, our investment strategy is not proxying for momentum trading, because the momentum strategy does not outperform ours. In addition, even though the returns on the momentum strategy and ours are related, the nature of the information they exploit is different. The momentum strategy exploits returns continuations; our price contingent strategy exploits the correlation between the current prices and future returns. In the theoretical analysis of Admati (1985), the correlation between prices and subsequent returns can be positive, negative or insignificant. Indeed, the correlations we empirically estimate, and use in our price contingent strategies, have variable signs.

In the next section we present our theoretical framework. Section 3 presents the empirical analysis. Section 4 offers a further discussion of the results. Section 5 concludes.

2 Partially Revealing Equilibrium with Multiple Risky Assets

2.1 The Model

Our basic set-up is inspired directly by Admati (1985). As in Admati (1985) we consider:

- A two–date model in which portfolio allocation takes place at time t=0, while asset returns are earned and consumption takes place at time t=1.

- N risky assets with payoffs at time 1: $f_i, i = 1, ..., N$, and one riskless asset, which also serves the role of numeraire, and earns exogenous return r_f at time 1. (We adopt the

---

[3]See, e.g., Chan, Jegadeesh and Lakonishok (1996) and Lewellen (2002).

convention for the random variables we consider that lower case letters denote scalars, while upper case letters denote vectors.)

- A continuum of agents: $a \in [0,1]$, observing signals $y_{a,i} = f_i + \epsilon_{a,i}$. The precision of agent a's signal is denoted S_a (that is S_a is the variance–covariance matrix of the N noise terms $\epsilon_{a,i}, i = 1, ..., N$).

- The supply of asset i is random and equal to z_i. It is not observed by the investors. It is this noise which will prevent full revelation of the private information in equilibrium.

- All the random variables are assumed to be jointly normally distributed, and the noise terms, the aggregate endowments, and the payoffs are independent.

- The agents have constant absolute risk averse (CARA) utility. The absolute risk tolerance coefficient of agent a is denoted ρ_a. The average risk tolerance: $\int_a \rho_a da$ is denoted ρ.

In this context, as shown by Admati (1985, Theorem 3.1, p. 637), there exists a linear rational expectations equilibrium, whereby:

$$P = A + BF + CZ,$$

where A, B, and C are constant vector and matrices, while P is the $(N, 1)$ vector of prices, F is the $(N, 1)$ vector of cash flows, and Z is the $(N, 1)$ vector of aggregate endowments.

In this equilibrium, as in the standard CAPM, prices are equal to expected cash flows minus a risk premium related to the supply of the risky assets. Because there is a continuum of informed agents with signals equal to the final cash flow plus a noise term, prices, which aggregate the investors' information, reflect the final cash flow (F). However, because the supply shocks and correspondingly the aggregate supply of the risky assets are not known by the agents, prices are not fully revealing. In this context, investors condition their portfolio decisions on prices, but must also use their signals. Thus, unlike in a standard CAPM, investors do not follow buy-and-hold strategies, as they alter their portfolio holdings to react to their signals as well as prices. Note that the random supply shocks imply that the market portfolio is not observed by the agents.[4]

In the following we elaborate on the formal similarity in terms of pricing between the standard CAPM and the partially revealing equilibrium of the present model. We then explain in more detail to what extent the two models make different predictions in terms of portfolio choices.

---

[4]This unobservability of the market portfolio by the investors is central to the theoretical foundation of the model. It is consistent with the unobservability for the econometrician discussed in the Roll Critique (1977).

2.2 Equilibrium Prices And Returns

Here, we analyze in more detail equilibrium prices and returns in the linear rational expectations equilibria (REE) characterized in Admati (1985). For simplicity, we focus on the case where all agents have identical precisions, where, as shown below, an "aggregate CAPM" holds.

In the linear REE, all variables are jointly normal. Let p_i denote the price of asset i, and Q_a the vector of demands of agent a for the N assets. The program of the agent is:

$$Max_{Q_a} E[U_a(\sum_{i=1}^{N} q_{a,i}(v_i - p_i(1 + r_f)))|I_a],$$

where I_a denotes the information set of agent a, which consists of private signals as well as prices: (Y_a, P). Since the agents have CARA utility functions, and since, as shown by Admati the prices and values are jointly normal, this simplifies to a mean–variance program. Hence the demand of agent a for asset i, $q_{a,i}$, is:

$$q_{a,i} = \rho_a \frac{E(f_i|I_a) - p_i(1 + r_f)}{V(f_i|I_a)} - \sum_{j \neq i} q_{a,j} \frac{cov(f_i, f_j|I_a)}{V(f_i|I_a)},$$

Let

$$E^m(f_i) = \int_a \frac{\rho_a}{\rho} E(f_i|I_a) da,$$

denote the average across agents of the conditional expectations of the cash flows of asset i. It deserves emphasis that $E^m(f_i)$ is *not* equal to the expectation of the value of the asset conditional on the union of the information sets of all the agents, i.e.,

the equilibrium returns equals:

$$E^m(r_i) - r_f = \frac{cov^m(r_i, r_m | I_a)}{V^m(r_m | I_a)}(E^m(r_m) - r_f).$$

The proposition states that equilibrium prices are identical to those which would obtain in a homogeneous information–representative agent economy, where i) the market portfolio (and the supply z_j) would be known by the representative agent, ii) his expectation of the cash flows would be $E^m(f_i)$, and iii) his perception of the variances and covariances would be equal to $cov(f_i, f_j | I_a)$. This representative agent holds the market portfolio. Hence, *the standard CAPM return-covariance relationship holds, from the perspective of the representative agent.*

Note however that the view of the representative agent contrasts with that of the actual agents in the model, who do not know the random aggregate endowments, and who act based upn their individual signals as opposed to the average market opinion. The point of view of the representative agent also differs from that of the econometrician, in particular because the market portfolio is difficult to observe, as emphasized in the Roll (1977) critique. Hence, the equivalence between the prices established in the hetereogenous agents economy and those set in the representative agent economy cannot be directly relied upon in the econometrics to estimate a representative agent model. Indeed, the beliefs and endowments of that representative agent are not observable by the econometrician. Thus we take another route to confront our model to the data, as explained in the next section.

2.3 Comparison with the literature

Admati (1985) The pricing relation stated in Proposition 1 differs from that stated in Corollary 3.5 in Admati (1985). Admati's characterizes the ex–ante expected price, computed by averaging across all realizations of the random variables, while the pricing function described here holds for each realization of the random variables. Correspondingly, Admati (1985) shows that an aggregate CAPM obtains on average across possible realizations of the random variables. This contrasts with the equilibrium relationships described in the present paper, which holds in every possible state of the world.

The difference between our Proposition 1 and Corollary 3.5 in Admati (1985) stems from our assumption that the different signals observed by the agents have the same precision. Correspondingly, in our analysis investors agree on conditional covariances, limiting the extent to which they look at their mean-variance pictures differently.[5] For example, using the mean-

---

[5] The first-order condition from agent i's portfolio problem is a linear relationship between price and quantity, reflecting the investor's own signals. Because the variances and covariances are identical across agents, the individual linear relationships aggregate to the CAPM.

variance geometry the minimum variance portfolio is identical for all agents.

Wilson (1968) Our aggregation result with heterogeneous beliefs reflects in part our assumption that agents have exponential utility. This is in line with the Gorman aggregation results obtained by Wilson (1968) (see also Huang and Litzenberger, 1988, p. 146-148). In these analyses, however, beliefs are exogenous. In ours, aggregation obtains with endogenous beliefs.

DeMarzo and Skiadas (1998) DeMarzo and Skiadas (1998) also offer a theoretical analysis of a CAPM with heterogeneous information, but our model differs from theirs. On the one hand they allow for a more general class of utility functions than we do. On the other hand, a key ingredient in our model is that the aggregate supply is unknown by the agents, which prevents prices from being fully informative. In contrast, the CAPM result obtained by DeMarzo and Skiadas (1998) reflects their assumption that the aggregate supply of each of the risky assets is common knowledge for all the agents.[6]

2.4 Portfolio Choices

For simplicity, consider the case where there are only two risky assets $(i = 1, 2)$. Agent a's holdings of asset 1 are:

$$q_{a,1} = \rho_a \frac{E(f_1|I_a) - E^m(f_1) - \frac{1}{\rho}cov(f_1, f_m|I_a)}{V(f_1|I_a)} - q_{a,2}\frac{cov(f_1, f_2|I_a)}{V(f_1|I_a)}$$

while his holdings of asset 2 are:

$$q_{a,2} = \rho_a \frac{E(f_2|I_a) - E^m(f_2) - \frac{1}{\rho}cov(f_2, f_m|I_a)}{V(f_2|I_a)} - q_{a,1}\frac{cov(f_1, f_2|I_a)}{V(f_2|I_a)}$$

After simple manipulations, we obtain the following characterization of the agents' equilibrium holdings.

Proposition 2 *When agents have identical precision, in the case where there are only two risky assets, agent a's equilibrium holdings of asset i are:*

$$q_{a,i} = \frac{\rho_a\left[\frac{E(f_i|I_a) - E^m(f_i)}{V(f_i|I_a)} - \frac{cov(f_i, f_j|I_a)}{V(f_i|I_a)}\frac{E(f_j|I_a) - E^m(f_j)}{V(f_j|I_a)}\right]}{1 - corr(f_1, f_2|I_a)^2} + \frac{\rho_a}{\rho}z_i,$$

where j denotes the other asset than i, and where $corr(.,.)$ denotes the correlation coefficient.

---

[6]Proposition 6 in DeMarzo and Duffie (1998) establishes that a CAPM holds in equilibrium. It is obtained in the context of their definition of a Linear Risk Tolerance Economy. The definition of a Linear Risk Tolerance Economy (Definition 4, pages 138 and 139) states that the endowment of agent i is $e_i = a_i + b_iV$, where V is the value of the asset and a_i and b_i are coefficients such that: $a = \sum_i a_i$ and $b = \sum_i b_i$ are common knowledge to all the agents. Hence, in this economy, the aggregate endowment of the risky assets is common knowledge.

This proposition has two implications: portfolio separation fails, and there is a winner's curse problem.

To understand why portfolio separation fails, note that the equilibrium holdings of agent a are expressed in terms of deviations from the market portfolio (asset i contributes z_i to the market portfolio). On average, agents hold the market portfolio ($E[f_i|I_a] - E^m(f_i)$ averages out across agents for $i = 1, 2$, and hence, the first term averages out to zero), so that supply equals demand. But agents do not observe the market portfolio and invest in portfolios that deviate systematically and individually. This implies that portfolio separation fails, unlike in the standard CAPM.

Importantly, the expression in the proposition reveals a winner's curse problem: agent a invests *more* than the market portfolio in asset i when his expectation of the cash flow $E[f_i|I_a]$ is *greater than* the average expectation $E^m(f_i)$, while he invests *less* otherwise. The differences $E[f_i|I_a] - E^m(f_i)$ will be larger as the prediction error of the agent increases.

The error of agents' signals can be interpreted as *estimation risk*. In the past (e.g., Kandel and Stambaugh [1996]), estimation risk has been studied under homogeneous information, in which case it only adds to variance. In our setting, information is heterogeneous, and therefore, estimation risk also yields a winner's curse. Consequently, our analysis introduces a new dimension to the nature of estimation risk.

The major empirical difference between the standard CAPM and the partially revealing linear REE, therefore, is the failure of buy-and-hold to be optimal. Agents must change the composition of their portfolio as a function of prices change and signals. To pave the way for the empirical analysis, we now consider the demand of a marginal agent with no private signal.

Corollary 3 *In our model, the demand of a marginal investor, with no private signal is:*

$$q_{a,i} = \frac{\rho_a \left[\frac{E(f_i|P) - E^m(f_i)}{V(f_i|P)} - \frac{cov(f_i,f_j|P)}{V(f_i|P)} \frac{E(f_j|P) - E^m(f_j)}{V(f_j|P)} \right]}{1 - corr(f_1,f_2|P)^2} + \frac{\rho_a}{\rho} z_i.$$

The corollary illustrates that the uninformed agent's demand reflects the information content of prices. Thus, the marginal uninformed agent does not find it optimal to buy and hold a fixed portfolio. This is because he must adjust the quantity of shares he holds in response to the information content of price changes, to cope with the winner's curse induced by asymmetric information. Of course, since he does not observe any private signal, he faces an even stronger adverse selection problem than the informed agents.

3 Econometric approach

The theoretical section showed that in the noisy rational expectations equilibrium, prices do reveal information, but because the revelation is only partial, buy-and-hold strategies are not optimal. We now wish to assess the empirical relevance of such an analysis of financial market equilibrium. The exercise will shed light on important questions with respect to market efficiency. Do prices reflect information to a significant extent? Do they reflect all available information? Or is some significant amount of information only partially revealed, as in the model of the previous section?

To answer these questions, we compare the performance of buying and holding to a portfolio allocation strategy that uses information contained in prices to predict expected returns.[7] If indeed a significant amount of information is revealed in prices but only partially so, then price-contingent portfolio allocation strategies will outperform buy and hold.[8]

In itself, rejection of the optimality of buy and hold may not seem like a new result. It has long been known that proxies for the market portfolio have been inferior historically (see Fama and French [1996] and Davis, Fama and French [2000]). The inferiority has been obtained purely on an *ex post* basis, however. That is, proxies of the market portfolio have been found to be mean-variance suboptimal relative to some ex-post determined combination of, in particular, three specific "factor portfolios," namely, the market proxy itself, a portfolio long in small firms and short in large firms, and a portfolio long in value stock and short in growth stock. Cooper, Gutierrez and Marcum (2002) have recently shown that if one uses only information in prior returns to determine optimal combinations, Fama and French's factor portfolios do not improve on buy and hold. This still leaves open the possibility that price-contingent allocation strategies may outperform buy and hold. Our model offers a theoretical rationale for this. Because the distinction between ex–ante and ex–post analysis has proved to be important, our empirical analysis will be ex–ante, i.e., to form portfolios at date t, we will use only data observed prior to t. The ex–post analysis will also be presented – only to enable comparison with Fama and French's results.

---

[7]Our empirical analysis takes the view of the econometrician, who has no information whatsoever, and verifies whether this person ought to follow price-contingent strategies. We deviate from the Hansen and Singleton [1982] methodology, though, because they assume the aggregate investor is observable and test whether he/she optimally invests. In our setting the aggregate investor is not observed.

[8]While the formal portion of our theoretical analysis is based upon a static (one-period) framework in which each investor's position is a function of his vector of signals, our empiricial analysis can be motivated by replicating the setting over time. Multiperiod extensions of the Admati (1985) model can be found in Brennan and Cao (1997) and Grundy and Kim (2002).

It deserves emphasis that we are going to compare buy and hold against a specific allocation strategy suggested by theory, as opposed to embarking on an exhaustive exercise, whereby one searches for past information that could have been used to outperform buy and hold. Without the discipline that theory imposes, such an exercise runs into the danger of data snooping.[9]

Before we carry out the empirical analysis, several issues have to be addressed. Which asset universe should we consider and what period? What securities should we include in the construction of our price-contingent trading strategy? How should we measure relative prices? How should we measure performance? And how can we tell whether the superior information is statistically significant? We will address these issues in turn.

3.1 The Data

We focus on monthly returns on U.S. common stock listed on the NYSE, AMEX and NASDAQ, as recorded by CRSP. The span of our analysis is limited by CRSP, namely, 7/1927 till 12/2000. We take the value-weighted CRSP index to be our buy-and-hold portfolio. This index has been used as the market proxy in previous empirical studies.

Against buying and holding the CRSP index, we study the performance of price-contingent portfolios. In principle, one could construct those portfolios by combining individual stocks. This would require, however, that one handle thousands of different stocks, correlating their returns to their prices, a computationally challenging exercise. A more parsimonious approach is to use groups of stocks as building blocks for our portfolios.

A natural choice for these groups of stocks is to focus on the six portfolios which have been used extensively in the empirical asset pricing literature. These are specific portfolios constructed from a double sort of the securities based on size of the issuing firms as well as the ratio of book value to market value. Together, they make up the three Fama-French factor portfolios mentioned before. We will refer to them as the six FF benchmark portfolios. Monthly returns are taken from Ken French's web site. We use the returns that are adjusted for the substantial transaction costs caused by flows of individual assets in and out of the portfolios. Such flows are the result of changes in firm size, book and market values.

Table 1 displays descriptive statistics on the monthly returns of the six FF benchmark portfolios. Portfolio 1 selects stocks of large companies with low ratio of book to market value. Portfolio 2 also selects large companies, but with medium book to market value. Portfolio 3 is

---

[9]The information we use, namely, relative prices, has never been explicitly conditioned on before. This information may have been *implicitly* conditioned on before, in particular, in momentum investment. We discuss the difference between our (price-contingent) strategy and momentum investment in the next section.

comprised of large value companies. Portfolios 4 to 6 are analogous to portfolios 1 to 3, but for small firms only. All portfolios are value-weighted. Details can be found on Ken French's website. Both the value and size effects are obvious from Table 1: the mean monthly return increases with the book-to-market ratio, and decreases with size. Notice also that the returns exhibit substantial kurtosis.

It is not obvious how to measure the relative prices on which our portfolio allocation strategy will be based. We do not have the valuations of the six FF benchmark portfolios. We opted to use the weights in a buy-and-hold portfolio of the six FF benchmark portfolios, reinvesting dividends into the FF portfolio that generates them. More specifically, let $R_{i,t}$ denote the rate of return on FF benchmark portfolio i ($i = 1, ..., 6$) over month t. ($t = 1$ corresponds to 7/1927). Let $p_{i,t}$ denote our measure of the relative price of portfolio i at the beginning of month t. It is computed as follows:

$$p_{i,t} = \frac{p_{i,t-1}(1 + R_{i,t})}{\sum_{j=1}^{6} p_{j,t-1}(1 + R_{j,t})},\tag{1}$$

$t > 0$. We set: $p_{i,0} = 0.3, 0.25, 0.15, 0.13, 0.1$, and 0.07, respectively, for $i = 1, ..., 6$.[10] Notice that $\sum_{i=1}^{6} p_{i,t} = 1$, so our prices are effectively portfolio weights in the buy-hold portfolio that starts out with \$1 at the end of 6/1927, originally invested across the six FF benchmark portfolios as in the $p_{i,0}$s above, with dividends reinvested in the components that generated them. Our proxies for relative prices are thus weights in a value-weighted portfolio. Appendix II offers more detailed information on these weights and compares them to the weights in the CRSP index.

Figure 1 plots the evolution of our construction of relative prices over time. Notice the high level of persistence in the series. The size and value effects in stock returns cause the relative prices of small and high value firms to increase gradually, although substantial variation in the two effects is apparent.

One could be concerned about the persistence in the prices, because our portfolio allocation strategy will be based on projections of a month's returns onto the vector of prices at the beginning of the month. The properties of estimated projection coefficients are known to be unusual when the explanatory variables exhibit persistence. In particular, the significance of the projection coefficients may be spurious. If not, the persistence is actually a virtue. Standard least squares projection coefficients are known to converge faster, so that estimation error can be ignored in the inference one makes subsequently, such as performance analysis of investment strategies based on the estimated coefficients. We will come back to these issues later, when we document that there is indeed persistence, but that the correlation between returns and prices is not spurious.

---

[10]These initial values are picked arbitrarily. The results are robust to changes in initial values.

3.2 Portfolio Allocation Strategy

Our portfolio allocation strategy is based on simple mean-variance optimization, which is consistent with our theoretical model. For each month in the sample, referred to as the *target month*, we determine the composition of the portfolio that promises the highest expected return for a volatility equal to that of the benchmark CRSP index. In accordance with our theory and extant empirical studies, short-sale constraints are not imposed.

Determining this portfolio requires estimating expectations and variances. We follow our theory and estimate mean returns by projecting returns onto prices. Variances and covariances are estimated from the errors of these projections. The resulting portfolio, therefore, implements an optimal, price-contingent allocation strategy.

To determine the optimal portfolio for any target month, we use observations from the sixty-month period *prior to* the target month. That is, our analysis is entirely ex–ante, i.e., only based on information that investors had available at the beginning of the target month.

Generalized Least Squares (GLS) was used to estimate the coefficients in projections of returns onto prices, to adjust for the substantial autocorrelation in the error. It sufficed to adjust for first-order autocorrelation. No further adjustments were made, although one obviously could think of many potential improvements (Iterated Least Squares, higher-order autocorrelation in the error term, autoregressive heteroscedasticity, etc.).[11]

3.3 Performance Evaluation

After obtaining the optimal portfolio for each target month, we determine whether it outperforms buying and holding of the CRSP index (our market proxy). Note that our testing strategy is similar to the well-known Fama-MacBeth [1973] strategy to test the CAPM: both are two-step procedures, whereby information over the prior sixty months generates the input for the second step.[12] The second step[13] is executed for each month ("target month") in the sample – starting obviously at month 61. Testing[14] is based on the time series of target months. Because both

---

[11]Since the error terms in the return-price projections will be correlated, one may want to use Seemingly Unrelated Regressions (SUR). Because the regressors are the same for each of the six projections, however, SUR boils down to ordinary Least Squares.

[12]In our case, the information generated by the prior sixty months is an optimal portfolio; in the case of Fama-MacBeth, betas constitute the information.

[13]Namely, comparison of the performance of the optimal portfolio relative to the market, in our case; cross-sectional projection of returns onto betas, in the case of Fama-MacBeth.

[14]The test verifies whether the optimal portfolio beats the market, in our case; whether the intercept and slope coefficients in the monthly projections are significant, in the case of Fama-MacBeth.

procedures use a fixed window of sixty months of prior information as input for the second (testing) step, nonstationarities are accommodated.[15]

With mean-variance preferences, the Sharpe ratio (ratio of average excess return over volatility) is the appropriate performance measure to determine whether our optimal portfolio outperforms buying and holding of the CRSP index. Because our optimal portfolio is constrained to generate the same (historical) volatility as the CRSP index, the comparison of Sharpe ratios boils down to a comparison of mean returns. This facilitates statistical inference: a test of the significance of the difference in Sharpe ratio is merely a test of differences in mean returns, i.e., a standard $z-$test.

We investigate subperiods of ten years, but our performance plots allow the reader to gauge the influence of any single month on the overall significance. That is, we report *partial z-statistics*, from which the influence of outliers can be gauged, and from which significance levels can be deduced for any subsample.[16] The partial z-statistics are computed as follows. Let R_t^M denote the return on the CRSP over month t. Let R_t^o denote the month-t return on our optimal portfolio with the same volatility as the market. For a sample that starts at T_1 and ends at T_2, the partial z-statistics are computed from the partial sums of the difference between the return on the optimal portfolio and that on the market:

$$z_{T_1,T_2,t} = \frac{1}{\sqrt{T_2 - T_1}} \sum_{\tau=T_1+1}^{t} \frac{R_\tau^o - R_\tau^M}{\sigma}.$$

We estimate σ as

$$\hat{\sigma} = \sqrt{\frac{1}{T} \sum_{\tau=1}^{T} \{(R_\tau^o - R_\tau^M) - \frac{1}{T} \sum_{\tau=1}^{T} \{(R_\tau^o - R_\tau^M)\}^2}.$$

The partial z-statistics form a stochastic process on $[T_1, T_2]$, so they are easy to visualize. The functional central limit theorem predicts that, in large samples (meaning $T_2 - T_1 \to \infty$),

$$z_{T_1,T_2,t} \sim W(\frac{t - T_1}{T_2 - T_1}),$$

where W denotes a standard Brownian motion on $[0, 1]$. Note that the usual z-statistic over $[T_1, T_2]$ has $t = T_2$ and hence,

$$z_{T_1,T_2,T_2} \sim W(1),$$

---

[15]In our case, we can accommodate time variation in the return-price relationship, variances and covariances. In the case of Fama-MacBeth, nonstationarities in the betas are captured. We follow the tradition of using a sixty-month window. This is obviously arbitrary, but rudimentary experimentation with alternative window lengths (in particular, 12 and 120 months) produced inferior results.

[16]See Bossaerts (1995) for an earlier illustration of the use of partial z-statistics.

15

i.e., its asymptotic distribution is standard normal, in accordance with the central limit theorem. Confidence bands of 95% can readily be computed as:

$$\pm 1.97 \sqrt{\frac{t - T_1}{T_2 - T_1}}.$$

We provide plots of the partial z-statistics for $T_1 = 0$ (before the start of our sampling period, i.e., 6/1927), and $T_2 = T$ (the end of our sampling period, namely, 12/2000). That is, we report $z_{0,T,t}$. In that case, the 95% confidence intervals are given by:

$$\pm 1.97 \sqrt{\frac{t}{T}}.$$

One can compute confidence intervals starting at any $T_1 > 0$ and conditional on the partial z-statistic at that point, z_{0,T,T_1}. These derive from the fact that

$$
\begin{aligned}
z_{0,T,t} - z_{0,T,T_1} &= z_{T_1,T,t} \sqrt{\frac{T - T_1}{T}} \\
&\sim W(\frac{t - T_1}{T - T_1}) \sqrt{\frac{T - T_1}{T}}
\end{aligned}
$$

$(T_1 < t \leq T)$. Hence, the confidence interval starting T_1 and conditional on z_{0,T,T_1} equals

$$z_{0,T,T_1} \pm 1.97 \sqrt{\frac{t - T_1}{T}}.$$

We plot such conditional confidence intervals at ten-year intervals.

4 Empirical results

4.1 Our main results

The main results are displayed in Figures 2 and 3. Figure 2 shows the evolution of the difference in Sharpe ratio between the optimal price-contingent portfolio and the CRSP index. The average difference depicted is .117 while the average level of the Sharpe ratio of the index is .538, so that the price-contingent strategy adds substantially to the achievable return (the strategies have the same ex ante volatility). The Sharpe ratios are estimated as sixty-month moving averages centered around the target month. Figure 2 demonstrates that the price-contingent optimal allocation outperforms the CRSP index since the beginning of the sampling period.

Figure 3 displays the evolution of the corresponding partial z-statistic. It confirms that the outperformance has been significant. First, consider the evolution of the z–statistic from the beginning of the sample period (1927), to its end (2000). The hyperbola depicts the confidence bounds. The z–statistic crosses the confidence bound, indicating significant outperformance, as

soon as the 1930's. The final value of the statistic, at the end of the sample, reaches a highly significant 3. The gradual increase in the z-statistic indicates that the outperformance of the price-contingent strategy is not the effect of a few outliers.

Figure 3 also enables the reader to check the significance of the outperformance of the price–contingent strategy for any of the decades in our sample: 1930–1040, 1940–1950, ..., 1980–1990. The z–statistic is positive at the end of the decade in all but one ten-year subperiod; the corresponding p-level is 0.06.[17] The performance is significant at the 5% level in 3 out of 7 ten-year subperiods; the corresponding p level is less than 0.01.[18] That is, there is little doubt about the significance of the outperformance.

The results demonstrate that price-contingent allocation strategies significantly outperform buying and holding the CRSP index, confirming that prices reflect economically relevant information, while at the same time not fully revealing all of it, as in the Noisy Rational Expectations model inspired by Admati (1995) and presented in the theory portion of this paper.

4.2 Evaluation Of The Magnitude Of The Outperformance

To put the outperformance of our price-contingent strategy into perspective, Figure 4 displays the performance of a portfolio where expected returns are not estimated from correlation with prices, but simply as sample averages of returns over the sixty–month period prior to each target month. Except in one ten-year period, the procedure fails to significantly outperform the index. There are seven ten-year periods in the plot. There is as much as one chance in three of finding one or more significant periods out of seven at the 5% level. Figure 4 essentially confirms the findings of Cooper, Gutierrez and Marcum (2002), namely, when based on ex-ante *return* information only, the CRSP index has almost always been optimal in the past. In contrast, portfolios based on ex ante information on returns *and prices* do outperform the CRSP index.

While results based on a pure ex–ante analysis are more convincing (they rule out spuriousness, among other things), it has been traditional in empirical asset pricing to present only in-sample, i.e., ex–post results. In a stationary world, the distinction between ex–ante and ex–post is without consequence. But practice reveals that there usually is a substantial difference. This is also the case in our context, as we now demonstrate.

We first present results whereby we replicate our price-contingent portfolio strategy, but

---

[17]This p-level is based on a simple binomial test evaluating the probability of at least x positive outcomes (performance) in n independent trials (periods) when the probability of a positive outcome is 0.5. In the above, $x = 6$ and $n = 7$.

[18]This p-level is based on a simple binomial test evaluating the probability of at least x rejections in n independent trials (periods) when the probability of a rejection is 0.05. In the above, $x = 3$ and $n = 7$.

use ex–post information on the relationship between returns and prices. Specifically, consider returns in the sixty months centered around the target month, *excluding* the target month itself. Project them on the relative prices (using GLS, with an AR(1) error structure, as before). Then compute the expected return as the prediction for the target month. Variances and covariances of returns remain as before, estimated from the sample moments over the sixty months prior to the target month.

Figure 5 plots the evolution of the resulting partial z–statistic. One observes a dramatic improvement in outperformance of the price-contingent portfolio over the CRSP value-weighted index. As Figure 5 demonstrates, the z-statistic quickly moves above the 95% confidence interval. The results are not the effect of a few outliers or a few specific episodes, because the increase in the partial z-statistic is gradual and steady. From the end-point of the plot, one can infer that the z-statistic over the entire sample equals 12, almost four times as high as in the ex–ante analysis (Figure 3).

To put this finding in perspective, Figure 6 displays the evolution of the partial z-statistic for the optimal portfolio rebalanced on the basis of ex–post return (not: price) information. Means, variances and covariances are computed from returns over the sixty months centered around the target month. The figure reveals that a return-based strategy outperforms the CRSP index. It confirms recent findings by Fama and French (especially Fama and French [1996] and Davis, Fama and French [2000]).[19] Note that the optimal portfolio is determined from ex–post information, but only on returns. In Figure 5, ex-post information is also used, but only about the correlation between prices and returns. Everything else is ex–ante. Still, the latter's performance is more persistent and extreme than the former's. (To facilitate comparison, the scales in Figures 5 and 6 are the same.)[20]

---

[19]Figure 6 translates Davis, Fama and French's results into direct measurement of mean-variance inferiority of the CRSP index. In Fama and French (1996) and Davis, Fama and French (2000), the well-known Gibbons-Ross-Shanken procedure is used (see Gibbons, Ross and Shanken [1989]), which verifies ex-post mean-variance optimality of the CRSP index and benchmark portfolios constructed from size and value sorts. The procedure effectively searches a combination of the CRSP index and benchmark portfolios that brings one as close as possible to mean-variance efficiency. Closeness is measured in terms of the distance from the linear relationship between mean returns and betas that characterizes mean-variance optimality. Fama and French (1996) and Davis, Fama and French (2000) document that the best combination includes other portfolios than the CRSP index, thereby rejecting its ex-post optimality. Figure 6 shows this directly. Figure 6 also replicates the results in Jagannathan and Wang [1996], where the Fama-MacBeth procedure is used to determine ex-post mean-variance optimality of some combination of the CRSP, size-based and value-based benchmark portfolios. The Fama-MacBeth procedure effectively allows monthly changes in the weights on the benchmark portfolios that bring one as close as possible to mean-variance optimality.

[20]The evidence in Figure 5 is related to studies of the conditional (full-information) CAPM, but, to our knowl-

The outperformance obtained when using ex–post information on correlation between returns and prices is substantially greater than its counterpart based on ex–ante information only. The increase is, however, spurious. It stems from trends in returns that are spuriously picked up by trends in prices, as expected when one correlates trending variables.[21]

To demonstrate this, Figure 7 replicates Figure 5, but instead of correlating returns with observed prices, we correlate them with artificial prices constructed from six independent series of simulated, normally distributed returns with monthly mean equal to $0.15/12$ and volatility equal to $0.15/\sqrt{12}$. The outperformance is substantial: the all-sample z-statistic is about 9. This confirms that part of the outperformance of the strategy based on ex–post data is spurious. But the outperformance is lower than that obtained with observed prices, where the all-sample z-statistic was 12 (see Figure 5). The difference between the two z–statistics equals about 3, which is the value of the z–statistic obtained when using ex–ante information only (see Figure 3).

That the z–statistic obtained based on ex–ante information only is significant demonstrates that the correlation between returns and prices is not spurious. We now show this directly, by studying the behavior of the error term in the projection of returns on prices and verifying that it is stationary, meaning that returns and prices will never wander away from each other indefinitely (they are said to be co-integrated), unlike with uncorrelated trending series.

Rather than running sophisticated tests of stationarity on the error term in the projections of returns onto prices, we present an intuitive and simple test that is based on the spurious correlation that emerges among trending processes and which we referred to before. Specifically, we project the error onto the simulated price series we used to obtain Figure 7. If the error term is stationary, then the projections of it onto these simulated series ought to be insignificant. We have one sixty-month series of errors per target month. Each target month therefore generates an F-statistic corresponding to the projection of the errors (of the relationship between returns and prices) onto the simulated processes. In total, there are 822 target months, and hence, 822 F-statistics. In principle, under the null of no relationship between the errors and the simulated processes, the corresponding p-values should be draws from the uniform distribution between 0 and 1. Because of the nonstationarity of the regressors, however, the distribution of the p-values will tend to be skewed to the left, with more mass on high p levels (low significance). In contrast, under the alternative that the error term is nonstationary, the histogram should be skewed to

edge, relative prices have never been employed as conditioning information (instrument). Instead, instruments such as past returns, T-bill rates or dividend yields are generally used (see, e.g., Ferson and Harvey (1999)).

[21]While the trending of prices is obvious from Figure 1, it may seem surprising to discover trends in returns, which are generally assumed to be stationary. Our empirical results prove that the assumption is false.

the right relative to the uniform distribution.[22] Figure 8 displays the histograms of the 822 p-values for the six error terms (one for each FF portfolio). The shapes of the histograms are consistent with what one expects under the null that the error term is stationary.

4.3 The Momentum Effect In Disguise?

Our price-contingent investment strategy uses specific historical information (relative prices) that may be correlated with another instrument (relative past returns) which is known to generate superior performance. It has indeed been observed that stocks whose return over the past twelve months is low relative to that of others tend to underperform, so that a strategy whereby one shorts these losers and invests the proceeds in recent winners generates superior performance. This strategy has become known as the momentum strategy. It has been analyzed in depth, among others, in Chan, Jegadeesh and Lakonishok (1996). To the extent that past performance and relative price levels are correlated, the outperformance of the momentum strategy may translate into outperformance of price-contingent strategies like ours.

The import of our findings would obviously be far less if all we accomplished was merely to exploit momentum by using a proxy for past relative performance, namely, relative prices. If this is the case, however, then a comparison of the performance of the momentum strategy against our price-contingent strategy would reveal the superiority of the former. That is, if momentum provided a cleaner signal of future expected returns while prices correlate with future returns only because prices correlate with momentum, then the momentum strategy should outperform the price-contingent strategy. The comparison is provided in Figure 9, which displays the evolution of the partial z-statistic for the difference between the return on a generalized version of the momentum strategy and the return on our price-contingent strategy. Our version of the momentum strategy does not mechanically short the losers of the past twelve months, going long the winners, but instead uses the average returns over the past twelve months and the variances and covariances estimated over the past sixty months to determine the *mean-variance optimal* portfolio. Both the momentum and the price-contingent portfolio are chosen to have the same volatility as the CRSP value-weighted index over the prior sixty months. Figure 9 demonstrates that the momentum strategy does not outperform our price-contingent strategy. In the first ten-year subperiod, it even underperforms significantly.[23] Consequently, the success of our price-contingent strategy is not merely the result of proxying for past relative returns.

---

[22]The draws are not independent: they have a moving-average structure, because there is overlap between the 822 sixty-month time series of errors.

[23]Like the price–contingent strategy, our momentum strategy outperforms buying and holding the CRSP value-weighted index. The overall z-statistic (covering the period 7/1927 till 12/2000) equals 2.5.

There is a strong relationship between the return on the momentum and price-contingent portfolios, however. The sample correlation equals 0.57, and the price-contingent return explains a significant proportion of the momentum return that is not accounted for by the return on the CRSP value-weighted index. The least squares projection coefficient equals 0.40, with a standard error of 0.03. (The least squares projection coefficient for the return on the CRSP value-weighted index equals 0.22, with a standard error of 0.04.)

The success of our price-contingent strategy is based on the structure of correlation between prices and future returns. Absence of or even negative correlation between a security's price level and its future returns could be a source for a successful price-contingent strategy. Only if prices and future returns are sufficiently positively correlated would buy-and-hold potentially become optimal. In that case, price changes would generate automatically portfolio weight changes in the same direction as the adjustments needed to keep one's portfolio optimal in the face of the changes in expected returns. But note that if correlation becomes extremely positive – recent price run-ups signal even better future returns – then portfolio weights may still have to be adjusted beyond the automatic change following the price movements. The resulting price-contingent strategy would look like momentum investment, provided price level and historical performance are correlated: one would end up buying winners, and selling losers.

In fact, we find that prices and future returns do not exhibit the extreme positive correlation that is required for successful price-contingent strategies to mimic momentum investment. We find no or even negative correlation. Below is a list of the average slope coefficients in least-squares projections of returns on the FF portfolios onto their own price. Each sixty-month estimation period prior to a target month generates one estimate. Standard errors (in parentheses) are computed as the sample standard deviation of the estimated slope coefficients. Overlap between the sixty-month estimation periods is corrected for because the sample standard deviation is computed only on the sub-sample of the slope coefficients estimated for target months that are sixty months apart. FF portfolios are identified as holding stock in big firms (B), small firms (S), high-value (H), medium-value (M) or low-value firms (L).

BL	BM	BH	SL	SM	SH
-0.02	0.30	-0.56	-1.32	-1.43	-2.08
(0.32)	(0.27)	(0.33)	(0.49)	(0.42)	(0.82)

The correlations between future returns and price levels are indeed non-positive, unlike what would be expected if our price-contingent strategy were merely momentum in disguise. This finding suggests that relative price level must refer to performance over a longer horizon than the twelve months used in successful momentum investing, and that price-contingent strategies

21

look more contrarian, exploiting mean-reversion – although our actual strategies are far more complex, as will be pointed out shortly.[24]

The above correlations provide an apt illustration of our theoretical model. They underscore that uninformed investors should use price-contingent strategies, to offset the noise in prices. For we know that if noise is absent, i.e., in the fully revealing Rational Expectation equilibrium, buying and holding is optimal, which must mean that prices do all the adjusting to changes in expected returns. This would imply that prices and expected returns are positively correlated: the more promising a security, the higher its price, and hence, the higher its weight in one's portfolio. In contrast, when there is noise, the level of correlation between prices and returns is incorrect to keep one's portfolio optimal. For instance, there may be no correlation between prices and returns. In that case, portfolio weights do not automatically change when expected returns change, and hence, the investor has to actively intervene by buying or selling in order to generate the desired portfolio weight.

If the correlation is nonpositive, holding on to securities whose price has increased would entail investing too much in it – a winner's curse. The winner's curse is all the more extreme as the correlation between prices and returns becomes negative. In that case, prices increases signal a decrease in expected returns, and the investor has to sell, not only to offset the price increase, but also to adjust for the change in expected returns. That is, an active price-contingent strategy is called for, to lessen the winner's curse.

In this respect, it is not surprising to discover that significant negative correlation between prices and returns is mainly recorded for small firms, where one would indeed expect the winner's curse to be worst. This fact provides additional, albeit indirect, evidence in favor of our model.

The price-contingent strategy that we actually implemented was more complex than the contrarian strategy suggested in the foregoing discussion, because it was based on multivariate correlation analysis between a security's future returns and *all* relative prices, not just its own. Our price-contingent strategy therefore escapes simple categorization as momentum or contrarian because the complexity of the relationshop between prices and future returns. This complexity is also consistent with the theoretical model (see Admati (1985, pp. 643–646) for further elaboration).

---

[24]For recent evidence of the extent to which stock prices mean revert, see Lewellen (2002). Again, our findings are not mean reversion in disguise, because Lewellen reports that mean reversion is strongest in large stocks, whereas the above correlations would induce contrarian price-contingent strategies mostly in small stocks.

5　Conclusion

We find that a price–contingent allocation strategy significantly outperforms buying and holding the CRSP value-weighted index. This is consistent with Noisy Rational Expectations Equilibrium, where prices convey information but are not fully revealing, and where optimal portfolios need to be adjusted when relative prices change, to lessen the winner's curse caused by automatic portfolio weight changes induced by (noisy) price movements. The returns on our price–contingent investment strategy are correlated with those on the (very profitable) momentum strategy, but the latter does not outperform the former, so that the theory underlying our price-contingent strategy can be viewed as an equilibrium explanation for the profitability of momentum investment.

There is still ample scope for improving the performance of price–contingent strategies. Our results are based on rather crude groupings of stocks. Less aggregate groupings should be contemplated, as well as other groupings (e.g., industry-based portfolios). Our estimation of the correlation between returns and prices is based on simple linear generalized least squares. We did not investigate more sophisticated specifications or estimation strategies, such as nonlinear least squares or conditional heteroskedasticity. No attempt was made to estimate the optimal window size on which to estimate the correlation between prices and returns. Refining the statistical analysis along those and other lines may yield more powerful information extraction and consequently superior performance.

The significant outperformance we uncover suggests that the price–contingent investment approach is a valuable complement to fundamental and quantitative investment analysis. It should be emphasized that our results are out of sample, so that the outperformance we obtain is based on information that was available to the investors at the time portfolio allocation decisions had to be made. Our results suggest that value can be created not only in traditional ways, by designing optimal portfolios (quantitative investment analysis) or estimating cash flows (fundamental investment analysis), but also by studying price formation in the marketplace and using the results to infer information about future returns that only competitors observe directly. Our setting provides a reconciliation between the philosophies of active and passive portfolio management as investors tilt their portfolios in favor of the assets for which they are particularly optimistic and in that sense follow active strategies.

References

Admati, A., "A Noisy Rational Expectations Equilibrium for Multiple Asset Securities Markets." *Econometrica* 53 (1985): 629–657.

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

Brennan, M., and H. Cao, "International Portfolio Investment Flows," *Journal of Finance* 52 (1997): 1851-1880.

Chan, L., N. Jegadeesh, and J. Lakonishok, "Momentum Strategies," *Journal of Finance* 51 (1996): 1681-1713.

Cho, J.-W., and M. Krishnan, "Prices as Aggregators of Private Information: Evidence from S&P 500 Futures Data," *Journal of Financial and Quantitative Analysis* 35 (2000): 111-126.

Cooper, M., R. Gutierrez, and W. Marcum, "On the Predictability of Stock Returns in Real Time," *Journal of Business* 75 (2002), forthcoming.

DeMarzo, P., and C. Skiadas, "Aggregation, Determinacy, and Informational Efficiency for a Class of Economies with Asymmetric Information," *Journal of Economic Theory* 80 (1998): 123-152.

Davis, J., E. Fama, and K. French, "Characteristics, Covariances, and Average Returns: 1929 to 1997," *Journal of Finance* 55 (2000): 389–406.

Diether, K.B., C.J. Malloy, and A. Scherbina, "Differences of Opinion and the Cross-Section of Stock Returns," *Journal of Finance* 57 (2002): 2113–2141.

Fama, E., and K. French, "Multifactor Explanations of Asset Pricing Anomalies," *Journal of Finance* 51 (1996): 55–84.

Fama, E., and J. MacBeth, "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy* 81 (1973): 607–636.

Ferson, W., and C. Harvey, "Conditioning Variables and the Cross Section of Stock Returns," *Journal of Finance* 54 (1999): 1325–60.

Gibbons, M., S. Ross, and J. Shanken, "A Test of the Efficiency of a Given Portfolio," *Econometrica* 57 (1989): 1121-52.

Grundy, B., and Y. Kim, "Stock Market Volatility in a Heterogeneous Information Economy," *Journal of Financial and Quantitative Analysis* 37 (2002): 1-27.

Grossman, S., and J. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review* 70 (1980): 393–408.

Hansen, L., and K. Singleton, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica* 50 (1982): 1269-86.

Hellwig, M., "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory* 22 (1980): 477-498.

Huang, C., and R. Litzenberger, *Foundations for Financial Economics.* New York: North–Holland (1988).

Jagannathan, R., and Z. Wang, "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance* 51 (1996): 3-53.

Kalay, A., and A. Wohl, "The Information Content of the Demand and Supply Schedules of Stocks," Tel Aviv University working paper (2001).

Kandel, E., and R. Stambaugh, "On The Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance* 51 (1996): 385-424.

Lewellen, J., "Temporary Movements in Stock Prices," MIT Sloan School working paper (2002).

Lo, A., and J. Wang, "Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory," *Review of Financial Studies* 13 (2000): 257-300.

Lo, A., and J. Wang, "Trading Volume: Implications of an Intertemporal Capital Asset Pricing Model," working paper, MIT, October 5, 2001.

Roll, R., "A Critique of the Asset Pricing Theory's Tests. Part I: On Past and Potential Testability of the Theory," *Journal of Financial Economics* 4 (1977): 129–176.

Wilson, R., "The Theory of Syndicates," *Econometrica* 36 (1968): 119-132.

Appendix I: Proof of Proposition 1

Integrating the first-order condition across agents, using the market–clearing condition and the assumption that agents have identical precisions, we obtain the following equality:

$$z_i = \frac{\int \rho_a E(f_i|I_a)da - \rho p_i(1+r_f)}{V(f_i|I_a)} - \sum_{j \neq i} z_j \frac{cov(f_i, f_j|I_a)}{V(f_i|I_a)},$$

where ρ is the average rate of risk tolerance: $\int \rho_a da$.

Hence, the equilibrium price of asset i is:

$$p_i = \frac{1}{1+r_f}\left[\int \frac{\rho_a}{\rho}E(f_i|I_a)da - \frac{z_i V(f_i|I_a)}{\rho} - \frac{\sum_{j \neq i} z_j cov(f_i, f_j|I_a)}{\rho}\right].$$

It can be rewritten as:

$$p_i = \frac{1}{1+r_f}\left[E^m(f_i) - \frac{\sum_{j=1}^{N} z_j cov(f_i, f_j|I_a)}{\rho}\right],$$

which directly yields the price equation stated in the proposition.

To rewrite this equilibrium price function in terms of returns divide both sides by the price. After simple manipulations this leads to:

$$E^m(r_i) - r_f = \frac{1}{\rho}cov^m(r_i, f_m|I_a).$$

Applying this equation to the portfolio generating f_m (the market portfolio):

$$E^m(r_m) - r_f = \frac{1}{\rho}cov^m(r_m, f_m|I_a) = \frac{p_m}{\rho}V^m(r_m|I_a).$$

Hence:

$$\frac{1}{\rho} = \frac{E^m(r_m) - r_f}{p_m V^m(r_m|I_a)}.$$

Substituting in the equation for $E^m(r_i)$ the expected return condition stated in the proposition directly obtains.

QED

Appendix II: Comparison between our proxies for relative prices and weights in the CRSP index

While our proxies for relative prices are weights in a value-weighted portfolio, they differ from the weights in the CRSP value-weighted index. First, the initial weighting is relatively arbitrary and unrelated to the CRSP weights (but the choice does not affect our empirical results). Second, the CRSP weights are determined from the relative valuations of the component stock, and not as buy-and-hold weights whereby dividends are reinvested in the stock that generates them. CRSP effectively re-invests dividends in all stock, proportional to the relative valuations of the stock. Third, the CRSP index is periodically extended through new issues, and it shrinks when firms go bankrupt, stock is repurchased, or merged into privately–held companies. These effects are adjusted for indirectly through similar adjustments in the FF benchmark portfolios – but such adjustments occur only on a quarterly basis.[25]

---

[25]Reinvestment of dividends, portfolio adjustments because of new issues, mergers, acquisitions and delistings imply that neither the CRSP index nor our portfolio of the FF benchmark portfolios are really buy-and-hold portfolios.

Table 1: Descriptive Statistics, Monthly Returns in percentage points, 6 FF benchmark portfolios, 7/1927 to 12/2000.

Portfolio	Mean	Standard Deviation	Kurtosis	skewness
Big, Low Value	0.965	5.465	8.3	-0.1
Big, Medium Value	1.001	5.844	19.3	1.4
Big, High Value	1.252	7.447	21.7	1.7
Small, Low Value	1.066	7.743	12.6	0.9
Small, Medium Value	1.283	7.226	18.5	1.5
Small, High Value	1.446	8.443	22.9	2.1
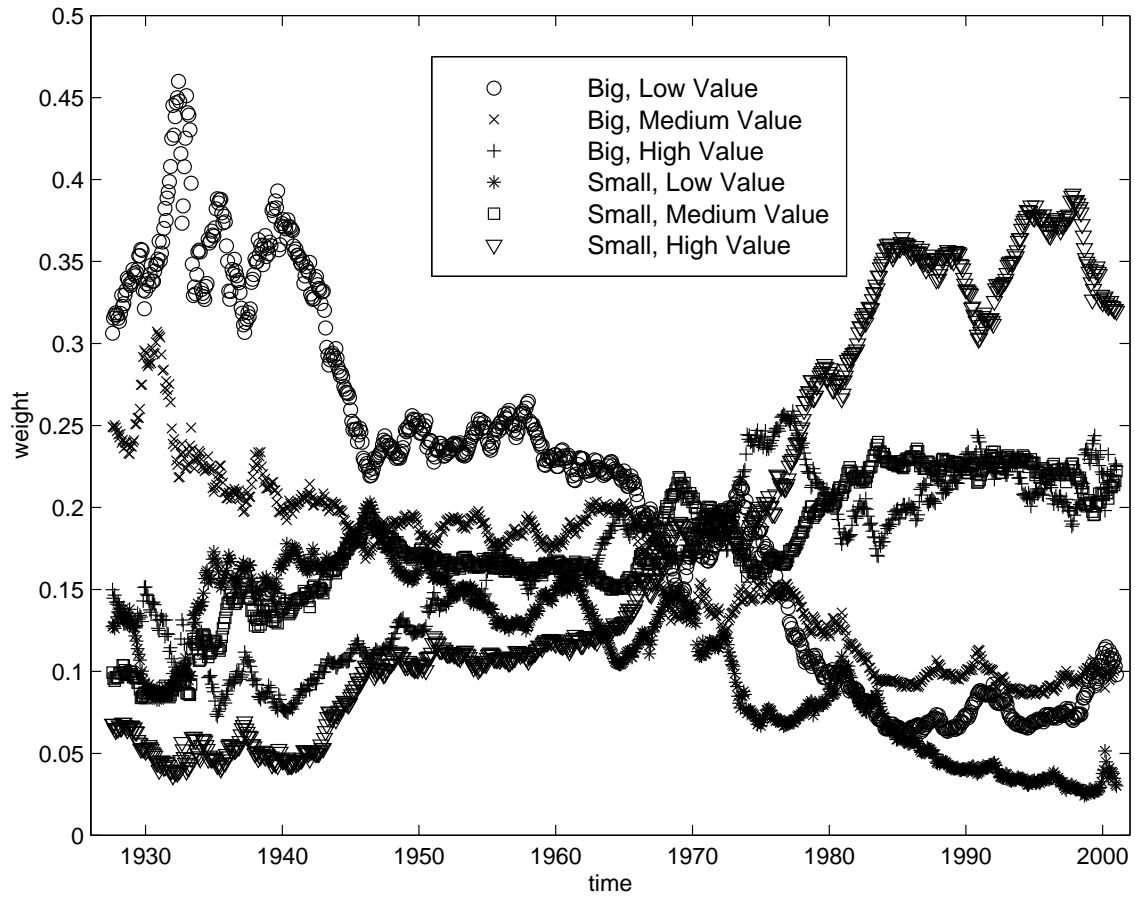
Figure 1: Evolution of the relative values of the six FF benchmark portfolios, 7/1927-12/2000. Relative values are computed as weights in a buy-and-hold portfolio, with dividends reinvested in the component portfolios that generated them. See Equation 1.
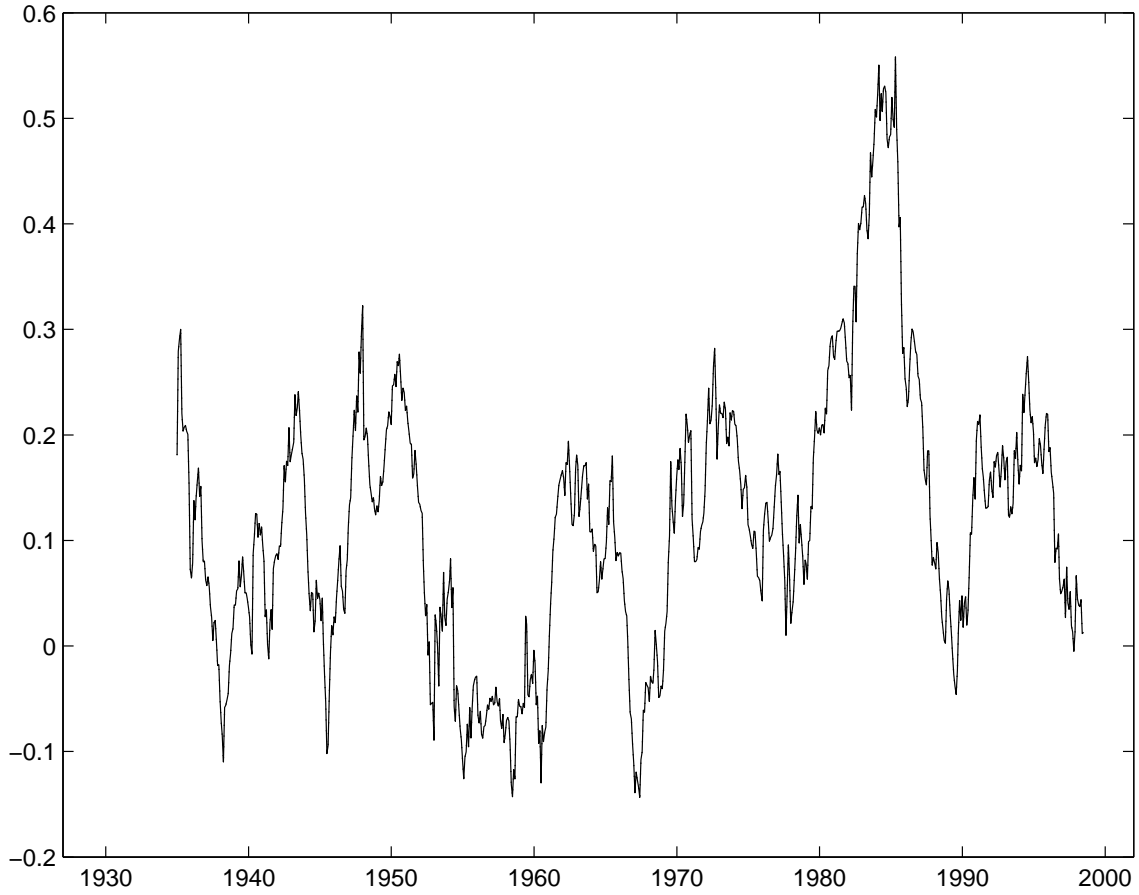
Figure 2: Evolution of the difference between the Sharpe ratios of: (i) the optimal price-contingent portfolio whereby the return-prices relationship is estimated from the sixty months prior to the target month (weights change as a function of (a) expected returns based on relative prices and the estimated price-return relationship and (b) variances and covariances of returns estimated on the sixty months prior to the target month), and (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same ex–ante volatility as the CRSP index. The difference in Sharpe ratios is estimated on the basis of a moving, fixed-length window of sixty months centered around the target month.
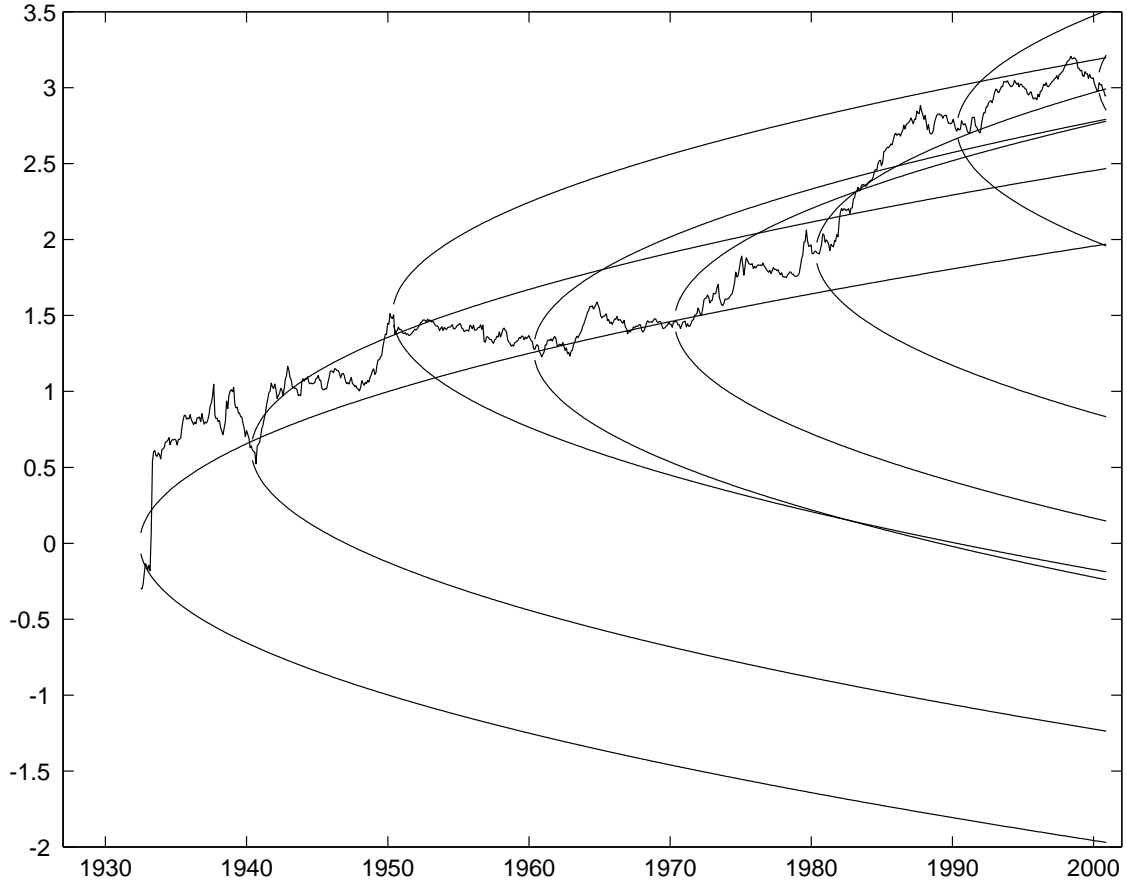
Figure 3: Evolution of the partial z-statistic of the difference in return between (i) the optimal price-contingent portfolio whereby the return-prices relationship is estimated from the sixty months prior to the target month (weights change as a function of (a) expected returns based on relative prices and the estimated return-prices relationship and (b) variances and covariances of returns estimated on the sixty months prior to the target month), and (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same ex–ante volatility as the CRSP index. The price-contingent portfolio outperforms the CRSP value-weighted index when the partial z-statistic is positive; the performance is significantly different from zero in a given ten-year period if the partial z-statistic moves outside the 95% confidence region bounded by the parabola anchored at the beginning of the ten-year period.
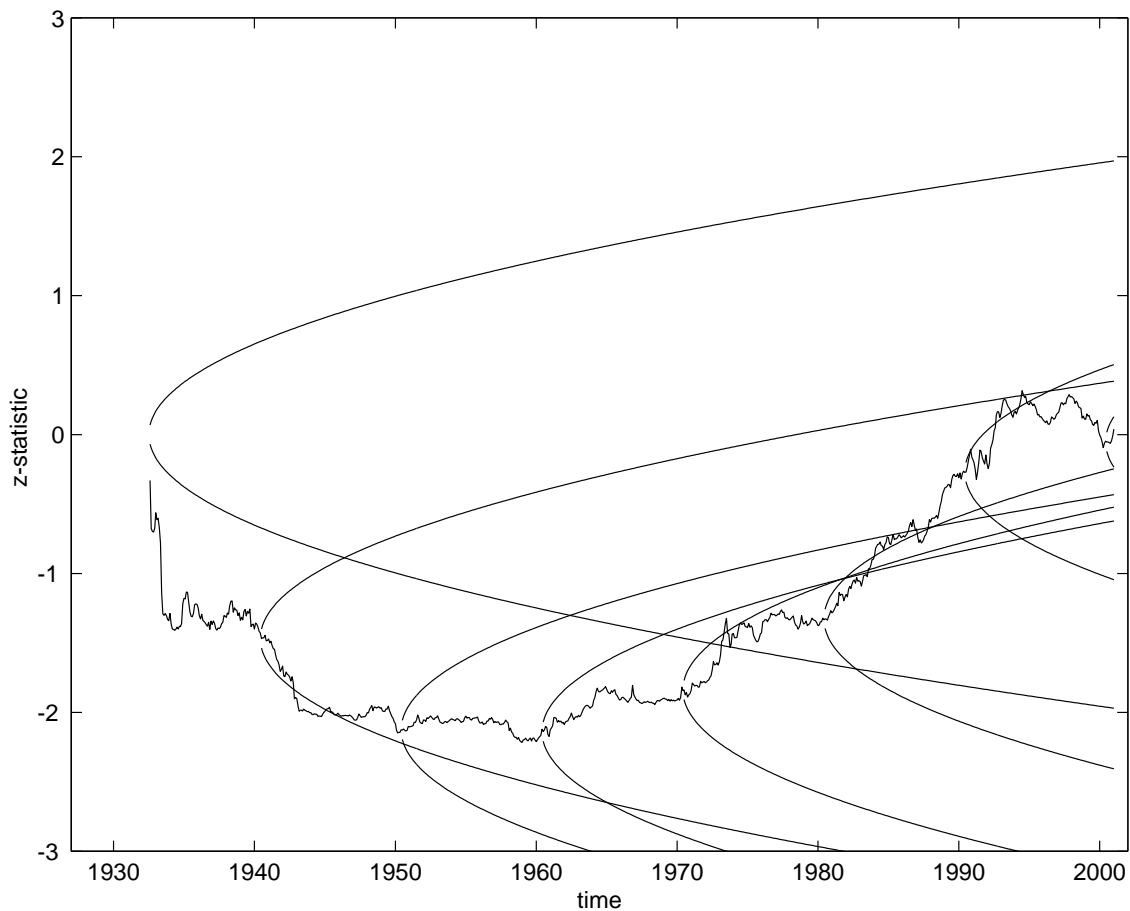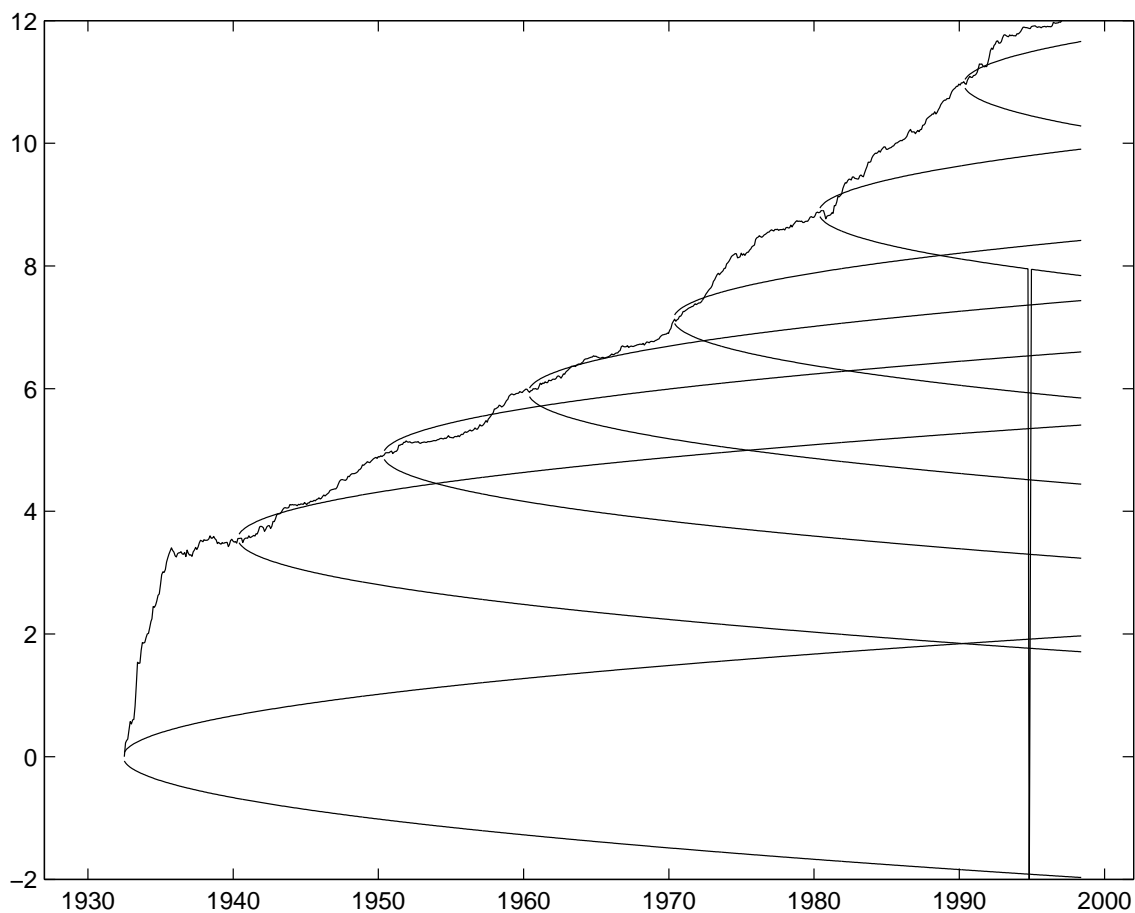
Figure 4: Evolution of the partial z-statistic of the difference in return between (i) the optimal portfolio whereby weights change as a function of averages, variances and covariances of returns over the sixty months prior to the target month, and (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same ex–ante volatility as the CRSP index. Portfolio (i) outperforms portfolio (ii) when the partial z-statistic is positive; the performance is significantly different from zero in a given ten-year period if the partial z-statistic moves outside the 95% confidence region bounded by the parabola anchored at the beginning of the ten-year period.

Figure 5: Evolution of the partial z-statistic of the difference in return between (i) the optimal price-contingent portfolio whereby the return-prices relationship is estimated on the basis of the sixty months straddling the target month (weights change as a function of (a) expected returns based on relative prices and the estimated return-prices relationship and (b) variances and covariances of returns estimated on the thirty months prior to the target month), and (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same volatility as the CRSP index. The price-contingent portfolio outperforms the CRSP value-weighted index when the partial z-statistic is positive; the performance is significantly different from zero in a given ten-year period if the partial z-statistic moves outside the 95% confidence region bounded by the parabola anchored at the beginning of the ten-year period.
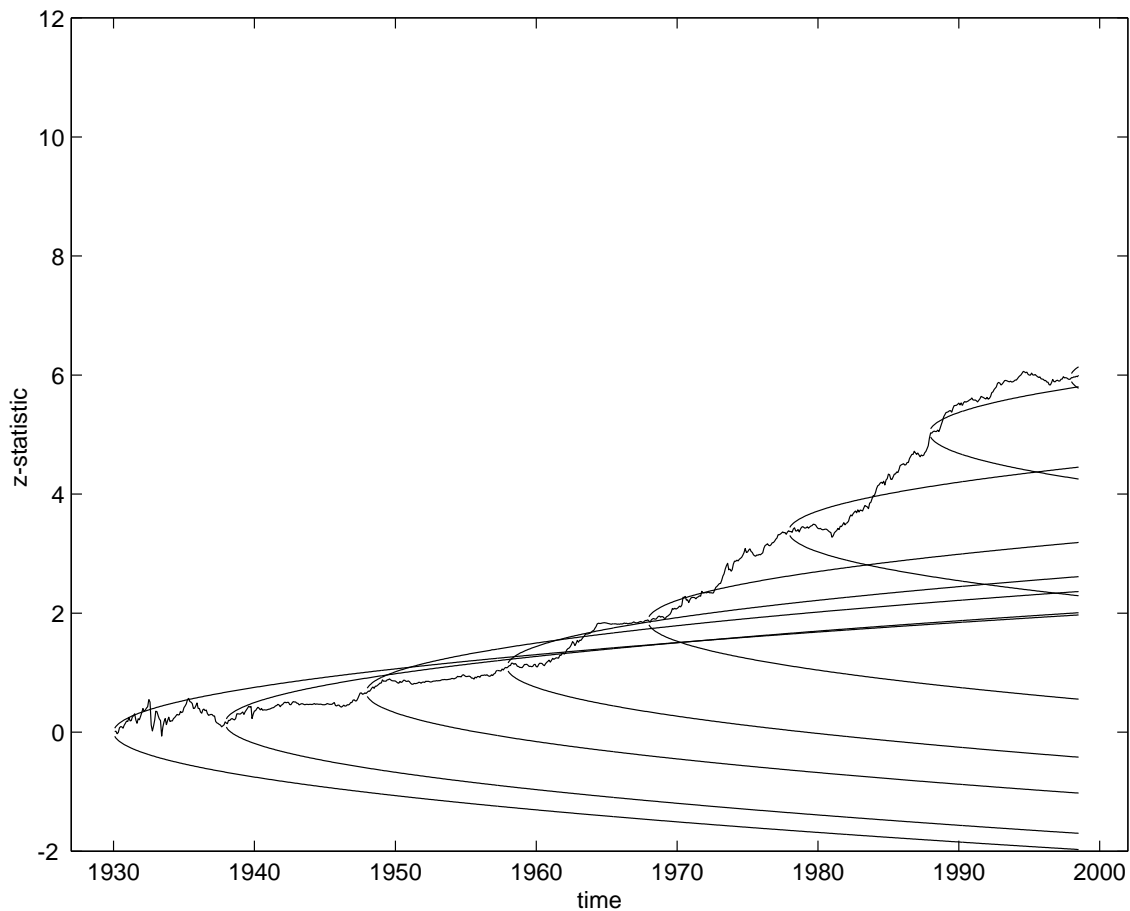
Figure 6: Evolution of the partial z-statistic of the difference in return between (i) the optimal portfolio whereby weights change as a function of averages, variances and covariances of returns over the sixty months centered around the target month, and (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same volatility as the CRSP index. Portfolio (i) outperforms portfolio (ii) when the partial z-statistic is positive; the performance is significantly different from zero in a given ten-year period if the partial z-statistic moves outside the 95% confidence region bounded by the parabola anchored at the beginning of the ten-year period.
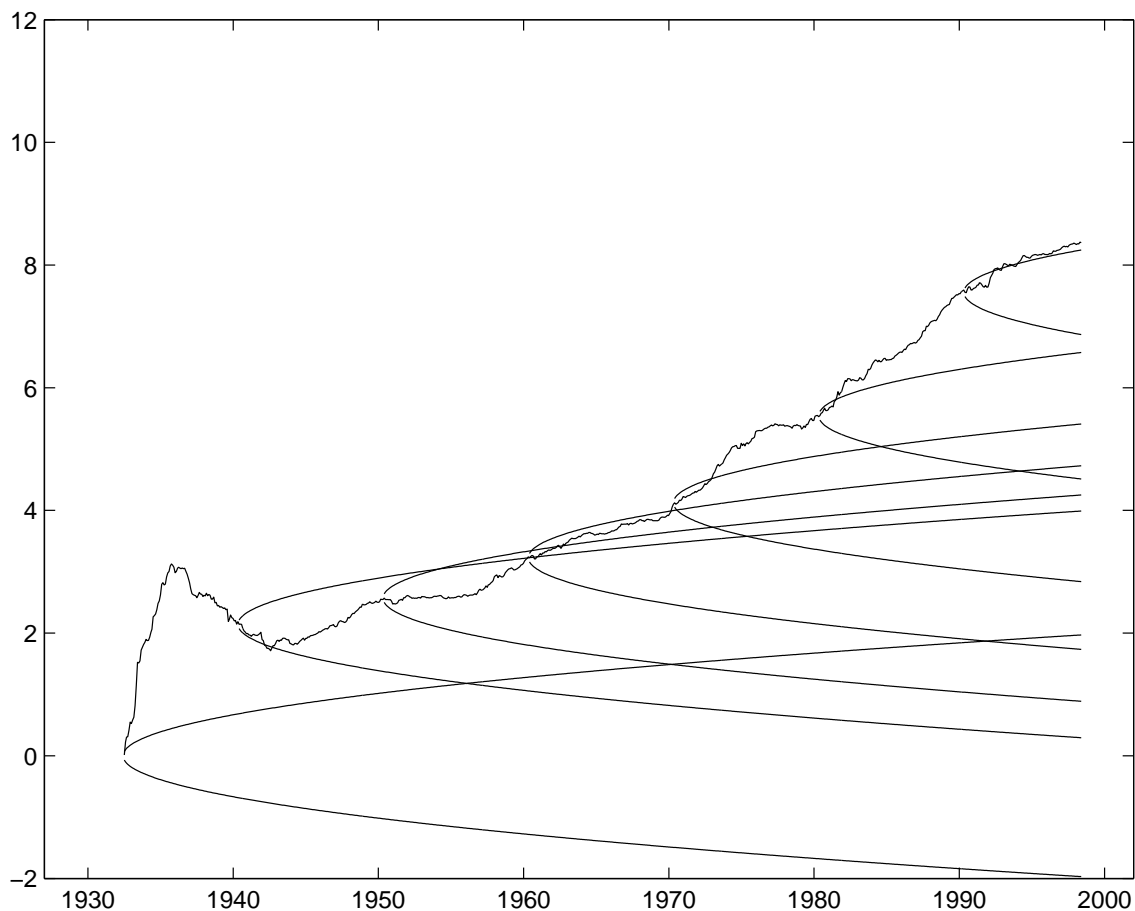
Figure 7: Evolution of the partial z-statistic of the difference in return between (i) the optimal portfolio whereby the return–simulated prices relationship is estimated on the basis of the sixty months straddling the target month (prices are *simulated*, i.e., independent of observed prices; weights change as a function of (a) expected returns based on simulated prices and the estimated return – simulated prices relationship and (b) variances and covariances of returns estimated on the thirty months prior to the target month), and (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same volatility as the CRSP index. Portfolio (i) outperforms portfolio (ii) when the partial z-statistic is positive; the performance is significantly different from zero in a given ten-year period if the partial z-statistic moves outside the 95% confidence region bounded by the parabola anchored at the beginning of the ten-year period.
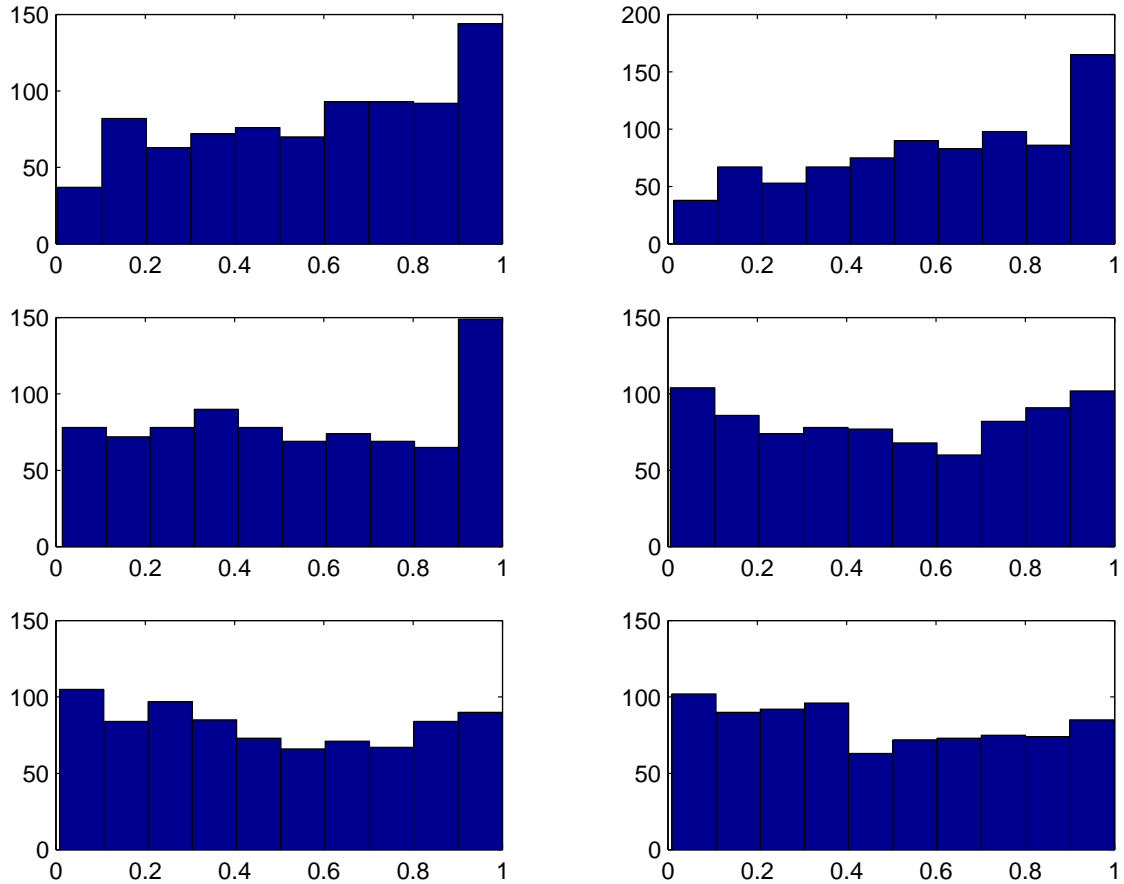
Figure 8: Histograms of the p level of the F-statistic in projections of the error term in the estimated relation between returns and actual prices onto the simulated prices used to generate Figure 7. Each sixty-month period straddling a target month generates one p level. One histogram per FF portfolio, row-wise from FF portfolio 1 to 6. Under the null that the error term is stationary, the histogram of the p levels should be skewed to the left relative to the uniform distribution (more mass at high p levels, i.e., low significance); under the alternative that the error term is nonstationary, the histogram should be skewed in the opposite direction.
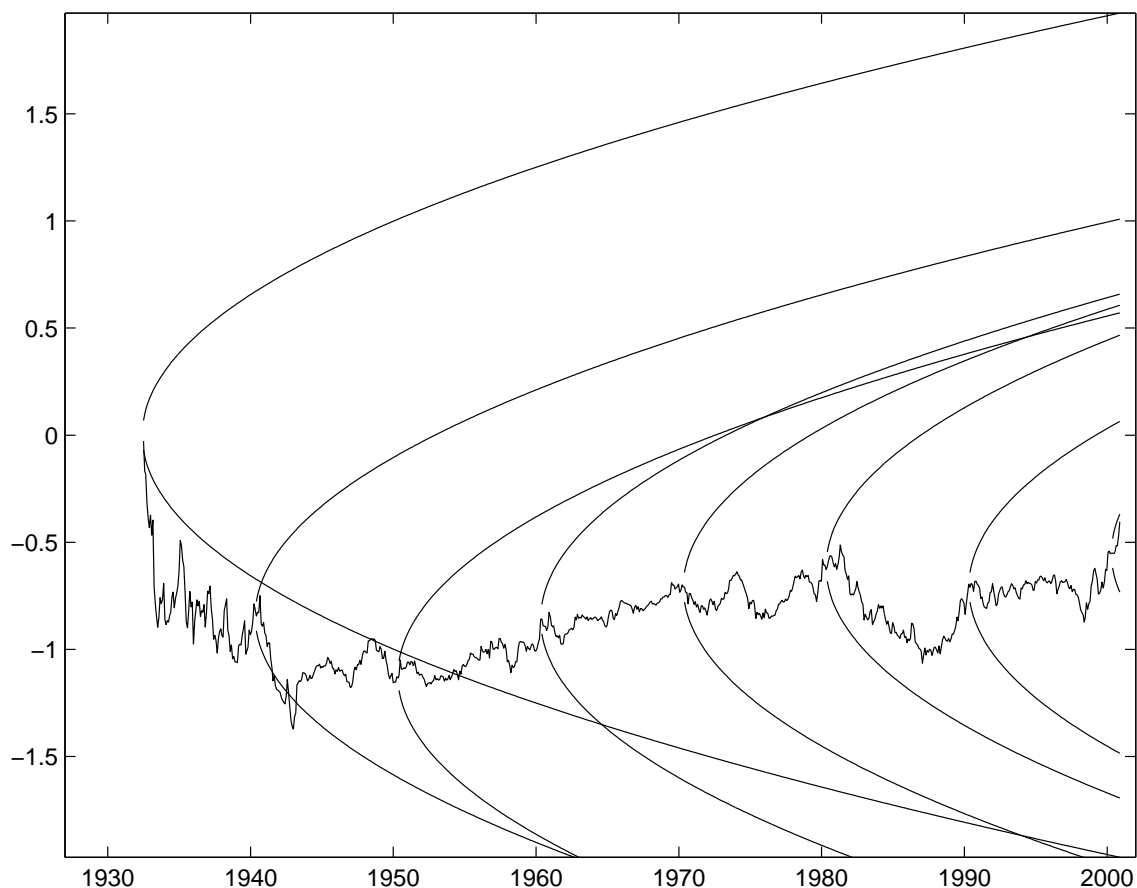
Figure 9: Evolution of the partial z-statistic of the difference in return between (i) an optimal momentum portfolio (weights change as a function of (a) expected returns estimated as the sample average over the twelve months prior to the target month and (b) variances and covariances of returns estimated on the sixty months prior to the target month), and (ii) the optimal price-contingent portfolio whereby the return-prices relationship is estimated from the sixty months prior to the target month (weights change as a function of (a) expected returns based on relative prices and the estimated return-prices relationship and (b) variances and covariances of returns estimated on the sixty months prior to the target month); 7/1927-12/2000. The optimal portfolios are chosen to have the same volatility as the CRSP value-weighted index. The momentum portfolio outperforms the price-contingent portfolio when the partial z-statistic is positive; the performance is significantly different from zero in a given ten-year period if the partial z-statistic moves outside the 95% confidence region bounded by the parabola anchored at the beginning of the ten-year period.