

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

March 22, 2002

We gratefully acknowledge financial support from the Institute for Quantitative Research in Finance (the Q-Group) and from a Moore grant at the California Institute of Technology.

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 mitigate a winner's curse. We implement empirically such portfolio allocation strategies and show they outperform economically and statistically passive/indexing buy-and-hold strategies. We thus demonstrate empirically how prices *reveal* information, in contrast with the homogeneous information CAPM. However, our findings also imply that, consistent with Noisy Rational Expectations Equilibrium, prices do *not fully* aggregate information.

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 examine some of the implications of partially revealing (noisy) rational expectations equilibria for asset pricing and asset allocation, and to gauge 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. Note that the randomness of supply shocks implies that investors do not know the exact structure 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. This 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 true value of the asset, i.e., they hold the market portfolio plus a "tilt" portfolio that reflects their private signal. Investors partially mitigate 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 investors who do not receive a private signal as they 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 first-price auction theory. In first-price auctions, bidders adjust for the winner’s curse by shaving the price they bid. In our setup, markets are competitive, and hence, price adjustments are ruled out. Instead, agents can and will make quantity adjustments.

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. In this case, price changes cause portfolio weights to change not only as a result of information on future cash flows, but also because of supply noise. This noise induces a winner’s curse which has to be offset by deliberate rebalancing.

To evaluate the empirical relevance of this theoretical analysis, we study the performance of price-contingent portfolio allocation strategies 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 prices. We use the corresponding expected returns and variance-covariance matrix to construct mean-variance optimal portfolios. We then compare the performance of these portfolios, measured by their Sharpe ratios, to that of passively buying and holding the value-weighted CRSP index. We find that the optimal price-contingent portfolios outperform economically and statistically the buy-and-hold strategy. 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.¹

It should be emphasized that the optimal price-contingent portfolio allocation strategies we analyze are 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 optimal price-contingent allocation strategies outperform the buy-and-hold indexing strategy differs from the 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

¹Interesting applications of the noisy rational expectations framework include Brennan and Cao (1997), who study the implications for international investment flows. Grundy and Kim (2002) study the implications for volatility of partially revealing rational expectations equilibrium. For an application of this type of equilibrium to the study of market microstructure, see Kalay and Wohl (2001).

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.

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: $f_i, i = 1, \dots, N$, and one riskless asset, which also serves the role of numeraire, and earns exogenous return r_f at time 1. (We adopt the 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 $\epsilon_{a,i}, i = 1, \dots, N$).
- 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 and the prices. Note that the random supply shocks imply that the market portfolio is not observed by the agents.²

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.

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).³ For simplicity, we focus on the case where all agents have identical precisions, when an "aggregate CAPM" holds.⁴

In the linear REE, all variables are jointly normal. Let p_i denote the price of asset i . Since the agents have CARA utility functions, the demand of agent a for asset i , $q_{a,i}$, is such that:

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

where I_a denotes the information set of agent a , which consists of private signals as well as prices: (Y_a, P) .

Let

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

²This unobservability of the market portfolio by the investors is central to the theoretical foundation of the model and differs from the unobservability for the econometrician discussed in the Roll Critique (1977).

³DeMarzo and Skiadas (1998) also examine CAPM issues in the context of a Rational Expectations specification. In contrast to our approach, aggregate supplies are commonly known in their model.

⁴In the general case, 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 (i.e., conditional on the signals).

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., $E(f_i | \bigcup_{a \in [0,1]} I_a)$. Also, let

$$f_m = \sum_j z_j f_j,$$

denote the total cash flow generated in the economy, equal to the sum of the cash flows generated by the different assets. Let

$$\text{cov}^m(f_i, f_m) = \sum_{j=1}^N z_j \text{cov}(f_i, f_j | I_a),$$

be the sum of the conditional covariances, $\text{cov}(f_i, f_j | I_a)$, taken from the perspective of agent a , multiplied by the realizations of the random supplies, z_j . The latter are not in the information set of the agent.

The next proposition gives the equilibrium price and expected return of asset i .

Proposition 1 *When the informed agents have identical precision, the equilibrium price of asset i ($i = 1, \dots, N$) is:*

$$p_i = \frac{1}{1 + r_f} [E^m(f_i) - \frac{1}{\rho} \text{cov}^m(f_i, f_m | I_a)];$$

the equilibrium returns equals:

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

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 $\text{cov}(f_i, f_j | I_a)$. This representative agent holds the market portfolio. Consequently, from his perspective, the standard CAPM return-covariance relationship holds. This is clear from the second part of the proposition. The view of this fictitious agent contrasts with that of the actual agents in the model, who do not observe the market portfolio, since it reflects the stochastic and unobserved aggregate endowments (see footnote 2).⁵

⁵Our representative agent effectively aggregates demands in the economy. It is well known that such Gorman aggregation obtains with negative exponential utility even if beliefs differ (see Wilson (1968), Huang and Litzenberger (1988, p. 146-148)). In the traditional aggregation analysis, however, beliefs are exogenous. We essentially show that aggregation also obtains when beliefs are endogenous, through information revealed in prices. The analysis utilizes distributional assumptions (including identical precision on private signals).

Note that our pricing relation 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 stated here holds for each realization of the random variables. The difference between our result and Admati (1985) stems from our assuming identical precision across agents.⁶

Let us now turn to the analysis of equilibrium portfolio choices.

2.3 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:

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

while his holdings of asset 2 are:

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

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:*

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

where j denotes the other asset than i , and where $\text{corr}(\cdot, \cdot)$ denotes the correlation coefficient.

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

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 is 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.

⁶In our analysis investors agree on conditional covariances, limiting the extent to which the agents look at their mean-variance pictures differently. For example, using the mean-variance geometry the minimum variance portfolio is identical for all agents and independent of the agent signals.

Importantly, the expression in the proposition reveals a winner's curse: 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 to the empirical analysis, we now show that even a marginal agent with no private signal will not find it optimal to buy and hold.

His demand can be shown to be

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

The expression reveals that the uninformed agent's demand is a function of prices. Thus, buy-and-hold will not be optimal. This is because the uninformed agent is forced to use quantity adjustments to offset the winner's curse induced by price changes onto the composition of his portfolio. Of course, he cannot fully mitigate the winner's curse. The above expression reflects that the uninformed agent faces an even stronger winner's curse problem than the informed agents: he buys more of the asset than the market portfolio would entail when he is more bullish than the average informed agent.

3 Empirical Analysis

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?⁷

To answer these questions, we compare the performance of buying and holding to portfolio allocation strategies that use information contained in prices to predict expected returns. 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 et al. (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. The ex-post analysis will be presented only to enable comparison with Fama and French’s results.

It deserves emphasis that we are going to compare buy and hold against specific allocation strategies 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. The information we use, namely, relative prices, has never been explicitly conditioned on before. *A priori*, it is not clear that it is useful to outperform buy and hold.

Before we initiate our empirical analysis, several issues have to be addressed. Which asset universe should we consider and what period? What securities should we include in the

⁷Theorists would argue that there is a simple answer: prices must not fully reveal information because otherwise information would not be collected in the first place (Grossman and Stiglitz [1980]). That of course presumes that all information is costly to gather. It remains a purely empirical issue whether this is the case.

⁸While 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 empirical analysis can be motivated by replicating the setting over time and giving investors independent realizations of the various signals and shocks over time. This would provide a recursive formulation rather than a genuinely dynamic model that incorporates hedging demands. Multiperiod extensions of the Admati (1985) model can be found in Brennan and Cao (1997) and Grundy and Kim (2002).

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 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 stock 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 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.⁹ 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 strategies will be based. We do not have the valuations of the six FF benchmark portfolios. Instead, we

⁹Value-weighting makes portfolios implementable. Equal-weighting often implies positions in small firms that are impossible to acquire in practice. Equal-weighting is used in, e.g., Fama and French [1996].

compute them as the weights in a buy-and-hold portfolio that reinvests dividends into the component benchmark portfolios that generates them. More specifically, let $R_{i,t}$ denote the rate of return on FF benchmark portfolio i ($i = 1, \dots, 6$) 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:

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

$t > 0$. We set: $p_{i,0} = 0.3, 0.25, 0.15, 0.13, 0.1$, and 0.07 , respectively, for $i = 1, \dots, 6$.¹⁰ Notice that $\sum_{i=1}^6 p_{i,t} = 1$, so our prices are effectively portfolio weights in the buy-and-hold portfolio that starts out with \$1 at the end of 6/1927, originally invested across the 6 FF benchmark portfolios as in the $p_{i,0}$ s above, with dividends reinvested in the components that generated them.

While our relative prices are really portfolio weights, 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 the empirical results we are about to report). Second, the CRSP index is periodically extended through new issues, mergers with privately held companies, and it shrinks as dividends are paid, firms go bankrupt or stock is repurchased. These effects are adjusted for only on a quarterly basis, and only within each FF benchmark portfolio separately. (Incidentally, the latter adjustments imply that the portfolios are not really buy-and-hold portfolios – still, our returns are adjusted for the transaction costs that adjustments entail.)

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 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 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 persistence indeed, but that the correlation between returns and prices is not spurious.

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

3.2 Portfolio Allocation Strategies

Our portfolio allocation strategies will be 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.).

3.3 Performance Evaluation

With mean-variance preferences, the Sharpe ratio (ratio of average excess return over volatility) is the appropriate performance measure. 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.¹¹ 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 ,

¹¹See Bossaerts (1995) for an earlier illustration of the use of partial *z*-statistics.

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:

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

We estimate σ as

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

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,

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

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,

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

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

$$\begin{aligned} z_{0, T, t} - z_{0, T, T_1} &= z_{T_1, T, t} \sqrt{\frac{T}{T - T_1 + 1}} \\ &\sim W\left(\frac{t - T_1}{T - T_1}\right) \sqrt{\frac{T}{T - T_1 + 1}} \end{aligned}$$

($T_1 < t \leq T$). 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.

3.4 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 Sharpe ratios are estimated as sixty-month moving averages centered around the target month. Figure 3 displays the evolution of the corresponding partial z -statistic.

Figure 2 demonstrates that the price-contingent optimal allocation outperforms the CRSP index since the beginning of the sampling period. The partial z -statistic plotted in Figure 3 confirms that the outperformance has been significant. The all-sample z -statistic 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. The performance is positive in all but one ten-year subperiod; the corresponding p -level is 0.06. 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. 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 we presented in the theory part of the paper.

4 Discussion

To put the outperformance of price-contingent strategies 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 strategies, but use ex-post information on the correlation between returns and prices (only). Specifically, consider returns in the sixty months centered around the target month, *excluding* the target month itself. Project the deviation of the return from the sixty-month mean on all deviations of relative prices from their respective sixty-month means. Then compute the expected return as the sample average return over the thirty months prior to target month, plus the slope coefficients in the above projection times the deviation of the weights at the beginning of the target month from the mean weights over the prior thirty months. More specifically, let t denote the target month. Let p_τ denote the vector of five of the six FF portfolio values at the beginning of month τ . Perform the following projection:

$$R_{i,\tau} - \frac{1}{60} \sum_{\sigma=t-30, \sigma \neq t}^{t+30} R_{i,\sigma} = \beta_i^t \cdot \left(p_\tau - \frac{1}{60} \sum_{\sigma=t-30, \sigma \neq t}^{t+30} p_\sigma \right) + \epsilon_{i,\tau}.$$

In the above, \cdot denotes the vector product. Compute the expected return for target month t as follows:

$$\frac{1}{29} \sum_{\sigma=t-30}^{t-1} R_{i,\sigma} + \beta_i^t \cdot \left(p_t - \frac{1}{29} \sum_{\sigma=t-30}^t p_\sigma \right).$$

Variances and covariances are estimated as the corresponding sample moments over the sixty months prior to the target month. Again, note that only the information about correlation between returns and prices is *ex-post*.

Figure 5 plots the evolution of the resulting partial z -statistic. Merely estimating correlations between returns and prices on the basis of ex-post information generates a dramatic improvement in outperformance of price-contingent portfolios 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 more than 11, 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 [2001]).¹² Note that the optimal portfolio is determined

¹²Figure 6 translates Davis, Fama and French's results into direct measurement of mean-variance inferiority of

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.)

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. This substantial increase, however, is spurious. It stems from trends in returns that are spuriously picked up by trends in prices, as expected when one correlates trending variables.¹³

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 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, the CRSP index. In Fama and French (1996) and Davis, Fama and French (2001), the well-known Gibbons-Ross-Shanken procedure is used, 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 (2001) 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.

¹³While 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.

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 the right relative to the uniform distribution.¹⁴ 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 pretty much what one expects under the null that the error term is stationary.

5 Conclusion

Consistent with noisy rational expectations equilibrium, we find that prices convey information but are not fully revealing, so that price-contingent portfolio allocation strategies significantly outperform buying and holding the index.

There is still ample scope for improving the performance of price-contingent strategies. Our results are based on rather crude groupings of stocks. 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

¹⁴The draws are not independent: they have a moving-average structure, because there is overlap between the 822 sixty-month time series of errors.

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.

Because of the trading that naturally arises in our setting with heterogeneous information, our framework can potentially be adapted to examine the empirical determination of volume.¹⁵

¹⁵Brennan and Cao (1997) explore the potential implications of the Admati (1985) model for international capital flows. An alternative perspective on the nature of trading volume is given by Lo and Wang [2000], [2001]. They focus on liquidity rather than informational motives.

References

- Admati, A., “A Noisy Rational Expectations Equilibrium for Multiple Asset Securities Markets.” *Econometrica* 53 (May 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.
- 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 (February 2000): 389–406.
- Fama, E., and K. French, “Multifactor Explanations of Asset Pricing Anomalies,” *Journal of Finance* 51 (March 1996): 55–84.
- Grundy, B. and Y. Kim, “Stock Market Volatility in a Heterogeneous Information Economy,” *Journal of Financial and Quantitative Analysis* 37 (2002).
- Grossman, S., and J. Stiglitz, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review* 70 (June 1980): 393–408.
- 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 workin paper (2001).
- Kandel, E. and R. Stambaugh, “On The Predictability of Stock Returns: An Asset-Allocation Perspective,”” *Journal of Finance* 51 (1996): 385-424.
- 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 (March 1977): 129–176.

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

Appendix: 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:

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

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

Hence, the equilibrium price of asset i is:

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

It can be rewritten as:

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

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:

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

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

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

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

**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.




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 (i) relative prices, (ii) averages, variances and covariances of returns estimated on the sixty months prior to the target month), (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 (i) relative prices, (ii) averages, variances and covariances of returns estimated on the sixty months prior to the target month), (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.




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, (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.




Figure 5: Evolution of the partial z -statistic of the difference in return between (i) the optimal price-contingent portfolio whereby the return-prices correlation is estimated on the basis of the sixty months straddling the target month (weights change as a function of (i) relative prices, (ii) averages, variances and covariances of returns estimated on the thirty months prior to the target month), (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same volatility as the CRSP index.




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, (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same volatility as the CRSP index.

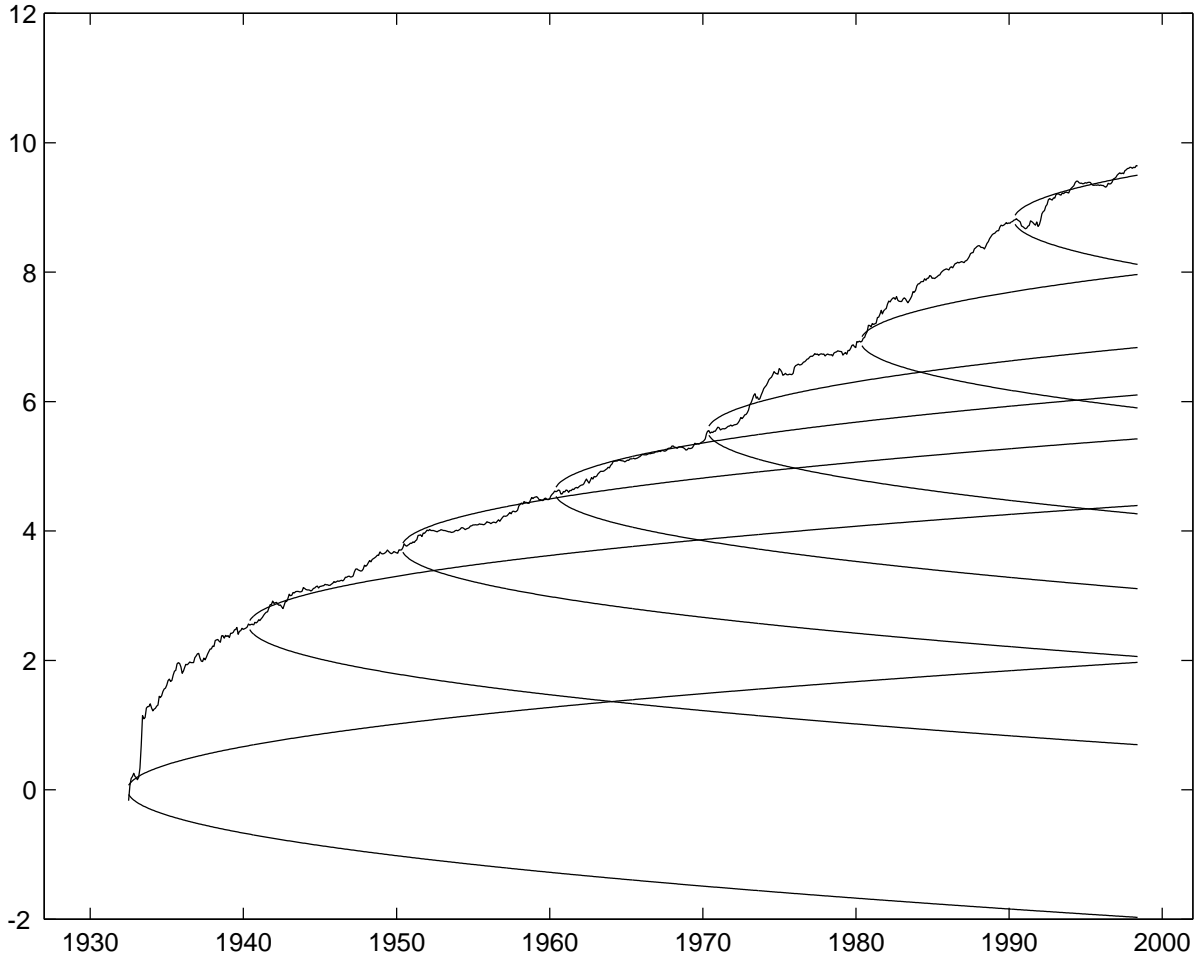


Figure 7: Evolution of the partial z -statistic of the difference in return between (i) the optimal portfolio whereby the return-prices correlation is estimated on the basis of the sixty months straddling the target month (weights change as a function of (i) *simulated* prices, independent of observed prices, (ii) averages, variances and covariances of returns estimated on the thirty months prior to the target month), (ii) prior averages, variances and covariances of returns), (ii) the CRSP value-weighted index, 7/1927-12/2000. The optimal portfolio is chosen to have the same volatility as the CRSP index.




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 to the right.