Some New Filter Rule Tests: Methods and Results

Richard J. Sweeney

Journal of Financial and Quantitative Analysis, Volume 23, Issue 3 (Sep., 1988), 285-300.

Your use of the JSTOR database indicates your acceptance of JSTOR's Terms and Conditions of Use. A copy of JSTOR's Terms and Conditions of Use is available at http://www.jstor.ac.uk/about/terms.html, by contacting JSTOR at jstor@mimas.ac.uk, or by calling JSTOR at 0161 275 7919 or (FAX) 0161 275 6040. No part of a JSTOR transmission may be copied, downloaded, stored, further transmitted, transferred, distributed, altered, or otherwise used, in any form or by any means, except: (1) one stored electronic and one paper copy of any article solely for your personal, non-commercial use, or (2) with prior written permission of JSTOR and the publisher of the article or other text.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

Journal of Financial and Quantitative Analysis is published by University of Washington School of Business Administration. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at http://www.jstor.ac.uk/journals/uwash.html.

Journal of Financial and Quantitative Analysis
©1988 University of Washington School of Business Administration

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor@mimas.ac.uk.

©2001 JSTOR

http://www.jstor.ac.uk/
Mon Feb 26 21:15:53 2001

JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS VOL. 23, NO. 3, SEPTEMBER 1988

Some New Filter Rule Tests: Methods and Results

Richard J. Sweeney*

Abstract

Mechanical trading rules seem to have more potential than previous tests found. Fama and Blume (1966), looking at the Dow 30 of the late 1950s, found no profits for the best (½-percent) rule after adjusting for transactions costs. Fifteen of these stocks looked profitable in their sample, however; for the same rule, the surviving fourteen show statistically significant profits for 1970–1982 for transactions costs obtainable by floor traders. The test used here assumes constant risk premia, or more generally, that risk premia are on average approximately the same on days "in" as for the total period.

I. Introduction

The overwhelming majority of academic financial economists subscribes to the view that financial markets are at least "weak-form" efficient. Much of the evidence on which these views are based is from serial correlation and filter rule tests of the 1960s on data from the New York and American Stock Exchanges (see Fama (1970) for a well-known review article). In the 1970s, empirical work generally dealt with specific models such as the CAPM rather than with market efficiency, as Sharpe (1970) suggested in his discussion of Fama (1970). Even when the "anomalies" literature arose later in the 1970s, these were in terms of particular models; the underying view of efficiency was maintained while the particular model was questioned or rejected (JFE "Symposium," 1978). Indeed, many viewed the anomalies as not overly troubling in cases in which there did not appear to be opportunities for excess profits once transactions costs were included—exactly the view adopted in many papers on filter rules.

The filter rule evidence for weak-form efficiency is not, however, as obvious as it seems to many. The studies of the 1960s tended to understate filter rule returns relative to buy-and-hold and to do a poor job of selecting possible winners. In addition, these tests did not have statistical confidence bounds for judging significance.

This paper develops a test of statistical significance of filter rule profits and implements the test by reexamining some of the results from Fama and Blume (1966), perhaps the best known and most influential paper on mechanical trading

* Claremont McKenna College and Claremont Graduate School, Claremont, CA 91711. Thanks are due Marc Bremer for discussion and research assistance, members of the Columbia University Finance Workshop, two anonymous *JFQA* referees, and *JFQA* Managing Editor Peter Frost.

rules. The test selects a subset of the Fama and Blume stocks that looked most promising in their work and follows these stocks from 1970 through 1982. For this sample of stocks, it appears that significant profits can be made by investors with low but feasible transactions costs. In particular, floor traders who avoid use of specialists can achieve these profits while those who pay even the lowest commercial rates very likely cannot; borderline-significant profits may exist for institutional money managers. It is important to note, however, that these results depend on the assumption that closing prices are an unbiased estimate of the prices at which purchases and sales can be made (after adjusting for the bid-ask spread).

The approach used here implies higher filter returns relative to buy-and-hold than in Fama and Blume (1966). This paper's rule considers only long equity positions while Fama and Blume have the investor short a particular security whenever he is not long in it. The short positions perform poorly, so avoiding them raises the measured returns to the filter. In addition, avoiding short positions also saves on transaction costs. These different strategies have different effects on risk, but these effects are adjusted for by using a test statistic (X) that makes a comparison of profits under the filter to those under a buy-and-hold strategy, with an expected value of X of zero. Similarly, the comparison to buy-and-hold makes the test appropriate for both rising and falling markets.

Fama and Blume looked at all 30 Dow-Jones Industrial stocks for the late 1950s and early 1960s. They averaged filter profits across stocks or filters to adjust for random successes. This paper's approach looks instead for stocks that are "winners" in one period and asks whether there is persistence so that they remain winners in later periods. The winners that can be found among the Fama and Blume stocks seem to persist as winners into later decades. Focusing only on winners, rather than the entire Dow Jones 30 as Fama and Blume did, gives superior filter results.

In fact, when the individual stocks that looked like "winners" in the Fama and Blume study for the $\frac{1}{2}$ of 1-percent rule are reexamined on daily CRSP data for the period 1970–1982, *all* gave filter rule returns that were statistically significantly better than buy-and-hold when floor-trader transactions costs of $\frac{1}{20}$ of 1 percent are used for each one-way transaction. For an equally weighted portfolio, the filter significantly beats buy-and-hold even with transactions costs as high as $\frac{3}{20}$ of 1 percent; it might even be argued that the filter outperforms buy-and-hold when transactions costs are $\frac{1}{5}$ of 1 percent, levels that institutional investors can currently hope to approximate.

The results reported below, based on the same sample of stocks used by Fama and Blume but for a later period, give a substantially different impression of weak-form efficiency than does their work. The results reported here are not by any means the strongest that can be found with the approach used (Sweeney (1987a)), but are of interest because they are so closely tied to the influential Fama and Blume piece.

Section II develops the statistical test used below, and discusses issues in its application such as time-varying risk premia and the question of including shorting in the strategy. Section III discusses transactions costs for various classes of transactors. It also discusses various biases in the measured rates of return. Section IV reviews Fama and Blume's (1966) paper and finds 15 stocks that looked like winners in the late 1950s and early 1960s for a ½ of 1-percent filter. It then asks how these stocks did in the 1970s and early 1980s under the same rule. The answer is: very well for a floor trader. Using a portfolio test, the rule arguably made borderline statistically significant profits at transaction costs that money managers can attain. Section V offers some conclusions.

II. The Filter Rule Test[1]

This section develops the test for risk-adjusted filter rule profits. The test's key assumption is that rates of return are generated by processes with constant risk premia. Note that this assumption underlies most standard tests of portfolio performance, for example, those using conventional security market lines. The test is independent of particular asset pricing models but is consistent with (a) the Sharpe-Lintner CAPM with a beta on "the" market; (b) an intertemporal Merton (1973) CAPM with betas on the market and every independent state variable; (c) a Breeden (1979) CAPM with a single beta on an asset whose return is perfectly correlated with changes in aggregate consumption (or the portfolio with returns most highly correlated with changes in aggregate consumption); and (d) an APM with (potentially) multiple-priced factors, none of which need be the market. In each case, conditional on all information available at the start of the period, the expected excess return (beyond the risk-free rate) to being "in" a particular stock is equal to the sum of betas times risk premia of underlying factors—refer to the sum of these products simply as the risk premium PR_j for the investment j. As long as these risk premia are constant, the test developed below goes through for all of the models. The latter part of this section discusses qualifications to the test when risk premia are not assumed constant.

Compare a buy-and-hold strategy over N periods (here days) with a filter strategy that has the investor "in" the asset for Nin days and "out" for Nout days, $N\text{in} + N\text{out} \equiv N$. Each day, the investor uses the filter rule to decide whether to be in or out for the next day, based only on past and current information. At the end of the N days, he judges the filter's performance by looking at the record. The mean excess rate of return on buy-and-hold is

$$\overline{R}_{BH} = \frac{1}{N}\sum_{t=1}^{N}\left(R_{jt} - i_t\right), \tag{1}$$

where R_{jt} is the rate of return on stock j in period t, and i is the risk-free rate. The mean excess return on the filter is the weighted average of the mean excess returns on days "in" and "out," or

$$\begin{aligned}\overline{R}_F &= (1-f)\frac{1}{N\text{in}}\sum_{t\varepsilon I}\left(R_{jt} - i_t\right) + f\frac{1}{N\text{out}}\sum_{t\varepsilon O}\left(i_t - i_t\right) \\ &= (1-f)\frac{1}{N\text{in}}\sum_{t\varepsilon I}\left(R_{jt} - i_t\right),\end{aligned} \tag{2}$$

[1] For a similar development, see Sweeney (1986). The relation of the test proposed here to the different discussions in Praetz (1976, 1979) are covered in Sweeney (1980).

where $f \equiv N\text{out}/N$, $(1-f) \equiv N\text{in}/N$, I is the set of days in, and O the set of days out. Thus, the filter makes the excess return of $R_{jt} - i_t$ on each day in the asset, and $i_t - i_t = 0$ on each day out of the risky asset and in the risk-free asset.

Under any of the asset pricing models listed above, assume the expected excess rate of return for each stock equals a constant risk premium PR_j. The expected returns on the two strategies are $E\overline{R}_{BH} = PR_j$ and $E\overline{R}_F = (1-f)PR_j$. Clearly, $E\overline{R}_{BH} \neq E\overline{R}_F$ as long as $PR_j \neq 0$. Hence, use the adjusted statistic

$$X = \overline{R}_F - \overline{R}_{BH} + f\overline{R}_{BH} = \overline{R}_F - (1-f)\overline{R}_{BH} \tag{3}$$

to judge filter rule profits, where $EX = 0$. Since the expected values of the excess returns are the constant PR_j, the distributions of excess returns show no serial correlation (and only contemporaneous cross correlation over assets) and the standard deviation of X is

$$\sigma_X = \sigma\left[f(1-f)\right]^{1/2}/N^{1/2},$$

where σ is the standard deviation of $R_{jt} - i_t$, assumed constant over time.

Alternatively, similarly to Fama and Blume, one could use the average rate of return on days in,

$$\overline{R}_{FI} = \frac{1}{N\text{in}}\sum_{t \varepsilon I}\left(R_{jt} - i_t\right), \quad E\overline{R}_{FI} = PR_j \, .$$

Hence, for the statistic

$$x = \overline{R}_{FI} - \overline{R}_{BH} \, ,$$

$Ex = 0$ and $\sigma_x = \sigma[f/(1-f)]^{½}/N^{½}$. It can be shown that $X/\sigma_X \equiv x/\sigma_x$, and thus use of X or x is equivalent. An advantage of X is that it measures profits per day for each of the N days, while x measures profits per day in for each of the Nin days. Section V discusses the test for an equally weighted portfolio.

Use of X compares the filter's return with a measure of its equilibrium, risk-adjusted value, where this latter is measured by the actual return to buy-and-hold. This has the advantage of removing the need to measure any betas and factor risk premia. It also handles the problem of what to do when the sample average of a factor differs from its population mean.[2]

Use of the test requires three qualifications. First, the risk premia may not be constant, an issue taken up shortly. Second, transactions costs in running a mechanical strategy must be included, and third, there may be upward biases in measured rates of return. Section III discusses these latter two issues.

A. Time-Varying Risk Premia

While the derivation of X assumed a constant risk premium PR_j, this was, in

[2] The efficient markets view, with constant risk premia, would be that $E(\overline{R}_{FI} \mid I) = E(\overline{R}_{BH} \mid I)$, where I is the information set just before the sampling period begins. If information is also provided that the sample value of $\overline{R}_{BH}$ is a, then $E(\overline{R}_{FI} \mid I \text{ and } \overline{R}_{BH} = a) = E(\overline{R}_{FI} \mid a) = a$. Thus, even when it is known that $\overline{R}_{BH} = a$ before $\overline{R}_{FI}$ is calculated, $E(x \mid \overline{R}_{BH} = a) = 0$.

part, for convenience. Let the excess returns $R_{jt} - i_t$ equal a time-varying risk premium PR_{jt} plus a stochastic term y_{jt}. Take the null hypothesis of weak-form efficiency as saying that, conditioned on all past and current prices at t, $Ey = 0$. Then,

$$x = \frac{1}{Nin}\sum_{t \epsilon I} y_{jt} - \frac{1}{N}\sum_{t=1}^{N} y_{jt} + \left(\overline{PR}_{j,Nin} - \overline{PR}_{j,N}\right),$$

where $\overline{PR}_{Nin}$ and $\overline{PR}_N$ are the sample average risk premia over the days "in" and the total period. Clearly, if the investor has no particular insight about shifts in PR or doesn't exploit such insight, then $E(\overline{PR}_{Nin} - \overline{PR}_N) = 0$, and the x (and X) test goes through even with time-varying risk premia. On the other hand, the investor might have no insight about y, so $E[(1/Nin)\,\Sigma_{t\epsilon I} y_{jt} - (1/N)\,\Sigma_{t=1}^{N} y_{jt}] = 0$, but his rule may implicitly be forecasting PR to give $(\overline{PR}_{Nin} - \overline{PR}_N) > 0$. In this case, $x > 0$ and spuriously shows profits; the investor's measured profits reflect only the above average risk on days he chooses to be "in."

The results below have the investor "in" or "out" of the average stock just about once every three days. If PR_j varies but with long periodicity, then $\overline{PR}_{Nin} \cong \overline{PR}_N$ and the X test can be validly used. If PR_j has a period about three days, then a positive X could (but need not) arise through simply picking up premium shifts. (Other work using the same approach finds profits that would have to be explained by premia shifts with substantially different periodicity.)

When the x test is applied to R_j rather than the excess return $R_j - i$, the x' statistic can be used,

$$x' \equiv x + \bar{i}_N - \bar{i}_{Nin},$$

where $\bar{i}_N$ and $\bar{i}_{Nin}$ are the average of i over N and Nin. While R_j is nonstationary if $R_j - i$ is stationary and i varies, use of x' is adequate as long as $\bar{i}_N \cong \bar{i}_{Nin}$. As an empirical matter, the results reported below are unaffected by using x' rather than x. This illustrates how, if PR varies with long periodicity as i does, the x test may well be valid in face of time-varying premia.

In equilibrium asset pricing models, the overall premium on stock $j(PR_j)$ is the sum of betas on factors times the factor premia (pr) or $PR_j = \Sigma_i^K B_{ji} pr_i$, where there are K factors. Such models do not necessarily have constant PR, though very often this is assumed in empirical work. Varying PR arise because of varying B or pr. In results reported below, holdings of stocks get out of sync, in the sense that on a given day the investor is usually "in" some stocks and "out" of others. This argues that it is betas, not factor premia, that shift, since pr shifts that help the "in" stocks would hurt the "out" stocks relative to buy-and-hold (presuming all stocks have betas of the same sign for any given factor).[3]

[3] In various versions of the CAPM, betas on the market are positive for virtually all stocks. Asset pricing models with priced nonmarket factors can be divided into cases where the market is a factor and where it is not. When the market is a factor, the sum of the betas on any nonmarket factor, weighted with the same weights used in constructing the market index, must equal zero (Sweeney and Warga (1986a,b)). Hence, one might explain measured filter-rule profits by variations in nonmarket-factor premia with betas on the factors remaining constant. Results in Sweeney and Warga (1986a) suggest that premia on nonmarket factors are rather small relative to the premium on the market. When the market is not a factor, there are no obvious constraints on the sums of betas for a given factor. As an empirical matter, very often most stocks appear to have betas of the same sign on a given factor.

B. Long versus Long/Short Strategies

The X and x statistics are defined for long filter strategies, where the investor is either long in the stock or has the funds in the risk-free asset. Fama and Blume use a long/short strategy in which the investor goes short on days out. Thus, the filter's average rate of return on this long/short strategy is

$$\overline{R}_{F'} = \frac{1}{Nin}(1-f)\sum_{t\varepsilon I}\left(R_{jt}-i_t\right) - \frac{1}{Nout}f\sum_{t\varepsilon O}\left(R_{jt}-i_t\right),$$

since the excess rate of return to a short position is $-(R_{jt}-i_t)$. Alternatively, because the investor runs the strategy every day,

$$\overline{R}_{F'} = \frac{1}{N}\left\{\sum_{t\varepsilon I}\left(R_j-i_t\right) - \sum_{t\varepsilon O}\left(R_{jt}-i_t\right)\right\}.$$

Forming a difference similar to x, as Fama and Blume (1966) do, gives $\overline{R}_{F'}-\overline{R}_{BH}$. Taking expectations, $E(\overline{R}_{F'}-\overline{R}_{BH}) = (1-2f)PR_j-PR_j = -2fPR_j$. Thus, the Fama and Blume measure is heavily biased against finding a successful filter rule (cf. Praetz (1976)). Note that if $f \approx 0.5$, the expected rate of return to the long/short strategy $E\overline{R}_{F'}$ is approximately zero. Thus, the cases in which Fama and Blume found large positive filter returns, as in Table 1, are strong indications of profits.

The X statistic can be altered in the long/short case to

$$X' \equiv \overline{R}_{F'} - (1-2f)\overline{R}_{BH},\ EX' = 0.$$

It can be shown that $X'/\sigma_{X'} = X/\sigma_X$ before transactions costs, and hence tests of the long strategy reported in Table 1 have exactly the same results as those of the long/short strategy.

C. The Filter Rule

The statistical tests developed here are applicable for any buy-and-sell rule. The particular filter rule used here is similar to that of Alexander (1961, 1964) and Fama and Blume (1966): "Buy when the stock's price rises Y percent above its past local low and sell when it falls Z percent below its past local high."[4] The rule for a given stock starts the investor in the stock. If, at the close of the day, the price is low enough to trigger the Z-percent sell rule, the investor sells the stock at the closing price and invests the proceeds in the risk-free asset. He then follows the stock until a Y-percent buy signal at the close of a day, at which time funds are taken from the risk-free asset to purchase the stock at the closing

[4] The "price" series was calculated as $P_t = \pi_{j=1}^{t}(1+R_j)P_0$. This was necessary so that, for example, a stock's going ex-dividend would not trigger the filter (cf. Fama and Blume (1966), p. 232). The R_t are from the CRSP daily returns file. Missing observations were interpolated; the only ones missing were two for Allied and one for General Motors, and careful checking showed the interpolation made no difference to results.

price.[5] Thus, funds are always fully invested, sometimes in the stock, sometimes in the risk-free asset.

III. Transactions Costs

Fama and Blume (1966) argue that, while some of the filters they examine give returns above buy-and-hold, inclusion of transactions costs eliminates any profits.

In the filter strategies considered here, transactions costs consist of all the (opportunity) costs of buying and selling the stock and of getting in and out of the risk-free asset. Three classes of potential investors are considered: floor traders, money managers, and private transactors.

A. Floor Traders

For floor traders, the most important transactions cost, according to Fama and Blume, "is the clearing house fee, which varies according to the price of the stock but averages approximately [1/10 of 1] percent on each complete transaction (i.e., purchase plus sale . . .)" or 1/20 of 1 percent (0.05 percent) for each one-way transaction. Similarly, Phillips and Smith ((1980), p. 197) estimate an average of 2 cents per share for a floor trader; with average prices of \$30 to \$40 per share, this gives costs of 1/20 to 1/16 of asset value.

This is an overstatement of costs for the period after 1976. Beginning in 1977, the clearing house fee has been approximately \$3 per transaction. Hence, for any transaction of more than \$6,000, 1/20 of 1 percent is an overstatement of costs. Results below use 1/20 of 1 percent for the entire 1970–1982 period as a conservative assumption. Note that no account is taken of the costs of the floor trader's time and trouble.

---

[5] Settlements are not made on the day of security purchases or sales, but rather five trading days later. This means that for buy-and-hold over N periods, for the first five days $R_{BHt} = (R_{jt} - i_t) + i_t$, since the funds that will pay for the security will earn i_t for five days after the rights to R_{jt} have been bought. Similarly, for the five days after N, the investor earns $-i_t$ since he does not receive payment for selling the security for five days. Neglecting discounting and assuming i_1 and i_N are not drastically different, those two adjustments cancel. The same thing applies to the filter's initial and final positions in the security. This particular filter program starts the investor in the security; if the investor is in at time N, these adjustments both to the filter and buy-and-hold exactly cancel, and otherwise must be quite close to offsetting.

For intermediate days, the excess return on the filter should be similarly adjusted for any day on which a position is opened or closed. For example, if the initial open position is closed on day 20 and a new open position taken on day 22, the investor loses five days of interest by not receiving payment on the sale until day 25, but gains five days by not having to pay for the purchase until day 27. The effect of these two adjustments will depend on differences in interest rates over the two five-day periods and the fact that the later rates should be discounted back to that of the earlier sales.

This paper does not take explicit amount of these adjustments for intermediate days, because they should have negligible effect. First, with a small filter such as used here, purchases and sales are typically quickly reversed so the discounting effect is not important. Second, the interest rates are unlikely to be very different. Third, empirically, the behavior of the X statistic is strongly dominated by R and is quite insensitive to i. Suppose the adjustments detailed above are thought of as saying the i series used was mismeasuring the true i series; experimentation indicates this simply has negligible effects.

B. Money Managers

Some can trade through investment houses at 4–6 cents per share. With an average price of $30–$40 per share, this gives transaction costs of 1/5 to 1/10 of 1 percent for each one-way transaction.[6] Before "May Day," May 1, 1975, fixed commission schedules were much higher. This is clouded somewhat by the existence of "directed give ups" whereby the large transactor could recoup some of the excessive fees by having the broker provide services or transfer cash to third parties that provide services. Note that the filter profits reported below are fairly stable before and after May Day. For the latter part of the test period used here, Berkowitz, Logue, and Noser (1986) estimate one-way costs of 0.23 of 1 percent, including 0.05 of 1 percent due to moving the market.

C. Private Transactors

Toward the end of the period studied, Charles Schwab brokerage rate schedules implied one-way transactions costs of 4/10 of 1 percent of asset value for stocks in the $30–$40 range with round lot transactions of close to $70,000 (better rates are now obtainable).

D. The Risk-Free Asset

The other side of each transaction is getting in or out of the risk-free asset. The costs here may largely be opportunity costs. Suppose the investor habitually keeps in a sweep account any idle funds not invested in stocks. The costs of moving funds to and from a broker are those of depositing and writing checks, trivial at the margin. At another extreme, the investor habitually holds a risk-free asset with an explicit charge for transactions. Here the issue turns on his holding period; with a one-day holding period, any costs of using these assets in a technical strategy must be paid anyway and cannot be charged to the strategy.

If, however, say the holding period is one month, the cost of using the filter is the annual rate that could be made on a one-month risk-free investment versus a series of investments in overnight but otherwise comparable risk-free assets. If, for example, the one-month commitment yields 10.5 percent on an annual basis while the day-by-day expected yield is 10 percent, the cost is 1/2 of 1 percent per year, or on a 252-business-day year, approximately 0.002 percent per day. It seems likely that these transactions costs are lower in the latter part of the test period.

E. The Bid-Ask Spread

If the true price is P_t, the investor will deal with the specialist in a certain percentage of trades, raising transactions costs by about one-half of the bid-ask spread, on average, in these cases. Baesel, Shows, and Thorp (1983) report the specialist is involved in 11–12 percent of NYSE trades. If the bid-ask spread is 0.6 percent as Phillips and Smith (1980) estimate or 0.3 percent as Roll (1983b)

[6] The stocks discussed below tend to have higher than average share prices. On March 28, 1975, for example, their closing prices averaged $49.6 a share, with a high of $99 (Dupont) and a low of $10 (Chrysler).

estimates, this raises transactions costs by 0.036 percent (= (0.6 percent/2)0.12) to 0.0165 percent (= (0.3 percent/2)0.11), or, for the floor trader, from 0.05 percent to between 0.086 percent and 0.0665 percent. Note, however, that floor traders may receive a buy/sell signal early enough to "shop around" and hence avoid using the specialist if desired. The money manager with explicit transactions costs of $\frac{3}{20}$ to $\frac{1}{5}$ of 1 percent will face costs of from 0.1665 percent (= $\frac{3}{20}$ percent + 0.0165 percent) to 0.236 (= $\frac{1}{5}$ + 0.036 percent). The money manager can, however, place a limit order at P_t rather than a market order and thus avoid the higher costs of dealing with the specialist. The cost of this strategy is time, but in some percentage of cases, the trigger price will be hit long enough before the close to allow waiting.

In sum, then, transactions costs reasonably range from a minimum of $\frac{1}{20}$ of 1 percent for floor traders to possibly somewhat over $\frac{1}{5}$ of 1 percent for money managers getting best prices.[7,8]

F. Can the Investor Buy at the Closing Price?

Ideally, this study could have been done on trade-by-trade data. With these data, a price exceeding the trigger price (calculated from the percent filter and the previous local low) would be a buy signal for an investor who is out. The purchase would be calculated at one of the following observed transactions prices, with it being an open question of how soon after the signal the buyer could execute the transaction.

Since only daily closing prices were available, the study assumes that the closing price is both the price that triggers the transaction and the transaction price. This makes most sense if the trigger price is exceeded before the closing price. If the closing price is the only price that day that is above the trigger price, the transaction clearly cannot be made. Tests taking account of such "missed" transactions, however, may well be unaffected. Suppose that the average rate of return on missed transactions equals the average rate of return on days "in" that the filter signals. It is clear that x is unaffected by the "misses." Assuming the

---

7 Beebower and Priest (1970) give estimates from data on actual trades of the *net* effects of (1) having to use the market maker in some fraction of trades, and (2) moving the market against the investor. They find these net costs remarkably close to zero. Baesel, Shows, and Thorp (BST, (1983), p. 993, n.5) argue that the estimates in Beebower and Priest can be interpreted as implying that the net cost of market makers' services on a round trip averages out to zero; this is possible if market makers offer a discount at some times in order to encouarge trades (perhaps to adjust inventories) and the transactor can take advantage of this. Alternatively, BST (pp. 993–994) interpret the Beebower and Priest results as likely showing that the net cost is positive but so small that it cannot be uncovered from the data used. In either case, the trades Beebower and Priest consider evidently moved the market an indetectable amount, if at all.

Berkowitz, Logue, and Noser (1986) estimate that a one-way transaction costs a money manager an average 0.23 of 1 percent, of which 0.05 is due to moving the market. Note that the stocks considered below all trade in quite thick markets, reducing problems of moving the market.

8 The transactions costs discussed above are deducted, in principle, from the true rates of return. The measured or reported rates of return may not, however, equal true rates of return. Blume and Stambaugh (1983) discuss a bias due to closing prices not necessarily reflecting the interaction of market buy-and-sell-orders. Further, they report a bias due to nonsynchronous trading. As Sweeney (1987b) discusses, these two biases are not important for this study, both because the filter returns are adjusted for buy-and-hold (so the biases net out) and because of the large size of the Dow firms considered. Roll's (1983a) point about compounding biases is irrelevant in this case since simple returns are used (as in Blume and Stambaugh).

misses occur with the same frequency when the investor is in the stock as out, the ratio of days out to total days, f, is unaffected and, hence, so is σ_x. Thus, the test results under the null hypothesis would not be affected by this problem of missed trades (though the dollar profits from a profitable rule would be reduced).

When the price exceeds the trigger price before the end of the day, the test assumes the transactor buys at the closing price. Figure 1 illustrates the point. P_{LL} is the previous-local-low price, reached on day t_0. The trigger price is $P_T = (1+Y)P_{LL}$, where Y percent is the buy filter. The closing price on day t_1 is $P_c > P_T$. The graph assumes that, earlier in day t_1, as a best guess, the market price was P_c; this follows if percentage increments to price are zero mean and serially uncorrelated. Alternatively, the investor may, with equal likelihood, have bought above or below P_c by εP_c. Clearly EX remains zero.[9] It is important to note, however, that the only way to know how profitable a strategy is, is to try it.[10]

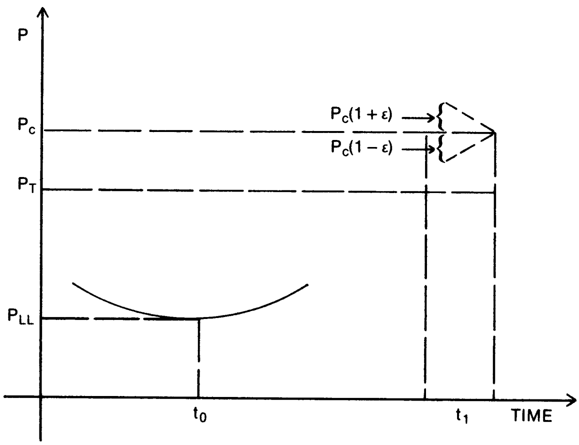

FIGURE 1
Is the Closing Price an Unbiased Estimate of the Actual Buy or Sell Price?

[9] One way of handling the problem of missed trades, caused by the closing price being the only one that day above the trigger price, is to assume all trades are made at the next day's closing price. (Note that this still leaves the next day's closing price as an *estimate* of the actual transactions price.) This turns the positive profits reported below into losses. To judge the sensitivity of the result, buying at half way between the first and second day's closing prices was also tried. To anticipate the discussion below, for zero transactions costs the portfolio's X statistic, in percent per day, is

	X
Same Day Closing Price	0.07480
Next Day Closing Price	−0.00233
Half Way between Closing Prices	0.02944.

[10] The above reasoning suggests the possibility of a downward bias for the test. If the filter rule can indeed detect trends, trading at the closing price seems likely to understate the profits to be made from the filter. With trends, a buy is likely to be recorded at a price higher than the one that triggers a Y-percent buy signal, though the investor may be able, on average, to actually buy somewhere between the trigger and closing prices, with similar opportunities for sell signals. Alternatively, using the closing price may give the floor trader time to "shop around" after a signal to avoid using market makers if their services add, on average, to transactions costs, and, similarly, the institutional investor can use limit orders for the same purpose.

IV. Later Performance of Stocks "Identified" in Fama and Blume

The results in Fama and Blume (1966) convinced many people that filter rules could not yield profits once account is taken of transactions costs. In particular, while there appeared to be profits for the ½ of 1-percent filter, they argue that transactions costs eliminate any opportunities even for the floor trader who has one-way transactions costs of only 1/20 of 1 percent.

Table 1 reproduces some results from Fama and Blume's Tables 1 and 3 for the ½ of 1-percent filter. As Panel A shows, before transactions costs are taken into account, the filter does better than buy-and-hold for 15 of the 30 stocks in the (then) Dow-Jones Industrial Average, and indeed, averaged over all stocks, the filter beats buy-and-hold by 11.5 percent per year compared to 10.4 percent. The average rate of return to the filter across all stocks, however, becomes −103.59 percent per year once transactions costs are taken into account. The 12,514 transactions come to approximately 84 per year per stock, or one every three days for a 252-business-day year.

TABLE 1
Results for ½ of 1-Percent Filter in Fama and Blume
(Percent Per Year)

Panel A: Average Returns for Individual Companies before Commissions

	F[a]	*B*[a]		*F*[a]	*B*[a]
Allied Chemical	15.5	6.8[b]	Int. Nickel	21.8	14.8[b]
Alcoa	40.1	2.5[b]	Int. Paper	20.5	1.0[b]
American Can	12.1	8.5[b]	Johns Manville	2.1	9.4
Amer. Tel. & Tel.	15.0	18.9	Owens-Illinois	0.8	11.3
Amer. Tobacco	16.5	17.0	Procter & Gamble	31.5	21.0[b]
Anaconda	28.8	4.7[b]	Sears	33.7	25.8[b]
Beth. Steel	8.2	3.2[b]	Std. Oil (Calif.)	7.6	9.3
Chrysler	3.1	0 4[b]	Std. Oil (N.J.)	3.6	7.7
Dupont	15.2	10.7[b]	Swift & Co.	1.0	4.7
Eastman Kodak	7.8	19.4	Texaco	17.2	18.8
General Elec	8.0	7.8[b]	Union Carbide	29.0	5.2[b]
General Foods	12 2	25.7	United Aircraft	−2.5	5.4
General Motors	10.7	8.8[b]	U.S. Steel	10.1	1.4[b]
Goodyear	−22.9	8.6	Westinghouse	0.8	3.8
Int. Harvester	−8.8	18.0	Woolworth	6.8	12.8
Average	11.5	10.4[b]			

Panel B: Average Return over All Companies

Before Commissions	After Commissions	Before Commissions: Long Positions	Before Commissions: Short Positions	Total Transactions
11.52	−103.59	20.89	0.97	12,514

[a] The *F* column gives the average rate of return (including both long and short positions) for the filter, *B* for buy-and-hold. Neither are adjusted for transactions costs.

[b] The filter return exceeds that of buy-and-hold.

Notes: Panel A is from Table 1, p. 229 of Fama and Blume (1966); the exact methods used in calculating the rates of return are described there on pp. 232–233. Panel B is from Table 3, p. 237, of Fama and Blume.

Fama and Blume's approach is to look at all stocks in the Dow-Jones 30 for a given time period, and then average filter results either across all stocks for a given filter (their Table 3), or across all filters for a given stock (their Table 2).

The averaging presumably reduces the importance of aberrations where a particular filter works for a given stock as a statistical fluke. The averaging can, however, serve to obscure filters that genuinely work for some but not all stocks. This suggests the approach below, of finding apparent "winners" in one period and then checking to see whether the rule and chosen stocks perform well in later samples.

In Table 1, for example, the filter's annual rate of return on Allied Chemical was 15.5 percent, versus 6.8 percent for buy-and-hold. Reducing this for 84 transactions per year, at 0.05 of 1 percent, gives a net annual filter return of 11.3 percent for the floor trader, or a net excess over buy-and-hold of 4.5 percent.[11]

This section uses the ½ of 1-percent filter rule to show that the stocks that seemed to give profits before commissions in Fama and Blume (1966) *all* give statistically significant profits in the approach used here with transactions costs of 1/20 of 1 percent, for the period 1970–1982 on daily data, using CRSP data for securities' rates of return and using the federal funds rates from the Board of Governors of the Federal Reserve System for the risk-free rate.[12] The approach here has three steps. First, securities' performances are examined to pick potential winners. The 15 stocks in which the filter beat buy-and-hold in Fama and Blume serve this function here. In this stage, the issue is not the performance of the initial portfolio chosen for scrutiny, but instead individual securities. The second step is to follow the prospective winners in subsequent periods to see how each performs. This is judged by its X statistic relative to its standard error. Third, the performance of an equally weighted portfolio of the securities in step two is judged by examining the portfolio's X statistic,[13] giving a more powerful test. A further step, not done here, is to revise the list of potential winners and test how well the revised lists do.

Table 2 shows 14 of the 15 stocks for which the filter beat buy-and-hold in Fama and Blume (Anaconda was acquired after the period of their study). Column (1) shows the X statistic, discussed above, without adjustment for transactions costs; its dimension is profits in percent per day. Column (2) shows the standard error of X. Column (3) shows the number of one-way transactions, for example 1132 for Allied; this is for 3284 business days. Column (4) shows the effect on X of adjusting for transactions. For example, for Allied with 1132 transactions, the average number of transactions per day is 0.345 (= 1132/3284). Assuming costs for a floor trader of 0.05 or 1/20 of 1 percent, then the average return per day to the filter must be reduced by 0.01725 percent (= 0.05 percent × 0.345). Hence the X value of 0.0788 percent for Allied is reduced by 0.01725

[11] The filter returns are for the combined long and short strategy, and hence tend to understate returns from a long strategy, as Section II discusses.

[12] The empirical results are quite insensitive to the proxy chosen for the risk-free rate. One way to see this is to note that for stock j, $x = \bar{R}_{j,\text{in}} - \bar{R}_j - (\bar{i}_{\text{in}} - \bar{i})$, where $\bar{R}_{j,\text{in}}$ is the sample average for R_j for days "in" and $\bar{R}_j$ for all days, and similarly for $\bar{i}_{\text{in}}$ and $\bar{i}$. As long as $\bar{i}_{\text{in}} - \bar{i} \simeq 0$, it makes no difference which i is chosen. Since the filter gets in and out of the average stock about every three days, this difference has to be very close to zero for any reasonable rate.

[13] The portfolio examined here is a simple average of the Xs of the individual stocks, exactly analogous to the averages reported by Fama and Blume (for example, their Table 3). In particular, there was no end-of-day rebalancing of portfolios in the results reported here, which would have caused higher transactions costs.

percent to 0.06155 percent. Column (5) divides (4) by (2); the significance level is below 0.05 for every security, often greatly.

TABLE 2

Tests for Individual Securities that Seemed Profitable in Fama and Blume ½ of 1-Percent Filter (Percent Per Day)

Security	(1) X	(2) Standard Error X	(3) Number of Trans- actions≠	(4) X Adjusted for Transactions Costs (0.05 of 1%)	(5) (4) ÷ (2). Adjusted X ÷ Standard Error	(6) X Adjusted for Transactions Costs (0.10 of 1%)	(7) (6) ÷ (2) Adjusted X ÷ Standard Error
Allied Chemical	0.0788	0.01675	1132	0.0615	3.67*	0.0442	2.64*
Alcoa	0.1247	0.01595	1048	0.1086	6.81*	0.0925	5.80*
American Can	0.0619	0.01170	932	0.0478	4.09*	0.0337	2.88*
Beth. Steel	0.0585	0.01621	1214	0.0401	2.47*	0.0217	1.34
Chrysler	0.0887	0.02553	1260	0.0695	2.72*	0.0503	1.97*
Dupont	0.0494	0.01299	1092	0.0328	2.53*	0.0162	1.26
General Elec	0.0576	0.01284	1044	0.0417	3.25*	0.0258	2.01*
General Motors	0.0606	0.01298	1062	0.0445	3.43*	0.0284	2.19*
Inter. Nickel	0.0828	0.01627	1168	0.0650	4.00*	0.0472	2.90*
Inter. Paper	0.1178	0.01486	1052	0.1018	6.85*	0.0858	5.77*
Procter & Gamble	0.0463	0.01108	986	0.0313	2.83*	0.0163	1.47
Sears	0.0601	0.01302	1106	0.0432	3.32*	0.0263	2.02*
Union Carbide	0.0891	0.01356	1066	0.0729	5.38*	0.0567	4.18*
U.S Steel	0.0715	0.01561	1120	0.0544	3.49*	0.0373	2.39*
Portfolio	0.0748			0.0582		0.0408	
Portfolio (per year)	18.85%			14.67%		10.28%	

≠ Total number of days: 3284.
* Significance level 0.05 or less.

These results are, of course, sensitive to the assumptions about transactions costs; 1/20 of 1 percent may be too large for the floor trade in the latter part of the period, but the opportunity costs both of getting in and out of the risk-free asset and of the trader's time and trouble may raise the costs above 1/20 of 1 percent. Columns (6) and (7) show the effects of doubling transactions costs for the floor trader from 1/20 to 1/10 of 1 percent. While all adjusted Xs are still positive, only 11 of 14 are significant.

Since returns tend to show nonstationarity of standard deviations over time for samples as large as that used here, the reader may well want to view all reported t-statistics as illustrative.

An Equally Weighted Portfolio

If transactions costs are 1/20 of 1 percent, the average X in column (4), and hence for the portfolio, is $X_p = 0.0582$. With n securities, the variance of X_p is

$$\sigma^2_{X_p} = \sum_j^n (1/n)^2 \sigma_{X_j} + \sum_j \sum_{j \neq h} \mathrm{Cov}\left(X_j, X_h\right)(1/n)^2 .$$

The covariances of the individual Xs depend on the covariances of the excess rates of return and on the covariances of the investor's positions in the stocks (Sweeney (1987b). Measured covariability of positions may arise from the covariability under the null, or may arise because the data depart from the null; the latter source should not be included in standard errors used to test the null. The author has not worked out what the covariances of positions should be under the null hypothesis. Table 3 shows standard errors assuming all covariances in positions (and hence in Xs) are zero under the null, and also assuming measured covariances contain no influence of departures from the null; the former may well understate the true standard errors, the latter overstate them. The average correlation of the Xs is 0.08 for this sample and around 0.015 over large samples of stocks and days; Table 3 also shows standard errors assuming an average correlation of 0.015 for the Xs.

Since it seems very likely that floor traders have transactions costs of less than 3/20 of 1 percent, Table 3 shows they would have made economically and statistically significant profits on a portfolio of filter strategies.[14] Money managers seem likely to face costs of at least 1/5 of 1 percent, giving profits to the portfolio of 1.89 percent per year that are only borderline significant if the covariances of the Xs are zero and are insignificant in the other cases.

Note that portfolio Xs are fairly robust across subsamples of the overall period. For example, for the periods 1970–1974, 1975–1978, and 1979–1982, the portfolio X_ps for zero transactions costs are 0.08297, 0.06602, and 0.07195, while for the overall period, the portfolio X_p is 0.07480 percent/day. Thus, profit rates for these filters seem insensitive to the changes in transaction costs associated with "May Day" or the decline in floor traders' costs.

TABLE 3

Effect of Transactions Costs on Performance of Portfolio

Transactions Costs (Fraction of 1 Percent)	(1) X_p	(2) X_p ÷ Standard Error (Assuming all Cov = 0)	(3) X_p ÷ Standard Error (Adjusted for Average Sample's Cov)	(4) X_p ÷ Standard Error (Adjusted for this Sample's Cov)
1/20	14.67%/year	14 19	12.98	10.06
1/10	10 28	10.14	9.27	7.19
3/20	6.10	6 09	5.57	4.32
1/5	1.89	2 04	1 87	1.45

Notes: The X_p in (1) are portfolio X statistics. The t-statistics in (2) are calculated assuming that all of the covariances of the Xs are zero. (3) and (4) use nonzero (typically positive) estimates of the covariances. (3) assumes the covariances are based on the typical correlation of Xs of 0.015 found over large samples of stocks and days (4) uses the actual sample covariances of the Xs for the 14 stocks studied; their correlation averages approximately 0 08.

V. Conclusions

Review of Fama and Blume (1966) shows 15 of the 30 securities they considered seem to offer potential profits for the 1/2 of 1 percent filter rule over the period 1956–1962. When the 14 available securities from this group are examined over the later period 1970–1982 with a test with statistical confidence bounds, each of these securities gives highly significant profits for a floor trader; for example, an equally weighted portfolio gives profits of over 14 percent per year. For a money manager who can reduce transactions costs to 1/5 of 1 percent,

[14] Table 3 converts from percent per day to percent per year by multiplying by 252 trading days per year.

the portfolio gives profits of a bit under 2 percent per year, with these profits arguably borderline significant (the issue turning on whether positive sample covariances of the *X*s can be attributed to inefficiency). Other work using the same approach on a larger universe of stocks, however, shows statistically significant profits for a ½ of 1 percent trading rule at transactions costs of ¼ to ⅕ of 1 percent (Sweeney (1987a)). Commercial transactors seem out of luck here.

These results are sensitive both to transaction costs and to whether the closing price is an unbiased estimate of the price at which one can buy or sell (after taking account of the bid-ask spread). Transaction costs, particularly the opportunity cost of the time and trouble of running the strategy, may be larger than assumed. Further, it is possible that one may systematically end up buying above and selling below the closing price (beyond the account taken above of the bid-ask spread). Argument cannot settle these issues. The only way to know is to try.

The interesting issue is why substantial profits still seem to be made at least by floor traders, that is, why the market seems weak-form inefficient at their level of transactions costs. One answer is that the cost of a seat on an exchange is just the (risk-adjusted) present value of the profits that could be made (Phillips and Smith (1980) make a similar argument for options markets). This does not explain, however, why current seat holders have not competed these profits to zero; there are too many of them to argue that a successful conspiracy is at work.

Another explanation is the opportunity costs involved in implementing this paper's rules and thus (as an unintended side effect) reducing the profits available. This would explain why current seat-holders have not competed away profits. Once the rule is known, however, a computer progiam that generates limit orders based on the rule can be created at trivial cost, and for any operation that already uses computers, the strategy can be implemented at negligible marginal cost.

A third argument says that the rule has been implemented by market participants to the extent justifiable on risk grounds. This view is nonoperational until a model of behavior under risk is specified. This view is untenable if one chooses any asset pricing model with constant risk premia, for under such models the expected value of X is zero, as Section II discusses. This paper has rejected an *ex ante* $X = 0$, at least at the level of the floor trader. Of course, these asset pricing models are also thereby rejected if risk premia are constant.

A fourth point often made is that attempts to capture the profits revealed above will eliminate them before the investor can execute the transactions, that is, the investor moves the market against himself while trying to consummate the trade. Scholes (1972) finds very large elasticities, but it might be argued that a successful trader would be so closely watched that his moves would shift market excess demand curves so that elasticity along a given curve is irrelevant. Beebower and Priest (1980) find little net evidence that the actual trades they examined moved the market in an unfavorable direction. Berkowitz, Logue, and Noser (1986) estimate that market pressures contribute on average 0.05 of 1 percent to the total costs of 0.23 of 1 percent faced by a money manager. Further, the 14 stocks examined here all trade in thick markets.

The basics of Markowitz portfolio theory suggest that the equally weighted portfolio used here is highly unlikely to be optimal. Thus, supposing there are inefficiencies in the market, other portfolios may consistently outperform the one used here, that is, reject the null hypothesis at higher significance levels.

All implementation requires knowing the rule and its properties. This may explain a fair portion of the profits, for the research reported here would have been quite expensive if done at commercial rates.

References

Alexander, S. S. "Price Movements in Speculative Markets: Trends or Random Walks." *Industrial Management Review*, 2 (May 1961), 7–26.

__________. "Price Movements in Speculative Markets: Trends or Random Walks, No. 2." *Industrial Management Review*, 5 (Spring 1964), 25–46.

Basel, J. B.; G. Shows; and E. Thorp. "The Cost of Liquidity Services in Listed Options: A Note." *Journal of Finance*, 38 (June 1983), 989–995.

Beebower, G., and W. Priest. "The Tricks of the Trade." *Journal of Portfolio Management* 6 (Winter 1980), 36–42.

Berkowitz, S. A.; D. E. Logue; and E. A. Noser. "The Total Costs of Transactions on the NYSE." Working Paper, Amos Tuck Business School, Dartmouth College (1986).

Blume, M. E., and R. F. Stambaugh. "Biases in Computed Returns." *Journal of Financial Economics*, 12 (Nov. 1983), 387–404.

Breedon, D. T. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *Journal of Financial Economics*, 7 (Sept. 1979), 265–296.

Fama, E. F., and M. E. Blume. "Filter Rules and Stock-Market Trading." *Journal of Business*, 39 (Jan. 1966), 226–241.

Fama, E. F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance*, 25 (March 1970), 383–417.

Merton, R. C. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, 41 (Sept. 1973), 867–887.

Phillips, S. M., and C. W. Smith, Jr. "Trading Costs for Listed Options: The Implications for Market Efficiency." *Journal of Financial Economics*, 8 (June 1980), 179–201.

Praetz, P. D. "Rates of Return on Filter Tests." *Journal of Finance*, 31 (March 1976), 71–75.

__________. "A General Test of a Filter Effect." *Journal of Financial and Quantitative Analysis*, 14 (June 1979), 385–394.

Roll, R. "On Computing Means and the Small Firm Effect." *Journal of Financial Economics*, 12 (Nov. 1983a), 371–386.

__________. "A Simple Measure of the Effective Bid/Ask Spread in an Efficient Market." Working Paper, Graduate School of Management, UCLA (1983b).

Scholes, M. "The Market for Securities: Substitution versus Price Pressure and the Effects of Information on Share Prices." *Journal of Business*, 45 (April 1972), 179–211.

Sharpe, W. "Comment" on Fama (1970). *Journal of Finance*, 25 (March 1970), 418–420.

"Symposium on Some Anomalous Evidence Regarding Market Efficiency." *Journal of Financial Economics*, 6 (June/Sept. 1978), 93–330.

Sweeney, R.J. "A Statistical Filter Rule Test: The Dollar-DM Exchange Rate." Claremont Working Paper (May 1980).

__________. "Beating the Foreign Exchange Market." *Journal of Finance*, 41 (March 1986), 163–182.

__________. "Evidence on Short-Term Trading Strategies." Claremont Working Paper (1987a).

__________. "Some New Filter Rule Tests: A More Extended Analysis." Claremont Working Paper (1987b).

Sweeney, R. J., and A. D. Warga. "The Pricing of Interest-Rate Risk: Evidence from the Stock Market." *Journal of Finance*, 41 (June 1986a), 393–410.

__________. "The Possibility of Estimating Premia in Asset Pricing Models." *The Financial Review* (1986b), 299–308.