



American Finance Association

Performance Evaluation with Transactions Data: The Stock Selection of Investment Newsletters

Author(s): Andrew Metrick

Source: *The Journal of Finance*, Vol. 54, No. 5 (Oct., 1999), pp. 1743-1775

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/222502>

Accessed: 24/11/2009 12:40

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

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.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Performance Evaluation with Transactions Data: The Stock Selection of Investment Newsletters

ANDREW METRICK*

ABSTRACT

This paper analyzes the equity-portfolio recommendations made by investment newsletters. Overall, there is no significant evidence of superior stock-picking ability for this sample of 153 newsletters. Moreover, there is no evidence of abnormal short-run performance persistence ("hot hands"). The comprehensive and bias-free transactions database also allows for insights into the precision of performance evaluation. Using a measure of precision defined in the paper, a transactions-based approach yields a median improvement of 10 percent over a corresponding factor model. This compares favorably with the precision gained by adding factors to the CAPM.

INVESTMENT NEWSLETTERS HAVE EXISTED since at least the early part of this century, and the current industry of more than 500 active letters has approximately 2 million subscribers (Hulbert (1996)). The typical newsletter is produced by a small staff and provides a wide range of advice targeted at the retail investor. Successful publications can earn millions of dollars in subscription revenue and allow their editors to become frequent speakers at investment seminars (Brimelow (1986)). This paper undertakes a detailed study of newsletters' equity recommendations. The dataset spans 17 years, is free of survivor and back-fill biases, and contains every recommended (long) transaction for 153 newsletters. These data allow for an analysis of two questions. First, do investment newsletters have stock-selection ability? Second, can transactions data be used to improve the precision of performance evaluation?

The *Hulbert Financial Digest* (HFD) has been tracking investment-newsletter recommendations since 1980. The only previous papers to use the HFD database are those by Graham and Harvey (1996, 1997) and Graham (1998), which focus on newsletters' timing between stocks and cash. The present paper is the first attempt to study the specific equity recommendations of

* Department of Economics, Harvard University and NBER. I thank John Campbell, Gary Chamberlain, Paul Gompers, John Graham, Campbell Harvey, Leslie Jeng, Susie Metrick, Jack Porter, Andrei Shleifer, René Stultz, Sheridan Titman, Richard Zeckhauser, an anonymous referee, and seminar participants at the NBER, the 1997 NBER Summer Institute, Harvard, MIT, and Chicago for helpful comments, Daniel Chen and Alex Tsai for editorial assistance, and James Choi and Tom Knox for excellent and tireless research assistance. Mark Hulbert was very generous with his time, advice, and data. This project was partially supported by the Richard Herrnstein Fund at Harvard University.

these financial advisors.¹ In this respect, the paper is similar to other studies of “expert” equity recommendations such as those by Barber and Loeffler (1993) (on *The Wall Street Journal*’s Dartboard column), Desai and Jain (1995) (on “Superstar” money managers in *Barron’s*), and Womack (1996) (on brokerage analysts). To carry out the analysis, the paper uses performance-evaluation methods developed in several recent studies of mutual funds (Mal-kiel (1995), Carhart (1997a), and Daniel, Grinblatt, Titman, and Wermers (DGTW) (1997)). Using these methods, I find no evidence of abnormal stock-selection ability in this sample of newsletters.

In addition to determining the investment value of newsletters’ stock selection advice, this paper also attempts to shed light on the power and limitations of various performance-evaluation methodologies. The transactions-level detail and bias-free construction of the HFD database allow for a comparison of results from several different methods and insight into their relative precisions. Though such comparisons can also be carried out on simulated data, such tests may miss crucial elements of actual managed portfolios and can lead to biased results. The HFD database provides a rare natural experiment. Using a measure of “precision” defined in the paper, I find that the transactions-based approach of DGTW (1997) yields a median improvement of 10 percent over the 4-factor model of Carhart (1997a), with the former approach providing more precise estimates of abnormal performance for more than 80 percent of the newsletters. This compares with a median improvement of less than 1 percent for the 4-factor model over the CAPM. These results have useful implications given the advent of transactions and holdings databases in studies of insider trading (Eckbo and Smith (1998), Jeng (1998)), of individual investors (Barber and Odean (1999), Grinblatt and Keloharju (1998)), and of mutual funds (DGTW (1997)).

Section I of the paper describes the HFD database in detail and provides summary statistics on newsletter recommendations and performance. Section II contains the analysis of stock-selection performance. Section III analyzes the consistency and relative precision of the three performance-evaluation models. Section IV looks for short-term persistence, or “hot hands,” in newsletter stock-selection ability. Section V concludes the paper. Two appendixes supplement the text: Appendix A discusses the calculation of newsletter returns and Appendix B describes the construction of the return series used in the DGTW (1997) characteristic-matching model.

I. Data

The *Hulbert Financial Digest* (HFD) has been tracking the performance of investment newsletters since July 1980. It accomplishes this task by subscribing to the newsletter print editions and regularly calling the free “tele-

¹ One newsletter contained in the HFD sample, *The Value Line Investment Survey*, has received substantial attention in the academic literature. A series of papers since the 1960s has alternatively identified and explained the “*Value Line Anomaly*,” the apparent superior performance of *Value Line*’s recommendations. See Shelton (1967), Black (1973), Copeland and Mayers (1982), Stickel (1985), Huberman and Kandel (1990), and Lewis, Rogalski, and Seward (1997) for some key papers in this debate.

phone hotlines” of the newsletters that have them. The sample includes many well-known newsletters whose editors are often quoted in the financial press or noted for their mutual fund management. Some examples (with editors given in parentheses) are *The Granville Market Letter* (Joe Granville), *Louis Rukeyser’s Wall Street* (Louis Rukeyser), *MPT Review* (Louis Navellier), *The Ruff Times* (Howard Ruff), and *The Value Line Investment Survey* (Value Line Publishing, Inc.).

Many newsletters in the HFD sample recommend more than one portfolio; in this case HFD tracks the recommendations for each portfolio separately. Although not all newsletters give specific advice about asset allocation, HFD has adopted a consistent methodology for translating vague advice into “model portfolios” for each newsletter. Details of this methodology are given in Appendix A. The resulting database includes every transaction in these model portfolios. These transactions data are the first useful property of the HFD sample.

The HFD does not track all newsletters. With a few exceptions, newsletters are added to the database only on January 1 of each year; these annual additions are based on suggestions from HFD subscribers and staff. When first added to the database, no data from prior years are “back-filled”; the holdings for each portfolio start on the day that HFD begins tracking it. The database is also free of survivor bias: It includes all portfolios that have ever been tracked by HFD, whether or not the portfolios still exist. Back-fill and survivor biases plague many studies of performance evaluation, and their absence here is the second useful property of the HFD sample.

Note that although newsletters must survive long enough to be noticed in order to get into the sample, this is not the same thing as “survivor bias.” Survivor bias refers to a specific statistical problem induced when data are omitted for historical members of the sample who are not alive at the end of the sample period. This is not the case for the HFD database; data exist for all of the historical members of the sample, whether or not they are alive at the end. Rather, the necessity of getting noticed in order to be tracked by HFD results in a nonrandom sample—this by itself would not induce bias unless some element of this nonrandom selection is correlated with subsequent performance and also ignored by the analysis. In this case, there would be a sample-selection bias, but not a survivor bias.² To the extent that

² An example of each kind of bias is the following. Take the universe of all newsletters that ever existed between 1970 and 1990. Now, the sample of such newsletters that survived until 1990 would suffer from survivor bias. Next, consider a random sample of newsletters that existed in 1980. Let the analysis of this sample be based only on their returns *subsequent* to 1980. As long as all of the newsletters initially in this sample (i.e., as of 1980) were included in the analysis, regardless of whether they survived all the way to 1990, then this sample would not have survivor bias but would possibly suffer from sample-selection bias. This latter bias would occur if survivorship until 1980 had predictive power for subsequent returns, and this relationship was otherwise unexplained by the performance-evaluation model. See Carhart (1997b) for a detailed discussion of these issues. Further problems can occur if newsletters come in and out of the sample and returns data exist only since their most recent entry; this “re-emerging bias”—which is not a problem for the HFD sample—is discussed in Goetzmann and Jorion (1996).

such a sample-selection bias does exist in the HFD database, the inclusion of the newsletters that did not get noticed would be likely to make newsletter performance look worse than it does in this paper.

The newsletter returns calculated for this paper are different from returns that would be obtained by actually trading on all newsletter recommendations. First, the analysis ignores all issues of allocation across broad asset classes and the timing of these allocations; only the equity portions of the portfolios are tracked and recommendations of all other assets are excluded. For this reason, the paper is complementary to Graham and Harvey (1996, 1997), who focus on the timing decisions between equities and cash. Second, newsletters' short positions are not followed; all portfolios are assumed to be long 100 percent in equities at all times.³ Third, transactions costs are not considered. For all of these reasons, the returns calculated for this paper do not represent what a real-world investor could expect to achieve by following the advice of the newsletters. The main goal of the paper is to evaluate raw equity performance, and the return calculations are designed for that purpose.

Table I illustrates the growth in HFD's coverage over time. For the entire sample period of July 1980 through December 1996, 153 separate newsletters are covered, with the average newsletter having returns for 81 months and holding 25 stocks in its portfolio. The second column shows the steady increase in HFD's coverage, from 15 newsletters in 1980 to 93 in 1996. The average return for all existing newsletters in any given year is listed in the third column. This average return is calculated in several steps. First, I calculate returns for every "model portfolio" recommended by every newsletter. To combine the returns of different model portfolios within a newsletter, I use an annual rebalancing procedure: For each newsletter, I assume that it invests equally in each of its model portfolios on January 1 of every year; no further rebalancing is done until the following January 1 and each model portfolio is assumed to reinvest all of its returns in its own portfolio. Thus, the weight on each model portfolio will change throughout the year. These steps yield a daily-return series for each newsletter. Finally, to calculate a return on the whole sample of newsletters, I use the same annual rebalancing procedure; each newsletter receives an equal weight on January 1 and reinvests its returns in its own holdings. This yields a return series for the whole sample on every day.⁴ Since the annual return on this series is just the average annual return for all active newsletters, I refer to it throughout the paper as the "average" return series. The fourth column contains the value-weighted-market return (VWM) for that year, where VWM is the total return on the CRSP NYSE/AMEX/Nasdaq index. The percentage of newsletters that beat the VWM for that year is given in the fifth column;

³ Since short positions are excluded, the analysis can only measure the performance of long recommendations; all the results of this paper should be considered with that caveat. But short positions make up less than 8 percent of all positions, so the bulk of the recommendations are still included. Appendix A discusses the reason for this exclusion.

⁴ Further details of these return calculations and their possible biases are discussed in Appendix A.

Table I
Summary Statistics on Newsletters

Summary statistics are given for the *Hulbert Financial Digest* (HFD) sample of 153 newsletters. The average newsletter is in the sample for 81 months and holds 25 stocks. Column 2 shows the total number of newsletters that existed for at least part of the year. Column 3 shows the annual return for an average of all newsletters. This series is rebalanced annually and includes all existing newsletters at all times. Weights are normalized to one for each newsletter on January 1 of each year, with newsletter returns reinvested in their own portfolios. Thus, at all times the weight on each newsletter will be equal to its “year-to-date” return. Column 4 shows the value-weighted-market (VWM) return for the year: the total return on the CRSP index (NYSE/AMEX/Nasdaq). (All returns for 1980 are calculated from July 1, when the HFD data begins, to December 31.) Column 5 shows the fraction of newsletters that beat the VWM return. If a newsletter was only in the market for part of the year, then its return is only compared to the VWM return for that part of the year.

Year	Newsletters Followed	Average Newsletter Return (%)	VWM Return (%)	Percentage of Newsletters Beating VWM Return
1980	15	32.57	22.63	80.0
1981	21	-5.60	-3.95	52.4
1982	24	29.03	20.42	70.8
1983	34	21.40	22.70	47.1
1984	44	-6.24	3.28	22.7
1985	57	32.72	31.47	52.6
1986	59	13.51	15.60	47.5
1987	65	0.28	1.81	35.4
1988	69	16.18	17.61	44.9
1989	69	26.45	28.45	49.3
1990	76	-9.43	-6.02	42.1
1991	79	47.77	33.59	67.1
1992	81	11.32	9.03	54.3
1993	88	15.75	11.48	60.2
1994	88	-3.23	-0.62	39.8
1995	90	36.34	35.73	47.8
1996	93	21.20	21.24	45.1

these results suggest that many newsletters follow a high-beta strategy, since fewer portfolios (usually) beat the market in low VWM years than in high ones.

What kinds of stocks are these newsletters recommending? To answer this question, I look at three different stock characteristics: size (market equity), book-to-market ratio, and momentum (11-month past return, lagged one month). For each characteristic, all stocks with the necessary data are ranked and placed into quintiles. I form quintile breakpoints using only NYSE stocks for size and book-to-market, with these breakpoints used to sort all CRSP-listed stocks into one of the five ordered categories.⁵ The momentum quin-

⁵ Details of these rankings are given in Appendix B. The use of NYSE breakpoints to sort *all* stocks into size and book-to-market quintiles, which results in a very different number of stocks across the size “quintiles,” is to maintain consistency with similar sorts done by other authors. See, for example, Fama and French (1993) and Lakonishok, Shleifer, and Vishny (1994).

Table II
Newsletters Sorted by the Average Size, Book-to-Market,
and Momentum Quintiles of Their Portfolios

Summary of the size, book-to-market ratio, and momentum characteristics of the stocks recommended by the 153 newsletters. Each July 1, all stocks with available data are placed in size and book-to-market quintiles, with quintile breakpoints made using only NYSE stocks. Momentum quintiles are formed monthly using past 11-month returns lagged one month; all stocks in CRSP are used to form momentum breakpoints. (Consult Appendix B for a more complete description of how these rankings are made.) Average quintiles for each characteristic and newsletter are calculated daily by multiplying each stock's portfolio weight (in a newsletter portfolio) by its (numerical) quintile, and summing across all stocks in the portfolio. Column 2 separates newsletters into low (small size, average quintile less than or equal to 2), medium (average quintile between 2 and 4), and high (large size, average quintile greater than 4). Column 3 performs a similar classification for the book-to-market quintile, and column 4 does so for momentum.

	Size	B/M Ratio	Momentum
Low (rank ≤ 2)	23	45	0
Medium ($2 < \text{rank} < 4$)	89	103	115
High (rank ≥ 4)	41	5	38

tiles are formed directly from all CRSP-listed stocks. Then, for each newsletter, I calculate a value-weighted quintile ranking for each characteristic on each day. For example, if on one day a newsletter has half of its weight in stocks that belong to size-quintile 5 and half of its weight in stocks that belong to size-quintile 4, the newsletter would receive a size ranking of 4.5 for that day. The time-series average of these daily rankings is the characteristic ranking for the newsletter. By this method, all newsletters receive a 1 to 5 ranking for each characteristic.

Table II summarizes the results for the 153 newsletters. The second column gives the results for the size rankings. Here, a ranking of 1 would indicate that all of a newsletter's recommendations are for stocks in the smallest quintile (by NYSE breakpoints), and a ranking of 5 would indicate that all the recommendations are for stocks in the largest quintile. The table shows that 23 newsletters have average size rankings less than or equal to 2 ("low"), and 41 newsletters have rankings greater than or equal to 4 ("high"). The remaining 89 newsletters have rankings between 2 and 4 ("medium"). By comparison, the S&P 500 has a size ranking of 4.93, a level exceeded by only two newsletters. Thus, even though newsletters are not focused on the smallest stocks, the average stock they recommend is still considerably smaller than the average stock in the S&P 500.

The third column of Table II shows the results for the book-to-market rankings: 45 newsletters concentrate on low book-to-market "glamour" stocks (average quintile less than or equal to 2), whereas only 5 newsletters concentrate on high book-to-market "value" stocks (average quintile greater than or equal to 4). Finally, the fourth column shows that 38 newsletters have

high momentum rankings, indicating a concentration on the stocks in the highest quintile of past returns, and not a single newsletter focuses on the out-of-favor stocks from the lowest quintile of past returns. The effect of these different strategies on newsletter returns plays an important role in the analysis of the next section.⁶

II. Evaluating Equity Performance: Methodology and Results

In order to properly evaluate stock-selection ability, it is necessary to define a performance-evaluation methodology. Unfortunately, the lack of consensus on the “right” model of expected returns puts performance-evaluation researchers in a quandary, for without a generally accepted model of expected returns it is obviously difficult to define abnormal returns and to quantify the value of investment advice. I deal with this problem by using several different models for expected returns and showing that the evidence on newsletter performance is qualitatively similar in each case.

A. The CAPM

The first model of expected returns is the standard CAPM:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i RMRF_t + \epsilon_{i,t}, \quad (1)$$

where $R_{i,t}$ is the return on newsletter i in month t , $R_{f,t}$ is the risk-free return in month t , and $RMRF_t$ is the month- t value-weighted market return minus the risk-free rate. The estimated intercept, α_i , is the key performance measure from the model. Although the last 20 years have witnessed significant evidence against this unconditional version of the CAPM, it is still used for performance evaluation, both by academics and practitioners.⁷ Thus, it provides a good base case for the analysis.

In the first test, I use the average newsletter return as the dependent variable. This return series is the same one as given in the third column of Table I; all existing newsletters receive equal weight on January 1 and then reinvest their returns in themselves throughout the year. The second column

⁶ It would also be interesting to examine a newsletter’s subscription base as a function of returns and/or strategy, but unfortunately, no reliable data on subscriptions are available.

⁷ The “anomalies” literature begun by Basu (1977) and Banz (1981) has now grown quite large. A recent summary can be found in Campbell, Lo, and MacKinlay (1997), Chap. 5.7. Some examples of the CAPM’s continuing role in performance evaluation can be found in Malkiel (1995), Morningstar (1996), and Shirk, Cuenca, and Carlson (1997). One way to deal with the anomalies is to add more factors or adopt a matching approach based on stock characteristics. These approaches are considered in Sections II.B and II.C, respectively. Another approach, not adopted in this paper, is to do performance evaluation using a conditional version of the CAPM (Ferson and Schadt (1996)).

of Table III gives the results of estimating equation (1) for this average return series. The R^2 of the regression is 0.884. The estimated α is -0.12 (i.e., -12 basis points per month), with a standard error of 0.13, so it is not significantly different from zero. The estimated β is 1.20, reinforcing the evidence of Table I which shows newsletters doing relatively better in bull markets than in bear markets.

I next estimate equation (1) separately for each newsletter, so that each newsletter has an estimate of α_i and β_i . The results are summarized in the second column of Table IV. Under the assumption that the returns from each newsletter are independent (an assumption that will be dropped later), one can perform tests on the percentiles of the α_i distribution.⁸ The idea here is to test whether too many newsletters seem to have extreme performances. For example, even if returns are generated completely by chance, it would not be surprising if one or two newsletters out of 153 have α_i estimates that are significantly higher than zero at the 99-percent level of confidence. But what would be surprising? More formally, under the null hypothesis that the "true" α_i equals 0 for all i , then each newsletter has a left-tail p_i -value corresponding to its α_i estimate. Then, for any value of p^* , the number of newsletters with p_i -values below p^* is distributed binomially (p^* , 153). We can then use the binomial distribution to compute the probability of observing the estimated number of newsletters above or below any p^* . Examples of this calculation are given below. Note the condition that p_i is smaller (greater) than 0.5 implies that α_i is smaller (greater) than 0; a p_i -value of 0.05 corresponds to a "95-percent significant" *negative* α_i , and a p_i -value of 0.95 corresponds to a "95-percent significant" *positive* α_i .

Table IV gives the results for $p^* = 0.001, 0.01, 0.05, 0.5, 0.95, 0.99$, and 0.999. Asterisks next to a table entry denote that there are too many newsletters that meet the relevant condition, where "too many" is defined as a two-tailed (binomial distribution) probability of less than 5 percent (*) or 1 percent (**). For example, under the null hypothesis that α_i equals 0 for all i , we would expect 1 percent of newsletters to fall in the $p_i \leq 0.01$ box (i.e., $0.01 \times 153 = 1.53$ newsletters). The actual sample has five such newsletters. To test whether this is too many, I compute the probability of obtaining five or more successes in 153 trials of a binomial distribution with parameter p^* set equal to 0.01. The binomial probability for this outcome is 0.02. This implies a two-tailed probability of 0.04 and a "rejection" that the number of entries is due to chance. Another rejection occurs for the $p_i \leq 0.05$ box, which has 17 entries. Here we would expect $0.05 \times 153 = 7.515$ newsletters to meet the condition. The probability of at least 17 successes with $p^* = 0.05$ and 153 trials is 0.002. No other box in this column of Table IV allows for rejection. Thus, the only evidence of abnormal performance is that there are too many newsletters with *negative* excess performance.

⁸ A similar procedure is used in Malkiel's (1995) Table III, Ferson and Schadt's (1996) Table III, and Graham and Harvey's (1996) Table II.

Table III
Tests on Average Returns from All Newsletters

Results are given of tests on the average of all newsletter returns. The series is rebalanced annually and includes all existing newsletters at all times. Weights are normalized to one for each newsletter on January 1 of each year, with newsletter returns reinvested in their own portfolios. Thus, at all times the weight on each newsletter will be equal to its “year-to-date” return. Column 2 gives the results for the CAPM (equation (1)); column 3 gives the results for the 4-factor model (equation (2)); column 4 gives the results for the characteristic-matching model (equations (3)–(6)). α is the regression intercept, and the next four rows give coefficients and standard errors (in parentheses) for the independent variables: RMRF, SMB, HML, and PR1. These variables are the returns to zero-investment portfolios designed to capture market, size, book-to-market, and momentum effects, respectively. (Consult Fama and French (1993) and Carhart (1997a) on the construction of these factors.) CS is the Characteristic-Selectivity measure, calculated here as the intercept of a regression of $R_{i,t}$ (the average newsletter return, given by equation (3)) on an intercept term and $R_{b(i),t}$ (the average return to the “matching bins”, given by equation (4)), where the coefficient on the latter is constrained to be one.

	CAPM	4-Factor Model	Characteristic-Matching Model
α	-0.12 (0.13)	0.03 (0.07)	
RMRF	1.20** (0.03)	1.09** (0.02)	
SMB		0.60** (0.03)	
HML		-0.19** (0.03)	
PR1		-0.02 (0.03)	
CS			0.01 (0.06)
Adjusted R^2	0.884	0.970	0.973

** indicates two-tailed significance at the 99-percent level.

Table IV also gives the high and low extreme-values of p_i . Under the null hypothesis, in a sample of N (independent) newsletters, the probability that the highest p_i is at least p is given by $1 - p^N$. Similarly, the probability that the lowest p_i is at most p is given by $1 - (1 - p)^N$. The highest p_i is 0.9881; values at least this high would be expected 84 percent of the time in samples of this size. The lowest p_i is 0.0015; values at least this low would be expected 21 percent of the time. Neither the highest nor the lowest p_i seems unusual. If anything, the results of Table IV suggest that newsletters underperform relative to the CAPM.

The 153 separate CAPM regressions allow only the testing of percentiles one at a time, and the interpretation of these tests rely on the independence of excess returns across newsletters. For a more complete and formal test of newsletter performance under the CAPM, it is necessary to test all percen-

Table IV
Tests on All Newsletters Treated Independently

Results of several tests on the whole sample of individual newsletters are summarized. Column 2 gives the results for the CAPM (equation (1)); column 3 gives the results for the 4-factor model (equation (2)); column 4 gives the results for the characteristic-matching model (equations (3)–(6)). The p -values (p_i) are left-tail probabilities for t -tests on $\alpha_i = 0$ (in columns 2 and 3) and on $CS_i = 0$ (in column 4). The entries in rows 4 to 10 are the number of newsletters with p -values that satisfy the condition given in the first column. For each entry in rows 4 to 10, tail probabilities are computed from the binomial distribution (p , N). Row 3 gives the probability of observing a lowest p -value no higher than p_i (in row 2) in a sample of size N (from row 1). Similarly, row 12 gives the probability of observing a highest p -value of at least p_i (in row 11) in a sample of size N . The independent variables in the 4-factor model (equation (2))—RMRF, SMB, HML, and PR1—are the returns to zero-investment portfolios designed to capture market, size, book-to-market, and momentum effects, respectively. (Consult Fama and French (1993) and Carhart (1997a) on the construction of these factors.) The performance measure in the characteristic-matching model is calculated as the intercept of a regression of $R_{i,t}$ (the return to newsletter i , given by equation (3)) on an intercept term and $R_{b(i),t}$ (the return to the newsletter i 's "matching bins," given by equation (4)), where the coefficient on the latter is constrained to be one.

	CAPM	4-Factor Model	Characteristic-Matching Model
Number of newsletters (N)	153	151	153
Lowest p_i	0.0015	0.0076	0.0003
$1 - (1 - p_i)^N$ (lowest p_i)	0.21	0.69	0.04
$p_i \leq 0.001$	0	0	1
$p_i \leq 0.01$	5*	1	2
$p_i \leq 0.05$	17**	11	9
$p_i \geq 0.50$	66	78	82
$p_i \geq 0.95$	3	15*	11
$p_i \geq 0.99$	0	4	1
$p_i \geq 0.999$	0	0	0
Highest p_i	0.9881	0.9959	0.9920
$1 - p_i^N$ (highest p_i)	0.84	0.46	0.71

* and ** indicate binomial (two-tailed) probabilities of less than 5 percent and 1 percent, respectively.

tiles simultaneously and also to allow for return dependence. In order to calculate a full set of return covariances for such a test, a panel of newsletters with a sufficient number of overlapping months is needed. This requirement forces the exclusion of newsletters that have not survived for a long enough portion of the sample period. This exclusion could induce a positive survivorship bias for the remaining newsletters.⁹ If good stock-

⁹ In this case, the survivorship bias takes the form of a "look-ahead bias" as opposed to a "survivor bias." Look-ahead bias is similar to survivor bias, but is a result of a testing procedure (requiring a minimum return history) rather than a property of the overall sample. Most testing procedures will induce at least some look-ahead bias. See Carhart (1997b) for a detailed discussion of these biases.

Table V
Gibbons, Ross, and Shanken (1989) p -Values

Results are summarized from seemingly unrelated regressions (SUR) on the returns for a subset of newsletters. The table entries are the right-tailed p -values for the Gibbons, Ross, and Shanken (1989) F -test that all the intercept terms are jointly equal to zero. In each test, only the newsletters that have survived for the entire subperiod are included. Column 2 gives the number of newsletters that survived for the corresponding subperiod listed in column 1. Column 3 gives the results for the CAPM (equation (1)); column 4 gives the results for the 4-factor model (equation (2)); column 5 gives the results for the characteristic-matching model (equations (3)–(6)). The independent variables in the 4-factor model (equation (2))—RMRF, SMB, HML, and PR1—are the returns to zero-investment portfolios designed to capture market, size, book-to-market, and momentum effects, respectively. (Consult Fama and French (1993) and Carhart (1997a) on the construction of these factors.) The performance measure in the characteristic-matching model is calculated as the intercept of a regression of $R_{i,t}$ (the return to newsletter i , given by equation (3)) on an intercept term and $R_{b(i),t}$ (the return to newsletter i 's "matching bins," given by equation (4)), where the coefficient on the latter is constrained to be one.

Subperiod	Number of Newsletters	CAPM	4-Factor Model	Characteristic-Matching Model
1989–1996	36	0.92	0.82	0.98
1987–1996	29	0.94	0.51	0.99
1985–1996	25	0.79	0.45	0.85
1983–1996	16	0.34	0.37	0.88
1980–1996	6	0.83	0.66	0.59

selection performance aids survival over the necessary subperiod—a relationship established in Section IV—then the surviving newsletters look better than the whole sample. The implications of this bias are discussed below.

Using time periods of the last 8, 10, 12, and 14 years and the full sample length of 16.5 years, I form subsamples of newsletters that have survived for every month of the relevant period and estimate equation (1) in a seemingly-unrelated-regression (SUR) framework. Gibbons, Ross, and Shanken (1989) (GRS) derive an exact finite-sample F -test for the null hypothesis that the alphas of this estimation are jointly equal to zero. The third column of Table V gives the right-tailed p -values for the GRS test. Other things equal, the more extreme are the α estimations (both positive and negative), the lower will be the p -value for this test. A p -value of 0.05 would imply rejection at the 95 percent level of confidence. In these tests, the lowest p -value is 0.34 for the 14-year subsample (16 newsletters, 168 months), and all of the other subsamples have p -values of at least 0.79. Clearly, in no case would we reject that the CAPM adequately describes these subsamples of newsletter returns. Since this is a joint test of both the model and the null hypothesis on subsamples that may suffer from survivorship bias, the nonrejection is quite striking. It is possible, of course, that unbiased subsamples would result in rejection due to *underperformance*, and that the biased subsamples prevent this rejection. In any case, there is no evidence of superior performance.

B. Carhart's 4-Factor Model

One line of attack on the unconditional CAPM is that it cannot explain differences in returns for portfolios sorted by stock characteristics such as size, measures of "value" such as the price-to-earnings, cash-flow-to-price, or book-to-market ratios, or past returns (momentum).¹⁰ In light of such evidence, researchers have used many different multifactor models in performance-evaluation studies, and no industry standard has yet emerged. Given the evidence presented in Section I indicating the focus of many newsletters on low book-to-market stocks with high past returns, it is important for a study of stock-selection ability to make an attempt to adjust for differences that may result from these strategies. The 4-factor model introduced by Carhart (1997a) attempts to capture the above CAPM anomalies and proves useful in recent studies of mutual funds (Carhart (1997a), Chevalier and Ellison (1999), and DGTW (1997)). The model is estimated by

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,1}RMRF_t + \beta_{i,2}SMB_t + \beta_{i,3}HML_t + \beta_{i,4}PR1_t + \epsilon_{i,t}, \quad (2)$$

where $R_{i,t}$, $R_{f,t}$, and $RMRF_t$ are defined as in equation (1) and SMB_t , HML_t , and $PR1_t$ are the month- t returns to zero-investment factor-mimicking portfolios designed to capture size, book-to-market, and momentum effects, respectively.¹¹ Equation (2) is a straightforward extension of the Fama–French (1993) 3-factor model, a model that cannot adequately explain the returns of portfolios sorted by past return (Fama and French (1996)). Since the 4-factor model lacks a rigorous theoretical basis, one cannot necessarily give a risk interpretation to the factors. Thus, I view this regression as a method of performance *attribution* and interpret the estimated alphas as returns in excess of what could have been achieved by passive zero-cost investments in the "factors." It must be stressed that the debate over the risk interpretation of such multifactor models is ongoing, and the analysis here does not claim to add anything to it.

As with the CAPM, I estimate the 4-factor model for the average newsletter return, for each newsletter independently, and finally in an SUR framework. The third column of Table III shows the results for the average newsletter return. The R^2 of 0.970 indicates that the 4-factor model mimics these returns very closely. Here, newsletters perform slightly better than under the CAPM, with a positive but statistically insignificant point estimate of 0.03 for α . The negative coefficient on HML is consistent with the evidence of Table II and suggests an overall concentration on low book-to-

¹⁰ See Basu (1977) on P/E ratio, Banz (1981) on size, Fama and French (1993) on size and book-to-market, Lakonishok, Shleifer, and Vishny (1994) on several value measures, and Jegadeesh and Titman (1993) on momentum.

¹¹ I am grateful to Mark Carhart for supplying the data on factor returns. For details on how these factors are constructed, see Fama and French (1993) and Carhart (1997a).

market stocks. The positive coefficient on *SMB* may seem puzzling in comparison to the focus on mid-size to large-size stocks that is shown in Table II. However, the coefficient sign should be expected because even the second-largest quintile of NYSE stocks (quintile 4) has a significantly positive loading on *SMB*.¹² The insignificant coefficient on *PR1* is somewhat surprising, considering the number of newsletters found to have portfolios weighted toward stocks with high past returns; the (univariate) momentum focus suggested by Table II does not seem evident in the (multivariate) setting of the 4-factor model. Finally, even in this multifactor setting, the coefficient on *RMRF*, the market “beta,” is still significantly above one.

Next, I estimate equation (2) for each newsletter independently. Out of the full sample of 153 newsletters, 151 have sufficient return histories for this test.¹³ The results are given in the third column of Table IV. Here, newsletter performance looks better than it does under the CAPM. In particular, 15 newsletters have left-tailed p_i -values greater than 0.95. This is significantly more than we would expect to observe by chance; in a binomial test with 151 trials and a parameter of 0.95, the probability of 15 or more “failures” is just over 1 percent. Also, both the $p_i \geq 0.50$ and $p_i \geq 0.99$ boxes have more members than would be expected by chance, although neither number is significant at the 95 percent level. The extreme values, however, are not outside of their normal range; the highest p_i , 0.9959, would be expected to occur approximately 46 percent of the time in a sample of 151 (independent) newsletters. Nevertheless, this is the first suggestion of any superior performance. The problem, of course, is that the returns are not independent.

The SUR estimation attempts to control for the problem of return dependence. As in Section II.A, subsamples are formed by taking all newsletters that have survived for the last 8, 10, 12, 14, and 16.5 years, and then equation (2) is estimated in an SUR framework. The p -values for the GRS F -statistics are given in the fourth column of Table V. Despite the potential survivorship bias in these subsamples, the null hypothesis that the alphas are jointly equal to zero is not rejected for any of the subperiods. The closest the test comes to rejection is for the 1983 to 1996 subperiod, with a right-

¹² See Fama and French’s (1996) Table I. The reason for this somewhat paradoxical result is that the “B” part of *SMB*, which stands for “Small minus Big,” is a value-weighted average dominated by the very largest stocks, but “S” is a value-weighted average actually dominated by midsize NYSE stocks. This occurs because the breakpoints for these categories are formed using only NYSE stocks, with the “S” breakpoint at the median of the NYSE. Thus, midsize NYSE stocks have positive loadings on *SMB* and only the very largest stocks have negative loadings. Since the average quintile ranking of the newsletter sample is only slightly above 3 (midsize), the positive *SMB* loading is not surprising.

¹³ To perform inference on the five coefficients of the model, we need at least six months of data. This data requirement can induce survivorship bias for the remaining newsletters (see Carhart (1997b)). However, because only two newsletters are dropped and neither had extreme performance under the CAPM, this bias is likely to be small or nonexistent.

tailed p -value of 0.37. Thus, once we adjust for return dependencies and test the whole distribution of alphas, there is no significant evidence of superior performance.¹⁴

C. The DGTW Characteristic-Matching Model

In addition to the return anomalies, there are other difficulties with interpreting the alphas from factor-model regressions. As has been pointed out by several authors (Grinblatt and Titman (1995), Ferson and Schadt (1996), DGTW (1997)), estimated alphas and betas are biased when factor loadings are correlated with factor realizations. For example, if a newsletter were to consistently increase its holdings of small stocks (and, thus, its loading on *SMB*) just before periods of high returns for *SMB*, then this positive correlation would serve to bias upward its coefficient on *SMB*, with the effect on its estimated α depending on the sign of the average return to *SMB* over the sample period.

To solve this problem, we could attempt to correctly measure the factor loadings at all times and make adjustments accordingly. If the only available data are portfolio returns, then such a measurement is not possible. With the HFD sample, however, each transaction is known. Therefore, it is possible, in principle, to find “matching” portfolios at all times. In the example described above, such a matching portfolio would increase its holding of small stocks at precisely the same time as the underlying portfolio, and biases due to market timing might be reduced or avoided.¹⁵ A similar approach is often adopted for event studies, with stock-by-stock matching based on characteristics such as size, industry, or the book-to-market ratio.¹⁶ Daniel, Grinblatt, Titman, and Wermers (1997) extend this approach to a performance-evaluation study and use characteristic matching on every stock held by mutual funds in quarterly portfolio “snapshots.” I adopt a slightly modified version of the DGTW characteristic-matching model to obtain the paper’s third performance measure.

To obtain the DGTW measure, I begin by constructing 125 “bins” through a $5 \times 5 \times 5$ sort on size, book-to-market, and momentum quintiles. As in the calculations for Table II in Section I, NYSE breakpoints are used for size and book-to-market, and NYSE/AMEX/Nasdaq breakpoints are used for mo-

¹⁴ An earlier version of this paper (Metrick (1997)) uses the newsletters’ “model portfolios” as the main unit of study, rather than the approach taken by this paper of collapsing all of a newsletter’s portfolios down to a single return. Although performance looks better with portfolios as the unit, many of the best performers are drawn from a high number of nearly identical portfolios of two newsletters. There, SUR tests are not possible because the number of portfolios is too large to estimate return covariances. Since interpretation of those results is problematic, I use newsletter averages as the unit of analysis in this paper.

¹⁵ Another attempt to solve this problem is to adjust for changes in expected factor realizations rather than factor loadings. Ferson and Schadt (1996) take this approach and employ conditional asset-pricing models that use publicly available predictors of factor realizations.

¹⁶ See, for example, Brav and Gompers (1997), Desai and Jain (1995), Ikenberry, Lakonishok, and Vermaelen (1995), and Womack (1996) for recent examples using these characteristics.

mentum, with *all* NYSE/AMEX/Nasdaq stocks placed into quintiles on the basis of these breakpoints. These three characteristics correspond to the *SMB*, *HML*, and *PR1* factors, respectively. Next, I calculate a daily return for each bin; within each bin, stocks are equally weighted on the first day of each month and are assumed to reinvest their returns in their own stock throughout the month. Details on these sorts and return calculations are given in Appendix B. In this model, each stock falls into a bin, which I call its “matching bin.” Abnormal returns are then measured as the difference between a stock’s return and its matching bin’s return.¹⁷

For newsletters, the monthly measure of abnormal returns is calculated as the return on a zero-investment portfolio that is long in the newsletter’s actual portfolio and short in a portfolio constructed using equivalent weights in the matching bins. In effect, I just combine the monthly abnormal returns for each stock in the portfolio. Not all stocks will be included in this calculation; if a stock cannot be matched to a bin, then it is not included in the test. The two main reasons for a failure to match are, first, insufficient past returns for a momentum calculation and, second, the absence of a book-equity observation in COMPUSTAT. Each of these data requirements leads to new issues being deleted from the portfolios. To the extent that new issues underperform similar stocks, the result should be an upward bias in the estimated selectivity performance measure.¹⁸ If such an upward bias exists, it does not seem to have a significant effect on the results. For example, when I repeat the CAPM and 4-factor tests of Table III using returns calculated only from stocks that have bin assignments, the results are almost identical: the CAPM α is -14 basis points (two points lower than the corresponding α of -12 basis points in Table III) and the 4-factor α is five basis points (two points higher than the corresponding α of three basis points in Table III).

The assumption underlying this model is that all stocks in the same bin have exactly the same expected return. If this assumption is satisfied, then the performance measure has zero expected return *at all times*. Thus, newsletters that shift their portfolio composition conditional on expected factor realizations have no bias in their estimated performance measure.

To see how newsletter abnormal returns are calculated, consider the following example. Newsletter XYZ recommends holdings of 40 percent in IBM and 60 percent in Exxon on March 1 and does not recommend any further transactions in March. Suppose that this newsletter earns a return of 3 percent in March. During this time, IBM belongs to bin 122 and Exxon belongs

¹⁷ The choice of these characteristics for bin formation is an attempt to maximize the explanatory power for expected returns while also keeping the number of bins to a manageable number. Nevertheless, it must be acknowledged that the use of these 125 bins is arbitrary and they are used here mainly for consistency with the original model of DGTW (1997).

¹⁸ See Loughran and Ritter (1995) and Brav and Gompers (1997) for evidence on the new-issues bias. On a “bin-adjusted” basis, Brav and Gompers’ work suggests that this bias should not be large. See also Chan, Jegadeesh, and Lakonishok (1995) for a discussion of the bias induced by omitting stocks that do not have data in COMPUSTAT.

to bin 124. Then, form another portfolio on March 1 that holds 40 percent of its holdings in bin 122's stocks (equally weighted) and 60 percent of its holdings in bin 124's stocks (equally weighted). Suppose that this bin portfolio earns a return of 2 percent for March. In this case, the abnormal return for XYZ in March would be $3 - 2 = 1$ percent. If Newsletter XYZ had instead shifted its portfolio during the month, then the bin allocations would have been shifted at the same time.

To represent this formally, some notation is needed:

$d \in t$: the set of all days d in month t ;

$s \in i$: the set of all stocks s held by newsletter i ;

$b(s)$: bin b matched to stock s ;

$R_{s,d}$ = net return on stock s on day d ;

$R_{b(s),d}$ = net return on bin b matched to stock s on day d ;

$W_{s(i),d}$ = weight placed on stock s by newsletter i on day d ;

$R_{i,t}$ = net return to newsletter i in month t ;

$R_{b(i),t}$ = net return to the bins b matched to the stocks in newsletter i in month t .

Then, for each month, the net return on newsletter i is given by

$$R_{i,t} = \prod_{d \in t} \left(1 + \sum_{s \in i} (W_{s(i),d} * R_{s,d}) \right) - 1, \quad (3)$$

where $\sum_{s \in i} (W_{s(i),d} * R_{s,d})$ is the actual net return for newsletter i on day d . The net return to the bins b matched to the stocks in newsletter i in month t is given by

$$R_{b(i),t} = \prod_{d \in t} \left(1 + \sum_{s \in i} (W_{s(i),d} * R_{b(s),d}) \right) - 1, \quad (4)$$

where $\sum_{s \in i} (W_{s(i),d} * R_{b(s),d})$ is the net return that would be achieved on day d if all funds were invested in the bins matched to the recommended stocks of newsletter i . Now, define $CS_{i,t}$, the overall characteristic-selectivity measure for newsletter i in month t , as

$$CS_{i,t} = R_{i,t} - R_{b(i),t}. \quad (5)$$

The analogue to the factor models' α can be found by estimating a regression of $CS_{i,t}$ on an intercept term. The intercept from this regression, analogous to the α_i in the factor models, is called CS_i . There are two other equivalent ways to think about the CS_i calculation, each of which can aid intuition. First, instead of a regression of $CS_{i,t}$ on an intercept term, we can estimate a regression of $R_{i,t}$ on an intercept and $R_{b(i),t}$, where the coefficient on $R_{b(i),t}$ is constrained to be equal to one. The only point of this constrained

regression is that it allows an interpretation of R^2 as the fraction of the return variance explained by the characteristic-matching model. Second, we can interpret $CS_{i,t}$ as the monthly return on a trading strategy that is always long in the underlying newsletter stocks and short in the matching bins. Then, CS_i is just the average monthly return to this trading strategy, and can be calculated as

$$CS_i = \bar{R}_i - \bar{R}_{b(i)}, \quad (6)$$

where $\bar{R}_i$ is the average monthly return to the newsletter and $\bar{R}_{b(i)}$ is the average monthly return to its corresponding bins. All of these methods compute exactly the same estimates for CS_i .

As with the factor models, I use the characteristic-matching model on the average newsletter return, on each newsletter independently, and in an SUR framework using all newsletters that have been in existence for various subperiods. The results for the average return are presented in the fourth column of Table III. For purposes of computing an R^2 for this regression, I use the regression of $R_{i,t}$ on an intercept term and $R_{b(i),t}$, where the coefficient on the latter is constrained to be equal to one. The standard error of the intercept term, CS , is 0.06, which is lower than the equivalent standard errors on the α estimates in the CAPM and 4-factor model (0.13 and 0.07, respectively). This means that one *would* be able to reject the null of no excess performance if CS were found to be greater than 12 basis points in absolute value. The actual estimate of $CS = 0.01$ is not significantly different from zero.

I next analyze the individual CS_i estimates for each newsletter. The results are given in the fifth column of Table IV. The results appear “normal.” Under the assumption that returns are independent, none of the tails of this distribution seem unusually heavy. The worst performer has a t -statistic of -3.50 for 192 months of data. This implies a p_i -value of 0.0003; one would expect to see a p_i at least this low in only about 4 percent of samples of this size. The highest p_i is 0.9920; one would expect to see a p_i at least this high in about 71 percent of samples. None of these results suggest superior performance.

I next drop the assumption of independence for the subsamples of 8, 10, 12, 14, and 16.5 year survivors. The GRS p -values for the SUR estimations on these samples are given in the fourth column of Table V. As with the CAPM and the 4-factor model, I find that the null hypothesis of all alphas being jointly equal to zero cannot be rejected for any of the subperiods, with no p -value less than 0.59. These nonrejections occur despite the potential survivorship bias in the subsamples. Overall, under the characteristic-matching model there is no significant evidence of superior performance.

It could be argued that “real” stock-picking ability is demonstrated not just by earning excess returns relative to bins, but also through successful timing *across* bins. For example, a newsletter may not be a “timer” in the

traditional sense, but it may act on a belief that small stocks are undervalued relative to large stocks. If this belief turns out to be correct, then this skill would likely be missed by the *CS* statistic. DGTW define a measure called “characteristic timing” to capture such effects. I omit the formal definition and analysis of this measure because in unreported tests I find no evidence of any superior performance. If anything, newsletters exhibit below-average characteristic timing.

III. Comparing Performance-Evaluation Models

A. Consistency across Models

The results presented above do not support positive abnormal performance for any of the three models employed. In this subsection, I examine the consistency of the results across models. The Venn diagrams in Figure 1 illustrate the results for the extreme performers in each model. In Panel A, each newsletter with at least one extreme “good” performance is included. A good performance is defined as inclusion in at least one of the $p_i \geq 0.95$ boxes in Table IV: a left-tailed probability of at least 95 percent for a test that α_i (for either the CAPM or 4-factor model) or CS_i (for the characteristic-matching model) is equal to zero.

We can read the entries in Figure 1 to see that only a single newsletter has extreme good performances under all three models, six newsletters have extreme good performances under both the 4-factor model and the characteristic-matching model (but not the CAPM), and two newsletters meet the criterion for both the CAPM and 4-factor model (but not for the characteristic-matching model). By examining the reasons for disparate performance across the models, we can gain insight into the role that model choice plays in performance evaluation. The famous *Value Line Investment Survey* is an interesting case study. For the 168 months that this newsletter is in the sample (1/83 to 12/96), it achieved an annualized return of 21.9 percent.¹⁹ The CAPM α_i is 22 basis points per month, with a standard error for this estimate of 19 basis points. Thus, the t -statistic is not high enough to be significant by the $p_i \geq 0.95$ criterion. Under the 4-factor model, the *Value Line* newsletter portfolio shows a positive loading on *SMB*, a negative loading on *HML*, and a positive loading on *PR1*. As compared to the CAPM, these loadings slightly improve the measured α_i to 32 basis points per month, and the improved fit of the regression (adjusted R^2 of 0.91 versus 0.84 for the CAPM) causes the standard error on this estimate to fall to 15 basis points. The resulting t -statistic for a one-tailed test on $\alpha_i = 0$ is high enough to land *Value Line* in the $p_i \geq 0.95$ box in Table IV. Here, it is both the effect on the point estimate of α_i and its improved precision which cause *Value Line*’s performance to look better under the 4-factor model than under the

¹⁹ *Value Line* is an example of the “no-backfilling” rule of the HFD sample. *Value Line* has been providing recommendations on this portfolio since 1965, but only the years since 1983 (when HFD began contemporaneous coverage) are included in my study.



Panel A. Extreme "Good" Performers: $p \geq 0.95$ **Panel B. Extreme "Bad" Performers: $p \leq 0.05$** 

Figure 1. Model comparison for extreme performers. The number of newsletters with extreme performance across the three performance-evaluation models is illustrated. The Venn diagram in Panel A shows the extreme "good" performers: those newsletters with a left-tailed probability greater than or equal to 95 percent for t -tests on $\alpha_i = 0$ (in the CAPM and 4-factor model) or $CS_i = 0$ (in the characteristic-matching model). This cutoff corresponds to the eighth row of Table IV ($p_i \geq 0.95$), and the total in each circle in Panel A equals the entry for the corresponding column in the eighth row of Table IV. Panel B analogously shows the extreme "bad" performers: ($p_i \leq 0.05$ on any of the tests.) This cutoff corresponds to the sixth row in Table IV. The independent variables in the 4-factor model (equation (2))—RMRF, SMB, HML, and PR1—are the returns to zero-investment portfolios designed to capture market, size, book-to-market, and momentum effects, respectively. (Consult Fama and French (1993) and Carhart (1997a) on the construction of these factors.) The performance measure in the characteristic-matching model is calculated as the intercept of a regression of $R_{i,t}$ (the return to newsletter i , given by equation (3)) on an intercept term and $R_{b(i),t}$ (the return to newsletter i 's "matching bins," given by equation (4)), where the coefficient on the latter is constrained to be one.

CAPM. For the characteristic-matching model, the estimated CS_i is 21 basis points per month, but the standard error of this estimate is 11 basis points per month. This added precision allows *Value Line* to demonstrate significantly positive performance with its CS_i measure.

Panel B of Figure 1 shows the analogous categories of extreme “bad” performers under the three models. A bad performance is defined as inclusion in one of the $p_i \leq 0.05$ boxes in Table IV. Again, it is useful to look at a case where the measures might lead to different conclusions. The *RHM Survey of Warrants* has significantly negative abnormal performance under the CAPM, with an estimated α_i of -99 basis points per month (t -statistic of -2.26) over the 156 months that it is in the sample. As it turns out, the data suggest that this newsletter focused on relatively low-momentum small growth stocks—a terrible strategy over its lifespan. When its return history is analyzed using the 4-factor model, it is found to load positively on *SMB*, negatively on *HML*, and negatively on *PR1*—all three of these loadings were bad for realized returns over the relevant period. Once adjusted for these loadings, the newsletter has a 4-factor α_i of -60 basis points per month (t -statistic of -1.37); this is still low, but not low enough to land in the $p_i \leq 0.05$ category. The results for the CS_i measure are a bit better, with an estimate of -28 basis points per month (t -statistic of -0.75). Of course, *RHM* is still responsible for its poor (or unlucky) choice of strategy, but the analysis can remain agnostic as to whether this strategy earned lower returns for risk or nonrisk reasons. Nevertheless, the use of several different models does allow one to separate out the “equity style” from the “stock-selection” components of *RHM*’s performance.

B. Precision across Models

A model comparison on a case-by-case basis is helpful for understanding individual newsletter performance, but it cannot say anything general about the relative precision across models. There are certainly specific cases where precision seems to increase as one moves from the CAPM to the 4-factor model to the characteristic-matching model, but is this true in general? Should a researcher who has transactions data bother with a characteristic-matching model, or would a factor model be just as powerful? With the advent of transactions-based databases in performance evaluation, these questions have important practical implications.²⁰ It is possible to answer these questions using simulated portfolios, but this approach may miss important elements of actual managed portfolios. For example, managed portfolios often follow focused size, value, momentum, or industry-based strategies that may evolve over time. If simulated portfolios do not reflect these strategies, then the results would not accurately capture the relationship be-

²⁰ For recent examples of transactions databases in performance evaluation analyses, see Eckbo and Smith (1998) and Jeng (1998) for insider trading, Barber and Odean (1999) and Grinblatt and Keloharju (1998) for individual investors, and DGTW (1997) for mutual funds.

tween models; recall that the characteristic-matching model is designed specifically to capture such effects. It seems useful to analyze the precision of these models “in the field,” and that is the purpose of this section.

The goal of this analysis is to construct measures of precision that can be compared across models. To derive these measures, consider the test statistics for any given newsletter under each model. For the CAPM,

$$t_{CAPM_i} = \frac{\alpha_{CAPM_i}}{\sigma_{CAPM_i}}, \quad (7)$$

where σ_{CAPM_i} is the standard error of the α_{CAPM_i} estimate. For the 4-factor model,

$$t_{4F_i} = \frac{\alpha_{4F_i}}{\sigma_{4F_i}}, \quad (8)$$

where σ_{4F_i} is the standard error of the α_{4F_i} estimate. For the characteristic-matching model,

$$t_{CS_i} = \frac{CS_i}{\sigma_{CS_i}}, \quad (9)$$

where σ_{CS_i} is the standard error for the characteristic-selectivity measure, CS_i . To motivate the measure of precision, consider an experiment in which each stock for a given newsletter has a small and equal increase in its return for every month it is held by the newsletter. This induces an equivalent increase in the monthly newsletter return, $\bar{R}_i$. What would be the effect of this experiment on the test statistic for each model? In other words, what would be the derivative of the test statistics with respect to $\bar{R}_i$?

Since the factor returns are calculated from a large number of stocks, one can assume that the increase in the returns of a single newsletter's stocks would have only a negligible effect on the factor returns. Therefore, for both factor models, an increase in an average newsletter return, $\bar{R}_i$, induces a one-for-one increase in α_i , so the effects on the test statistics are

$$\frac{\partial t_{CAPM_i}}{\partial \bar{R}_i} = \frac{1}{\sigma_{CAPM_i}} \quad (10)$$

for the CAPM, and

$$\frac{\partial t_{4F_i}}{\partial \bar{R}_i} = \frac{1}{\sigma_{4F_i}} \quad (11)$$

for the 4-factor model. These measures of precision are simply the inverses of the standard errors for the respective α_i estimates. For the characteristic-matching model, the procedure is slightly more complicated. First, equation (6) can be substituted into equation (9) to yield

$$t_{CS_i} = \frac{\bar{R}_i - \bar{R}_{b(i)}}{\sigma_{CS_i}}, \quad (12)$$

where $\bar{R}_{b(i)}$ is the mean monthly return of the bins matched to the stocks in newsletter i . To compute the derivative of this t -statistic with respect to $\bar{R}_i$, we need to consider separately each term in the numerator. The derivative of $\bar{R}_i$ with respect to itself is obviously one, but what about the second term? A logical criticism of DGTW's characteristic-matching approach is that the use of "too many" bins effectively renders CS_i to be meaningless. In the extreme, of course, this criticism must be true. If we were to choose as many bins as there are stocks, then by definition all $\bar{R}_i$ would be equal to $\bar{R}_{b(i)}$ and all CS_i would be zero. But is 125 bins "too many"? If the bins do not contain many stocks, or if a newsletter owns most of the stocks in any one bin, then the effect of our experiment on the bin returns could be substantial. Consider the effect of the experiment for a "representative" first day of a month—thereby eliminating the need for time subscripts and simplifying the weighting of stocks within bins.²¹ Let $n_{s,b}$ be the number of stocks in the bin b that contains stock s . If stock s has an increase in its return, then its effect on the corresponding bin's return, $\bar{R}_{b(i)}$, is proportional to $1/n_{s,b}$. Therefore, I write λ_i , the derivative of $\bar{R}_{b(i)}$ with respect to $\bar{R}_i$, as

$$\lambda_i = \frac{\partial \bar{R}_{b(i)}}{\partial \bar{R}_i} = \sum_{b=1}^{125} \sum_{s \in i} \left(W_{s(i)} * \sum_{s \in i} \frac{1}{n_{s,b}} \right), \quad (13)$$

where $W_{s(i)}$ is the weight placed on stock s by newsletter i . In an extreme case where there is only one stock in every bin (i.e., all $n_{s,b} = 1$ and $\bar{R}_i = \bar{R}_{b(i)}$), then λ_i is equal to one; if there are very few bins (i.e., large $n_{s,b}$), then λ_i is close to zero. With equation (13) in hand, we can differentiate equation (12) with respect to $\bar{R}_i$; this yields the total effect of an increase in the return for every stock in newsletter i on the t -statistic for CS_i :

$$\frac{\partial t_{CS_i}}{\partial \bar{R}_i} = \frac{1 - \lambda_i}{\sigma_{CS_i}}. \quad (14)$$

²¹ After the first of the month, weighting of stocks within bins will not be equal. Thus, the analysis below is only an approximation on those days. See Appendix B for a discussion of bin construction.

Equations (10), (11), and (14) are measures of precision for each test. I next compare these measures. The ratio of equations (11) to (10) leads to a natural measure of relative precision between the 4-factor model and the CAPM:

$$\text{relative precision (4-factor/CAPM)} = \frac{\frac{\partial t_{4F_i}}{\partial \bar{R}_i}}{\frac{\partial t_{CAPM_i}}{\partial \bar{R}_i}} = \frac{\sigma_{CAPM_i}}{\sigma_{4F_i}}. \quad (15)$$

For the 151 newsletters for which both models can be estimated, 81 have a relative precision greater than one; in these cases, the 4-factor model provides a more precise estimate than the CAPM. The median measure in this sample is just below 1.01. This relative precision is higher for newsletters with long return histories or with many stocks in their portfolios. In a regression of relative precision (or its logarithm) on length of return history and “average number of stocks held,” the coefficients on both independent variables are positive and significant at the 95 percent level. As an illustration of this effect, if we consider only those newsletters with at least 10 years of data, then the median measure of relative precision is 1.03, and 27 out of 40 newsletters have a measure greater than one. Similarly, among newsletters that hold an average of 20 or more stocks in their portfolios, the median measure is 1.11, with 43 out of 61 measures greater than one. Clearly, the benefits of the 4-factor model over the CAPM do not become apparent unless newsletters have long histories and/or hold many stocks.²²

I next turn to a comparison of the 4-factor and characteristic-matching models. First, it is necessary to redefine the set of stocks so that they are the same for both models. In Section II.B, the 4-factor regressions are estimated on newsletter returns using all stocks, but the estimates of CS_i in Section II.C are (necessarily) done only on the subset of stocks that have bin assignments. To make an apples-to-apples comparison, we must reestimate the 4-factor model on the same returns as used for the characteristic-matching model. The results are qualitatively similar to those presented in Section II.B. Dividing equation (14) by equation (11) yields a relative precision for the two models of

$$\text{relative precision (characteristic-matching/4-factor)} = \frac{\frac{\partial t_{CS_i}}{\partial \bar{R}_i}}{\frac{\partial t_{4F_i}}{\partial \bar{R}_i}} = \frac{(1 - \lambda_i) \sigma_{4F_i}}{\sigma_{CS_i}}. \quad (16)$$

²² This analysis ignores differences in the degrees of freedom of each test statistic. Since the 4-factor model estimates three more parameters than the CAPM, this makes the 4-factor model relatively less precise than the measure in equation (15) suggests. For newsletters with more than 30 months of data, these effects would usually be negligible.

An estimate of relative precision can be computed for 150 newsletters.²³ Estimates of λ_i are computed on July 1 of every year, and then averaged across all years to form a single estimate for each newsletter. For the whole sample, the median λ_i is 0.069. The median measure of relative precision for the 150 newsletters is 1.10, with 126 of the newsletters having ratios greater than one.²⁴ Overall, relative precision in equation (16) is not related to the length of the return history nor to the average number of stocks held. For example, among newsletters with at least 10 years of history, the median measure is 1.09 and 34 out of 40 newsletters have measures greater than one. Furthermore, there are no significant coefficients in a regression of relative precision (or its logarithm) on length of return history and average number of stocks held.

It is important to note that this analysis of precision does not imply anything about which model is “true.” For example, it is possible for the CAPM to be the true model of expected returns but for the 4-factor model to nevertheless provide more precise estimates. This can occur because the additional factors proxy for sample-period differences in the returns to various “styles.” Such styles may not actually affect *expected* returns, but are useful for performance attribution of *realized* returns over any given time period. For this reason, the researcher may care about precision while remaining agnostic about the true underlying model.

There are two main conclusions from this section. First, the gain in precision for the 4-factor model over the CAPM is small or nonexistent for newsletters with few stocks and short return histories, but grows with history length and number of stocks. Second, a characteristic-matching approach offers the potential for significant gains in precision over the 4-factor model. These gains occur across the board; they are not a function of return history or average number of stocks held by newsletters. This finding is worthy of further study on other real and simulated data, with extensions to a wider class of factor and characteristic-based methods.

IV. Do Newsletters Get “Hot Hands”?

The tests of Section II are designed to assess the stock selection of newsletters over their entire lifetimes. The results show no significant evidence of superior performance. However, what if some newsletters can get “hot hands” and make successful stock selections for short periods? Suppose that instead of picking a newsletter and sticking with it for its entire existence,

²³ When all stocks are used to compute returns 151 newsletters have adequate histories (Section II.B), but one newsletter drops to only four months of returns when the bin assignments are required.

²⁴ As in equation (15), the measure in equation (16) does not make degrees-of-freedom corrections. Such corrections would tend to make the comparison more favorable to the characteristic-matching model, which estimates only one parameter versus five in the 4-factor model. These effects would usually be negligible for newsletters with more than 30 months of history.

we were to choose only the best performers over the previous year. Could such a strategy earn excess returns for stock selection? Specifically, consider a trading strategy that divides the entire newsletter universe into deciles based on their raw returns from the previous year. Then, each January, reset the portfolio to include only the newsletters from the highest decile. The annualized return to this strategy would have been 15.7 percent for the 1981–1996 period, which is slightly higher than the 14.4 percent earned by the CRSP value-weighted index. Conversely, the annualized return to the strategy that invested in the lowest decile would have been only 8.9 percent over the same time period. The success of this simple short-term persistence strategy would appear to lend support to the existence of hot hands. But is this a real persistence of superior performance, or can it be explained by factor loadings?

A similar hot-hands effect is well documented in other contexts. Graham and Harvey (1996) find that the market-timing recommendations of newsletters whose previous prediction was “correct” are significantly better than those of newsletters whose previous prediction was “incorrect,” but they also find that this persistence is not strong enough to identify any long-term overperformance. Brown and Goetzmann (1995) demonstrate a pattern of short-run persistence in mutual funds, concentrated among poor performers, and provide evidence that this persistence is due to common strategies that are not captured by standard models. Hendricks, Patel, and Zeckhauser (1993) show that a strategy of buying the best past performers and selling the worst performers earned significant positive returns from 1974 to 1988 within a sample of no-load mutual funds. Carhart (1997a) confirms this result but shows that almost all of the difference can be explained by the 4-factor model (equation (2)), with the only remaining persistence occurring for the worst-performing funds. In this section, I adopt Carhart’s methodology and analyze the returns of decile-sorted past performers.

Before turning to the factor models, it is helpful to review newsletter survivorship and the persistence of the decile rankings themselves. Figure 2 shows a histogram of newsletter “death” frequencies in year t as a function of the year $t - 1$ raw-return decile. The overall trend suggests that good performance helps survival. The top half (deciles 1 through 5) have an average death rate of 5 percent, and the bottom half (deciles 6 through 10) have an average death rate of 10 percent. A logit estimation of survival (1 if yes, 0 if no) on the previous year’s decile yields a point estimate of 0.15 with a standard error of 0.05. This estimate can be translated into probability terms, where it implies a 1 percent survival advantage in year t for each higher decile in year $t - 1$.²⁵ These results are consistent with Graham and Harvey’s (1996) finding that good market-timing performance by newsletters increases their survival probabilities. Also, in samples of mutual funds,

²⁵ The point estimate from the logit regression can be (approximately) converted to probability terms by multiplying it by $p * (1 - p)$, where p is set to 0.93, the sample mean survival rate (see Maddala (1983)). This yields an answer of $0.15 * 0.93 * 0.07 = 0.0098 \approx 1.0$ percent.




Figure 2. Newsletter death frequencies. Newsletter death frequencies are plotted by year t as a function of the year $t - 1$ decile. At the beginning of each year t , all newsletters are placed into decile rankings by their raw returns in year $t - 1$. If a newsletter ceases publishing at any time before the end of year t , then this is counted as a “death.” For the entire sample, the death frequency for year t for each year $t - 1$ decile is computed and plotted in the figure.

the impact of performance on survivorship is large and significant.²⁶ Despite the fact that my performance measure is focused purely on stock selection, ignores transactions costs, and is not itself a published measure available to newsletter or HFD subscribers, there is still a correlation with survival.

Similarly, the decile rankings also show some persistence. The correlation of year $t - 1$ deciles with year t deciles is 0.09, a positive relationship that is significant at the 99 percent level. This result is consistent with the relative success of the trading strategy of buying the top-performing decile in each year. However, we still do not know if the persistence is explained by “betas” (factor loadings) or “alphas” (stock-picking skill).

Table VI summarizes the results of 4-factor regressions (equation (2)) on each of the hot-hands deciles, and also on a zero-investment portfolio long in the highest-performing decile and short in the lowest-performing decile.²⁷ The $PR1$ coefficients match well with intuition, with significant positive coefficients for deciles 1 and 3, and significant negative coefficients for deciles 9 and 10. The last row of Table VI shows the regression estimates when decile 1 minus decile 10 returns are used as the dependent variable. The

²⁶ For evidence on this point, see Brown and Goetzmann (1995), Malkiel (1995), and Carhart (1997b).

²⁷ The tests in this section are all carried out using rankings on past raw returns rather than rankings based on past performance measures. This is done in part to stay consistent with the literature, but also because past performance measures, if biased, would then bias the persistence tests as well. See Carhart (1997a) for a discussion of this point.

Table VI
4-Factor Regressions on “Hot Hands” Deciles

This table summarizes the results of 4-factor regressions (equation (2)) on decile portfolios of newsletters sorted by past return. Each January 1, all newsletters are ranked and placed into deciles based on their previous year's raw return. Decile 1 has the highest past performance. Decile returns are calculated as the average of their constituent newsletters' returns, with weights on each day equal to a newsletter's “year-to-date” returns. The resulting 10 return series, plus an additional series consisting of decile 1 minus decile 10, are then regressed on the four factors of equation (2). Column 2 gives the adjusted R^2 for each regression. In column 3, α is the regression intercept, and the next four columns give coefficients on the factors: RMRF, SMB, HML, and PR1. These variables are the returns to zero-investment portfolios designed to capture market, size, book-to-market, and momentum effects, respectively. (Consult Fama and French (1993) and Carhart (1997a) on the construction of these factors.)

Decile	R^2	α	RMRF	SMB	HML	PR1
1 (best)	0.81	0.09	1.14**	0.97**	-0.40**	0.26**
2	0.71	0.32	1.04**	0.68**	-0.26*	0.08
3	0.86	0.04	0.99**	0.63**	-0.22**	0.13*
4	0.89	-0.01	1.15**	0.46**	-0.19**	-0.00
5	0.83	-0.04	1.02**	0.42**	-0.05	-0.01
6	0.88	-0.19	1.09**	0.88**	0.08	0.04
7	0.83	-0.42*	1.11**	0.59**	-0.12	0.02
8	0.84	0.15	1.13**	0.68**	-0.28**	-0.11
9	0.81	0.15	1.15**	0.67**	-0.19*	-0.15*
10 (worst)	0.60	0.25	0.96**	0.45**	-0.61**	-0.29*
1 - 10	0.08	-0.16	0.18	0.52**	0.22	0.55**

* and ** indicate two-tailed significance at the 95- and 99-percent levels, respectively.

estimated α is -16 basis points per month. Under the 4-factor model, decile 10 appears to outperform decile 1. What happened to the excess performance reported at the top of this section? A large part of it is attributable to the momentum factor. The average return to PR1 over this sample period is 85 basis points *per month*. A loading of 0.55 on PR1 attributes 47 basis points per month by itself. Moreover, decile 1 has a higher loading on RMRF and HML than does decile 10; both these factors have positive returns over the sample period. Of all the α estimates, only the α for decile 7 is significant at the 95 percent level. In an SUR framework, the p -value for the GRS F -statistic for the 10 deciles is 0.51—well above any critical level. Although there is mild short-term persistence in raw performance, there is no evidence that newsletters exhibit *abnormal* short-term persistence relative to the 4-factor model.

V. Conclusion

As greater numbers of individuals participate in equity markets, it becomes increasingly important for financial economists to form scientific opinions about the behavior and performance of retail investors and the advisors they follow. In such instances, our theories and intuition are often strong,

but the data availability is weak and empirical results are rare. This paper attempts to partially fill this gap by studying the rich and carefully constructed HFD database of 153 investment newsletters. The unique properties of the database allow for insights into the methodology of performance evaluation; these insights suggest significant benefits for using transactions data as a supplement to returns data.

Several different methodologies are employed to evaluate the stock-picking skills of newsletters: a pair of returns-based factor models, the CAPM and the 4-factor model of Carhart (1997a), and a transactions-based approach, the characteristic-matching model of Daniel, Grinblatt, Titman, and Wermers (1997). Each model is estimated on an average newsletter return, on each newsletter independently, and on dependency-adjusted subsamples of long-surviving newsletters. Overall, newsletters do not demonstrate significant abnormal performance: Average abnormal returns are close to zero, the best-performing newsletter under each model does not seem unusual given the sample size, and, for the CAPM and characteristic-matching model, there are not “too many” good-performing newsletters. Only under the 4-factor model do there appear to be too many good performers, and even this weak result disappears once return dependencies across newsletters are explicitly considered.

In addition to these tests on the whole history of newsletter returns, I also test for short-term persistence of returns: Do newsletters get “hot hands”? To test for persistence, all newsletters are placed into decile portfolios based on their previous year’s return. When the 4-factor model is estimated on 16 years of returns to these decile-sorted portfolios, the estimated set of alphas is insignificantly different from zero. Although a strategy of buying the best past performers and selling the worst ones would have earned positive raw returns over the sample period, the abnormal returns would have been negative and insignificantly different from zero. Overall, there is no evidence that newsletters have superior stock-selection skill, either over short or long horizons.

The main methodological contribution of the paper takes the form of a “field-study”; in this sample of newsletters, I measure the relative precision of the three models: CAPM, 4-factor, and characteristic-matching. I find that the latter approach, which requires transactions data, shows a median improvement in precision of 10 percent over the 4-factor model. This compares favorably with a median improvement of 1 percent for the 4-factor model over the CAPM. The 4-factor model shows its greatest improvement over the CAPM for newsletters with long histories and many stocks in their portfolios; even for this group the characteristic-matching model is more precise. These results suggest that researchers with transactions data should attempt to exploit this property when they do performance evaluation.

Appendix A. The Calculation of Newsletter Returns

Newsletters make recommendations in many different forms, and the HFD has devised a methodology to convert advice into “model portfolios.” Depending on the year, between one-half and two-thirds of the newsletters give

explicit advice about percentages or shares to be held of each security. The remaining newsletters do not give explicit advice, but use some system of ranking stocks by desirability, such as a numerical ranking or a “buy,” “hold,” “sell” classification. HFD takes such rankings and forms equally weighted portfolios from the stocks in the most highly recommended category, treating removals from this category as sells. The resulting model portfolios, some exact and some constructed, are the raw material for the paper. At this point, the HFD database leaves off and my analysis begins.

The first thing done with these model portfolios is to purge them of all nonequities and short-sales. Model portfolios are rebalanced on “transaction days” so that they include only the long equity positions. The idea here is to focus on pure equity positions and be able to abstract from market timing (equities versus other assets). Short positions are ignored because there is no obvious way to include them without allowing implicit market-timing bets. For example, consider a newsletter that recommends a portfolio that is 100 shares (long) in IBM and 100 shares (short) in Microsoft. If Microsoft is more expensive per share than IBM, then this results in a net short position in equities. If the Microsoft position is taken off, then the portfolio becomes net long in equities. The returns to the newsletter then reflect this market-timing strategy. Although the CS measure (Section II.C) would not be affected by this, the factor models (Sections II.A and II.B) would be affected, thereby complicating the comparison across models. The eliminated short recommendations comprise approximately 7.6 percent of all newsletter positions, and more than half of the newsletters never make a short recommendation. Therefore, the remaining analysis still includes the vast majority of equity recommendations.

To calculate returns to these portfolios, I assume that all trades take place at the first daily closing price after the recommendation was received. This methodology implies that a newsletter with a large market impact would have its returns underestimated. Unfortunately, it is not possible to do a direct test for this bias, because some stocks are recommended through telephone hotlines after major moves have already occurred on that day; any inclusion of “day zero” returns would be confounded by this effect. Since the methodology employed in this paper uses closing prices for the day that recommendations were received in the mail (or by phone before the market close), it is at least in principle an “implementable” strategy, albeit one without transaction costs. Only those stocks contained in the CRSP daily files are included in the return calculation. Because most stocks do not have delisting returns in CRSP, I substitute the estimated delisting returns suggested by Shumway (1997) and Shumway and Warther (1997): –30 percent for performance-related delists on the NYSE/AMEX and –55 percent on the Nasdaq. At the end of this procedure, each model portfolio has a daily return series for its whole lifetime.

One possible source of bias in portfolio returns is the differences in the ways stocks are recommended. Newsletters that recommend exact share amounts for every stock should have well-constructed returns. However, newsletters that recommend only portfolio weights present some problems. For

example, consider a newsletter that recommends a holding of 50 percent in IBM and changes the other 50 percent of the portfolio every day. Such a portfolio effectively gains the benefits of bid-ask bounce in IBM, with a daily rebalancing of its portfolio weight back to 50 percent. In the extreme, daily rebalancing can cause large differences in measured returns (see Canina et al. (1998)). To see if this bias is evident in the sample, I compare the *t*-statistics in each model for “percentage-type” versus “share-type” newsletters. To keep *t*-statistics comparable, I restrict this comparison to newsletters that have at least two years of data. The results are that, for each model, the share-type newsletters have *higher* mean *t*-statistics than the percentage-type newsletters, although these differences are never statistically significant. Thus, it seems unlikely that rebalancing bias is playing a large role. This is probably because very few percentage-type newsletters transact often enough to induce a measurable effect.

Since many newsletters have multiple portfolios, the next step is to create “newsletter returns” from these portfolio returns. This step is necessary because several newsletters have multiple portfolios with nearly identical holdings, and this duplication leads to difficulties in the interpretation of results on the full sample of portfolios. The results of this paper are quantitatively similar and qualitatively identical for several different methods of combining portfolio returns into newsletter returns. In the method used in the text, newsletter returns are calculated as an annually rebalanced average of their portfolio returns. That is, each newsletter is assumed to invest equally in all of its portfolios on January 1 of each year. All portfolios reinvest their own proceeds, so that weight on each portfolio shifts throughout the year. This yields a daily return series for each newsletter. The resulting annual newsletter return is the same as the average annual return across all of its portfolios.

Appendix B. Formation of the Quintiles and the 125 Bin Returns

Each July 1, all stocks are sorted into size and book-to-market quintiles using NYSE breakpoints. Size is computed at June month-end, and the book-to-market ratio is calculated as a stock’s most recent fiscal year’s book value of equity divided by the previous end-of-December market value of equity. Stocks that lack any of the necessary measures in CRSP (for market equity) or COMPUSTAT (for book equity) are excluded. REITs, ADRs, units of beneficial trust, and negative book-equity stocks are also excluded. The use of NYSE breakpoints causes the number of stocks in the smallest size quintile, and to some extent the second smallest quintile, to be greater than those in the other quintiles. The momentum quintiles are recalculated monthly based on prior 11-month returns lagged one month. This is the same definition of momentum as in Carhart (1997a) and DGTW (1997); as pointed out in the latter paper, the omission of the most recent month is done to avoid problems of bid-ask bounce and monthly return reversals. Momentum quintile-breakpoints are formed using *all* stocks. All stocks tied on a quintile breakpoint

are assigned to the lower of the two quintiles. A “bin” is the intersection of the three independent quintile sorts. Size and book-to-market quintiles do not change from July to June, but stocks can still change bins once a month when their momentum quintiles change. The returns to each bin are then computed from a monthly (*not* daily) equal-weighting procedure. Specifically, the weights for each stock are set to be equal on the first day of the month, and then a stock’s total return is assumed to be reinvested in its own shares through the month. Therefore, each stock’s weight is always proportional to its month-to-date return. This procedure is motivated by a suggestion in Canina et al. (1998); the goal is to minimize “within-bin” size effects without incurring the compounding bias of daily equally weighted returns. Where appropriate, I use the same estimated delisting returns as described in Appendix A. The bin returns are used as the $R_{b(s),d}$ in equation (3) and Section II.C. The resulting calculation of CS_i includes only stocks that have the data necessary to place them into bins. The omission of stocks that lack the necessary data should not be a source of bias: See Chan et al. (1995), Brav and Gompers (1997), and the CAPM and 4-factor regression results discussed in Section II.C.

REFERENCES

- Banz, Rolf, 1981, The relation between return and market value of stocks, *Journal of Financial Economics*, 38, 269–296.
- Barber, Brad M., and Douglas Loeffler, 1993, The “dartboard” column: Second-hand information and price pressure, *Journal of Financial and Quantitative Analysis* 28, 273–284.
- Barber, Brad M., and Terrance Odean, 1999, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance*, forthcoming.
- Basu, Sanjoy, 1977, The investment performance of common stocks in relation to their price-to-earnings: A test of the efficient markets hypothesis, *Journal of Finance* 32, 663–682.
- Black, Fischer, 1973, Yes Virginia, there is hope: Tests of the *Value Line* ranking system, *Financial Analysts Journal* 29, 10–14.
- Brav, Alon, and Paul Gompers, 1997, Myth or reality: The long-run underperformance of initial public offerings: Evidence from venture and nonventure capital-backed companies, *Journal of Finance* 52, 1791–1822.
- Brimelow, Peter, 1986. *The Wall Street Gurus* (Random House, New York).
- Brown, Keith C., W. V. Harlow, and Laura T. Starks, 1996, Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry, *Journal of Finance* 51, 85–109.
- Brown, Stephen, and William N. Goetzmann, 1995, Performance persistence, *Journal of Finance* 50, 679–698.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997. *The Econometrics of Financial Markets* (Princeton University Press, Princeton, N.J.).
- Canina, Linda, Roni Michaely, Richard Thaler, and Kent Womack, 1998, Caveat compounder: A warning about using the daily CRSP equal-weighted index to compute long-run excess returns, *Journal of Finance* 53, 403–416.
- Carhart, Mark, 1997a, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Carhart, Mark, 1997b, Mutual fund survivorship, Working paper, Marshall School of Business, University of Southern California.
- Chan, Louis K.C., Narasimhan Jegadeesh, and Josef Lakonishok, 1995, Evaluating the performance of value versus glamour stocks: The impact of selection bias, *Journal of Financial Economics* 38, 269–296.

- Chevalier, Judith, and Glenn Ellison, 1999, Are some mutual fund managers better than others? Cross-sectional patterns in behavior and performance, *Journal of Finance* 54, 875–899.
- Copeland, Thomas E., and David Mayers, 1982, The *Value Line* enigma (1965–1978): A case study of performance evaluation issues, *Journal of Financial Economics* 10, 289–321.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic based benchmarks, *Journal of Finance* 52, 1035–1058.
- Daniel, Kent, and Sheridan Titman, 1997, Evidence on the characteristics of cross sectional variation in stock returns, *Journal of Finance* 52, 1–34.
- Desai, Hemang, and Prem C. Jain, 1995, An analysis of the recommendations of the ‘superstar’ money managers at Barron’s annual roundtable, *Journal of Finance* 50, 1257–1274.
- Eckbo, B. Espen, and David C. Smith, 1998, The conditional performance of insider trades, *Journal of Finance* 53, 467–498.
- Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene, and Kenneth French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.
- Ferson, Wayne, and Rudi Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425–461.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.
- Goetzmann, William N., and Phillippe Jorion, 1996, Re-emerging markets, NBER Working paper #5906.
- Graham, John R., 1999, Herding among investment newsletters: Theory and evidence, *Journal of Finance* 54, 237–268.
- Graham, John R., and Campbell R. Harvey, 1996, Market timing ability and volatility implied in investment newsletters’ asset allocation recommendations, *Journal of Financial Economics* 42, 397–421.
- Graham, John R., and Campbell R. Harvey, 1997, Grading the performance of market timing newsletters, *Financial Analysts Journal* 53, 54–66.
- Grinblatt, Mark, and Matti Keloharju, 1998, Momentum investing and performance using Finland’s unique data set, Working paper, UCLA.
- Grinblatt, Mark, and Sheridan Titman, 1995, Performance evaluation; in R. A. Jarrow, V. Maksimovic, and W. T. Ziemba, eds.: *Handbook in Operations Research and Management Science: Vol. 9, Finance* (Elsevier, Amsterdam).
- Hendricks, Darryl, Jayendu Patel, and Richard Zeckhauser, 1993, Hot hands in mutual funds: The persistence of performance, 1974–1988, *Journal of Finance* 48, 93–130.
- Huberman, Gur, and Shmuel Kandel, 1990, Market efficiency and Value Line’s record, *Journal of Business* 63, 187–216.
- Hulbert, Mark, 1996, Mail-order portfolios, *Forbes* 157 (4), 118.
- Ikenberry, David, Josef Lakonishok, and Theo Vermaelen, 1995, Market underreaction to open market share repurchases, *Journal of Financial Economics* 39, 181–208.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.
- Jeng, Leslie A., 1998, Corporate insiders, market makers, and the window of opportunity, Working paper, Harvard University.
- Lakonishok, Josef, Andrei Shleifer, and Robert Vishny, 1994, Contrarian investment, extrapolation, and risk, *Journal of Finance* 49, 1541–1578.
- Lewis, Craig M., Richard J. Rogalski, and James K. Seward, 1997, The information content of Value Line convertible bond rankings, *Journal of Portfolio Management* 24, 42–52.
- Loughran, Tim, and Jay Ritter, 1995, The new issues puzzle, *Journal of Finance* 50, 23–52.
- MacKinlay, A. Craig, 1995, Multifactor explanations do not explain deviations from the CAPM, *Journal of Financial Economics* 38, 3–28.
- Maddala, G. S., 1983, *Limited Dependent and Qualitative Variables in Econometrics* (Cambridge University Press, Cambridge, U.K.).

- Malkiel, Burton, 1995, Returns from investing in mutual funds 1971 to 1991, *Journal of Finance* 50, 549–572.
- Metrick, Andrew, 1997, Performance evaluation with transactions data: The case of investment newsletters, Working paper, Harvard University.
- Morningstar, 1996, Ascent database, April.
- Shelton, John P., 1967, The *Value-Line* contest: A test of the predictability of stock-price changes, *Journal of Business* 40, 251–269.
- Shirk, George S. III, Julie Cuenca, and Greg Carlson, 1997, Platinum fund\$, *Mutual Funds* October, 87.
- Shumway, Tyler, 1997, The delisting bias in CRSP data, *Journal of Finance* 52, 327–340.
- Shumway, Tyler, and Vincent A. Warther, 1997, The delisting bias in CRSP's NASDAQ data and its implications for interpretation of the size effect, Working paper, University of Michigan.
- Stickel, Scott E., 1985, The effect of *Value Line* investment survey rank changes on common stock prices, *Journal of Financial Economics* 14, 121–143.
- Womack, Kent, 1996, Do brokerage analysts' recommendations have investment value?, *Journal of Finance* 51, 137–168.