

Differences of Opinion and the Cross Section of Stock Returns

KARL B. DIETHER, CHRISTOPHER J. MALLOY,
and ANNA SCHERBINA*

ABSTRACT

We provide evidence that stocks with higher dispersion in analysts' earnings forecasts earn lower future returns than otherwise similar stocks. This effect is most pronounced in small stocks and stocks that have performed poorly over the past year. Interpreting dispersion in analysts' forecasts as a proxy for differences in opinion about a stock, we show that this evidence is consistent with the hypothesis that prices will reflect the optimistic view whenever investors with the lowest valuations do not trade. By contrast, our evidence is inconsistent with a view that dispersion in analysts' forecasts proxies for risk.

IN THIS PAPER WE ANALYZE THE ROLE of dispersion in analysts' earnings forecasts in predicting the cross section of future stock returns. We find that stocks with higher dispersion in analysts' earnings forecasts earn significantly lower future returns than otherwise similar stocks. In particular, a portfolio of stocks in the highest quintile of dispersion underperforms a portfolio of stocks in the lowest quintile of dispersion by 9.48 percent per year. This effect is strongest in small stocks, and stocks that have performed poorly over the past year. Our results are robust to various risk-adjustment techniques, and are inconsistent with an interpretation of dispersion in analysts' forecasts as a proxy for risk.

We postulate that dispersion in analysts' earnings forecasts can be viewed as a proxy for differences of opinion among investors. Differences of opinion are typically modeled via dogmatic beliefs or asymmetric information sets, and have been included in numerous models that relax the standard

* Diether and Malloy are at the Graduate School of Business, University of Chicago; Scherbina is at Harvard Business School. This paper combines results from two earlier papers: "Analyst Disagreement and the Cross Section of Stock Returns" by Karl B. Diether and Christopher J. Malloy and "Stock Prices and Differences of Opinion: Empirical Evidence that Prices Reflect Optimism" by Anna Scherbina. The authors thank Nick Barberis, Kent Daniel, Doug Diamond, Eugene Fama, Rick Green (the editor), Kathleen Hagerty, Gilles Hilary, Ravi Jagannathan, Robert Korajczyk, Owen Lamont, Joel Lander, Robert MacDonald, Cade Massey, Lior Menzly, Toby Moskowitz, Jeremy Nalewaik, Lubos Pastor, Chris Polk, Todd Pulvino, David Robinson, Richard Thaler, Beverly Walther, Volker Wieland, Andy Wong, and an anonymous referee. Data on analysts' forecasts was provided by I/B/E/S Inc. under a program to encourage academic research. All remaining errors are our own.

assumption of homogeneous expectations.¹ When included in models in place of the standard assumption of homogeneous expectations, heterogeneous beliefs can change the stock market equilibria. For example, Williams (1977) and Goetzmann and Massa (2001) show that heterogeneous beliefs can affect aggregate market returns, while Miller (1977), Jarrow (1980), Mayshar (1982), Diamond and Verrecchia (1987), Morris (1996), Viswanathan (2000), and Chen, Hong, and Stein (2001) produce cross-sectional asset pricing predictions. And yet, to date there has been little empirical research to investigate how differences of opinion affect asset prices. Since many of the theoretical papers that incorporate differences of opinion produce conflicting cross-sectional implications, the debate can only be resolved with a careful empirical investigation. Our paper takes a step in this direction.

The theoretical debate centers on the intuitively appealing idea (set forth by Miller (1977)) that prices will reflect a more optimistic valuation if pessimistic investors are kept out of the market by high short-sale costs. In Miller's model (as in other price-optimism models such as Morris (1996), Chen et al. (2001), and Viswanathan (2000)), optimists hold the stock because they have the highest valuations. They suffer losses in expectation since the best estimate of the stock value is the average opinion. These price-optimism models suggest that the bigger the disagreement about a stock's value, the higher the market price relative to the true value of the stock, and the lower its future returns. By contrast, this predicted upward bias in prices disappears in models such as Diamond and Verrecchia (1987) and Hong and Stein (2000) that introduce influential rational agents. The Diamond and Verrecchia model depends on the existence of a perfectly rational market maker with unlimited computational abilities who can instantaneously estimate the unbiased stock value, conditional on all publicly available information. On the other hand, the Hong and Stein result relies on perfectly rational arbitrageurs that can eliminate mispricing, perhaps a questionable assumption in light of the burgeoning literature on the limits to arbitrage. For example, Shleifer and Vishny (1997), Gromb and Vayanos (2001), and Chen et al. provide compelling theoretical explanations why arbitrageurs may fail to close the arbitrage opportunity.

In light of the strong assumptions behind the models that predict unbiased prices when opinions diverge, perhaps it is not surprising that our results strongly support the basic Miller (1977) prediction. While we use dispersion in analysts' earnings per share forecasts as a proxy for differences of opinion, Chen et al. (2001) use breadth of mutual fund ownership as a measure of the magnitude of disagreement among investors and come to the same conclusion. The authors find that reductions (increases) in breadth of ownership lead to lower (higher) future returns, as Miller would predict.

¹ Harris and Raviv (1993) were the first to explicitly model investors who were dogmatic about their beliefs. However, their model attempts to explain trading volume rather than stock prices.

In addition, Lee and Swaminathan's (2000) finding that higher trading volume, another possible proxy for differences of opinion, predicts lower future returns is consistent with the Miller story as well.

On a more general level, given boundedly rational agents and limited arbitrage, any friction that prevents the revelation of negative opinions will produce the upward bias in prices that we find in the data. For example, we provide evidence that the incentive structure of analysts (which discourages analysts with poor outlooks from voicing their opinions) can be viewed as an alternative mechanism to the short-sale constraints emphasized in the Miller (1977) model.

Finally, our results strongly reject the interpretation of dispersion in analysts' forecasts as a measure of risk. While we show that dispersion is positively related to earnings variability, standard deviation in returns, and market β , the strong negative relation between dispersion and future returns is hard to reconcile within a risk-based framework. In fact, Gebhardt, Lee, and Swaminathan (2001) include dispersion in analysts' forecasts as one of their proxies for risk in their attempts to explain implied cost of capital estimates, and are surprised to find the same negative relation that we document here. However, they do not examine the source and robustness of this result in any depth. Other prior evidence of a dispersion effect is mixed. Cragg and Malkiel (1982) report a positive relation between dispersion in forecasts and future returns, but they do so using an extremely limited sample (including only 175 companies, with yearly forecasts from 1961 to 1969). By contrast, Ackert and Athanassakos (1997) provide evidence that standard deviation in analysts' forecasts is negatively related to future returns, but do so using a sample of only 167 firms.

The layout of the paper is as follows. In the next section, we discuss three competing hypotheses, the data, and our sample characteristics. Section II describes our procedure of forming portfolios on the basis of dispersion in analysts' forecasts and presents our initial results. Section III explores a risk-based explanation for our findings, while Section IV presents various robustness checks. Finally, Section V evaluates potential explanations for our results, and Section VI provides a summary and concludes.

I. Methodology

A. *Competing Hypotheses*

Our main objective in this paper is to analyze the role of dispersion in analysts' earnings forecasts in predicting the cross section of future stock returns. In particular, we test three competing hypotheses about the relation between dispersion in forecasts and future returns.

The first hypothesis views dispersion in forecasts as a proxy for differences of opinion among investors. This hypothesis centers on Miller's (1977) conjecture that whenever stock valuations differ, equity prices tend to reflect the view of the more optimistic investors, leading to low future returns.

Miller argues that when opinions diverge and short-sale constraints bind, investors with pessimistic valuations do not sell the stock, and optimistic demand pushes stock prices up. Miller's verbal model assumes that investors are boundedly rational in the sense that they are overconfident about their own valuations.² However, his argument can be easily modified to accommodate investors who are not overconfident, but rather make inaccurate inferences about others' signals. For example, if the short-sale constraint binds, an informed investor with a low valuation will not sell the stock. However, an informed investor may also stay out of the market if she agrees with the market price. Boundedly rational market participants may erroneously assume that the informed investor did not trade simply because her valuation equaled the market price. If so, boundedly rational traders will not revise their valuations downward after observing that the informed investor did not trade, and the stock price will remain unchanged. In this scenario, negative information will not be appropriately incorporated, and the market price will be upwardly biased.

The first hypothesis suggests that the larger the disagreement about a stock's value (either caused by overconfidence or by inaccurate inferences about others' signals), the higher the market price relative to the true value of the stock, and the lower its future returns. And although Miller's (1977) model focuses on short-sale constraints, any friction that prevents the revelation of negative opinions will produce this negative relation between dispersion in forecasts and future returns.

The second hypothesis also views dispersion in forecasts as a proxy for differences of opinion, but states that market prices will be unbiased when opinions about the correct valuation diverge, and so future returns will be independent of the current level of disagreement about the stock value. This prediction is stated in Diamond and Verrecchia (1987). As a price-setting mechanism, the authors model a rational market-maker who correctly sets bid and ask prices, conditioned on the possibility that some informed investors with low valuations are constrained from selling stock. The outcome is that stock prices are unbiased: No trading profits can be made based on the publicly available information. The key assumption of the model is that the market-maker has perfect knowledge of his economic environment and can perform Bayesian updating in the short time between two consecutive trades. Similarly, Hong and Stein (2000) achieve unbiased prices by introducing competitive, risk-neutral, and perfectly rational arbitrageurs who do not face short-sale constraints. The arbitrageurs correctly infer the expected stock value from the actions of informed overconfident short-sale-constrained investors in a price auction. The market clears at the price that is equal to expected stock value. Thus, both of these models predict no relation between observed dispersion in analysts' forecasts and future returns.

² See Gervais, Heaton, and Odean (2000) and Daniel, Hirshleifer, and Subrahmanyam (2001) for models that explicitly model overconfidence; in these models, investors are overconfident in the sense that they overestimate the precision of their private signals.

The final hypothesis views dispersion in analysts' earnings forecasts as a proxy for risk. Investors who are not well diversified will demand to be compensated for the idiosyncratic risk of the securities they hold (see, e.g., Merton (1987)). Since dispersion in analysts' forecasts likely indicates a more volatile, less predictable earnings stream, stocks with higher dispersion in analysts' forecasts should earn higher future returns, and dispersion in analysts' forecasts will hold explanatory power beyond the standard risk factors.

In summary, the first hypothesis predicts a negative relation between dispersion and future returns. The second hypothesis predicts no relation. And the third hypothesis predicts a positive relation.

B. Data and Sample Characteristics

Returns are drawn from the Center for Research in Securities Prices (CRSP) Monthly Stocks Combined File, which includes NYSE, AMEX, and Nasdaq stocks. For many of the tests in the paper, firms must also have data in COMPUSTAT on book equity for the fiscal year ending in calendar year $t - 1$.³ The data on analysts' earnings estimates are taken from the Institutional Brokers Estimate System (I/B/E/S).

The nature of our inquiry makes the standard-issue I/B/E/S data set unsuitable for our purposes because of the following reporting inaccuracy. I/B/E/S analysts' forecasts are adjusted historically for stock splits in order to produce a smooth time series of earnings per share estimates. For example, analysts' earnings per share estimates for December 1976 are reported on the basis of the number of shares outstanding as of today, rather than the number of shares outstanding as of December 1976. However, after dividing historical analysts' forecasts by a split adjustment factor, I/B/E/S rounds the estimate to the nearest cent. For example, for a stock that has split 10-fold, actual earnings per share estimates of 10 cents and 14 cents would be reported as 1 cent per share each. I/B/E/S would then include an adjustment factor of 10 in the Adjustment File, so that the "unadjusted" earnings per share estimates would be 10 cents each, rather than the correct values of 10 cents and 14 cents, respectively. The observed variance of analysts' forecasts would then be zero, when in fact it is positive. Observations with no dispersion in analysts' forecasts would thus contain ex-post information about the future success of the firm, embedded in the number of stock splits. We conducted our analysis on the raw forecast data, unadjusted for stock splits (provided by I/B/E/S on request). Had we utilized the commonly used I/B/E/S files, we would have found that as a group, stocks with the lowest levels of dispersion earned high abnormal returns, simply because we would have been picking up observations of firms that did well in the future.⁴

³ Limiting the sample for all tests to include only those firms that have data in COMPUSTAT on book equity does not significantly alter any of our results.

⁴ We believe we are the first to document this error. This finding has important implications for prior and subsequent empirical research that uses I/B/E/S data to investigate issues related to dispersion (e.g., herding among security analysts).

I/B/E/S data includes U.S. Detail History and Summary History data sets. The Summary History data set contains the summary statistics on analyst forecasts, such as mean and standard deviation values. These variables are calculated on the basis of all outstanding forecasts as of (ordinarily) the third Thursday of each month. The Detail History file contains individual analysts' forecasts organized by the date on which the forecast was issued. Each record also contains a *revision date*, that is, the date on which the forecast was last confirmed as accurate.

To address concerns that the U.S. Summary History file makes use of analysts' forecasts that are no longer current, we calculate forecast statistics from the Detail History file and match them to the numbers in the Summary History file. We compute month-end averages and standard deviations from the individual estimates in the Detail History file by extending each forecast until its revision date. For example, if the forecast was made in May and was last confirmed as accurate in July, it will be used in our computation of averages and standard deviations for May, June, and July. If an analyst makes more than one forecast in a given month, only the last forecast is used in our calculations.

In some records, a revision date precedes the actual forecast date, which constitutes an error on the part of I/B/E/S. In this case, the forecast will be assumed valid only for the month in which it was made. In addition, occasional aberrant observations in the Detail History file are not picked up by the Summary History file. Even given these inconsistencies, the mean and standard deviation values calculated from the Detail History file data closely track the values in the Summary History file. Generally, when compiling statistics from the Detail History file, we obtain fewer monthly forecasts than reported in the Summary History file because we assume that forecasts are no longer valid past the revision date, which is not the case for the Summary History file conventions. We compute our portfolio results in Section II using data in both the Summary History and Detail History files. Since the results are very similar, we report only the results obtained using the Summary History file throughout the paper.

For each stock in CRSP, we set the coverage in any given month equal to the number of I/B/E/S analysts who provide fiscal year one earnings estimates that month.⁵ Obviously each stock must be covered by two or more analysts during that month, since we define dispersion as the standard deviation of earnings forecasts scaled by the absolute value of the mean earnings forecast.

There is an unavoidable issue of data deletion when using the I/B/E/S database, since many firms listed on CRSP are not covered by I/B/E/S. However, LaPorta (1996) presents evidence showing that the performance of stocks in the I/B/E/S sample is almost identical to those in CRSP. The most salient feature of the sample of stocks in the intersection of CRSP, COM-

⁵ All of the tests in this paper were replicated for different fiscal periods (e.g., two years). Discussion of these and other robustness checks is presented in Section IV.

Table I
Summary Statistics on Analyst Coverage:
January 1976 to December 2000

Descriptive statistics for NYSE, AMEX, and Nasdaq stocks during the period from January 1976 to December 2000. Panel A reports statistics for all stocks, while Panel B limits the sample to only eligible stocks. A stock is “eligible” to be included in our analysis if it has a one (fiscal) year I/B/E/S earnings estimate, is covered by two or more analysts, and has a price greater than five dollars. Panels C and D report the same statistics as in Panels A and B for January 1983, and breaks the sample down by size decile (formed using NYSE-based market capitalization deciles).

Summary Statistics for 1976–2000						
Date	Panel A: All CRSP Stocks			Panel B: Eligible Stocks		
	Number of Firms	Mean Size (Millions)	Percentage of Firms Eligible	Number of Firms	Mean Size (Millions)	Mean No. of Estimates
01/1976	5,065	184.6	12.8	650	910.2	6.47
12/1979	4,781	267.5	29.2	1,395	621.5	7.47
12/1982	5,438	305.6	31.9	1,735	766.0	8.78
12/1985	6,227	362.4	35.4	2,203	938.5	9.95
12/1988	6,798	408.6	34.0	2,309	1,094.1	9.82
12/1991	6,643	610.1	35.5	2,361	1,590.1	9.23
12/1994	8,029	633.8	40.7	3,269	1,423.9	8.26
12/1997	8,839	1,241.2	46.8	4,140	2,473.1	7.19
12/2000	7,823	2,032.9	40.5	3,166	4,740.6	7.79

Summary Statistics by Size for 01/1983						
NYSE Breakpoints	Panel C: All CRSP Stocks			Panel D: Eligible Stocks		
	Number of Firms	Mean Size (Millions)	Percentage of Firms Eligible	Number of Firms	Mean Size (Millions)	Mean No. of Estimates
Decile 1	2,969	15.2	4.1	123	34.1	2.76
Decile 2	517	61.2	31.9	165	62.7	3.39
Decile 3	406	96.9	50.5	205	97.8	3.81
Decile 4	323	151.5	69.3	224	151.0	4.90
Decile 5	258	229.4	79.1	204	229.2	6.79
Decile 6	231	343.3	84.0	194	343.9	8.26
Decile 7	191	528.5	89.0	170	526.8	11.22
Decile 8	173	809.1	89.0	154	809.9	13.60
Decile 9	160	1,381.2	92.5	148	1,366.2	16.91
Decile 10	165	5,709.0	92.1	152	5,079.4	19.59

PUSTAT, and I/B/E/S is that it is heavily tilted towards big stocks. As Hong, Lim, and Stein (2000) note, the smallest firms are simply not covered. Table I provides an overview of the extent of analyst coverage for NYSE, AMEX, and Nasdaq stocks. As shown in Panel A, in January of 1976, only 12.8 percent of stocks are even eligible to be included in our sample. There is a marked

deepening of coverage by the end of 1979, with the fraction of eligible firms increasing to 29.2 percent. But Panel C shows that even by January of 1983 (the beginning point of our chosen data sample), only 4.1 percent of firms in the first decile are eligible for our sample. Additionally, fewer firms in our sample are delisted for cause relative to the entire CRSP universe.⁶ This suggests that financially distressed firms may be underrepresented in our sample.

We choose the time period of January 1983 through November 2000 for our portfolio tests in Section II for two reasons. First, by January of 1983, the cross section of stocks has substantial variation in size and book-to-market ratio. Second, the data in the Detail History file are not available prior to 1983. Since we replicate portfolio results from Section II with the Detail History file data, we need to use post-1982 data. Using the entire sample makes our results slightly stronger.

II. Portfolio Strategies

In this section, we assign stocks to portfolios based on certain characteristics, such as dispersion in analyst forecasts, in order to draw conclusions about average returns for these classes of stocks. This is a standard approach in asset pricing, which reduces the variability in returns. The methodology employed here was pioneered by Jegadeesh and Titman (1993). Following Jegadeesh and Titman (2001), stocks with share price lower than five dollars are omitted in order to ensure that the results are not driven by small, illiquid stocks or by bid-ask bounce.

Dispersion is defined as the standard deviation of earnings forecasts scaled by the absolute value of the mean earnings forecast. If the mean earnings forecast is zero, then the stock is assigned to the highest dispersion category. Excluding observations with a mean earnings forecast of zero does not significantly affect the portfolio returns. An alternative definition of dispersion as the ratio of standard deviation of earnings forecasts to the book equity per share does not significantly affect our results.

Each month, we assign stocks into five quintiles based on dispersion in analyst earnings forecasts as of the *previous* month. After assigning stocks into portfolios, stocks are held for one month. We calculate the monthly portfolio return as the equal-weighted average of the returns of all the stocks in the portfolio. The last column of Table II shows that this sort produces a strong negative relation between average returns and dispersion in analysts' earnings forecasts. The annual return on the $D1 - D5$ strategy is 9.48 percent, and is strongly significant. Of that $D1 - D5$ spread, 68.4 percent comes from the short side of the trade, that is, the difference between medium-dispersion stocks and high-dispersion stocks. This evidence is strongly supportive of the Miller (1977) hypothesis laid out in Section I.

⁶ For example, only 14.5 percent of the firms in our sample are delisted for cause (delist codes of 400 and above on CRSP), compared to 30.3 percent for the universe of CRSP stocks.

Table II
Mean Portfolio Returns by Size and Dispersion
in Analysts' Forecasts

Each month stocks are sorted in five groups based on the level of market capitalization as of the third Thursday of the previous month. Stocks in each size group are then sorted into five additional groups based on dispersion in analyst earnings forecasts for the previous month. Dispersion is defined as the ratio of the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts to the absolute value of the mean forecast, as reported in the I/B/E/S Summary History file. Stocks with a mean forecast of zero are assigned to the highest dispersion groups, and stocks with a price less than five dollars are excluded from the sample. Stocks are held for one month, and portfolio returns are equal-weighted. The time period considered is February 1983 through December 2000. The table reports average monthly portfolio returns; *t*-statistics in parentheses are adjusted for autocorrelation.

Mean Returns						
Dispersion Quintiles	Size Quintiles					All Stocks
	Small				Large	
	S1	S2	S3	S4	S5	
<i>D1</i> (low)	1.52	1.45	1.50	1.51	1.48	1.48
<i>D2</i>	1.12	1.40	1.41	1.18	1.35	1.36
<i>D3</i>	0.99	1.20	1.32	1.11	1.36	1.23
<i>D4</i>	0.76	1.07	1.18	1.33	1.33	1.12
<i>D5</i> (high)	0.14	0.56	0.83	1.03	1.20	0.69
<i>D1-D5</i>	1.37 ^a	0.89 ^a	0.67 ^b	0.48	0.29	0.79 ^a
<i>t</i> -statistic	(5.98)	(3.12)	(2.41)	(1.55)	(0.94)	(2.88)

Mean Dispersion						
Dispersion Quintiles	Size Quintiles					All Stocks
	S1	S2	S3	S4	S5	
	<i>D1</i> (low)	0.010	0.011	0.012	0.014	
<i>D2</i>	0.039	0.033	0.030	0.028	0.025	0.030
<i>D3</i>	0.081	0.062	0.053	0.047	0.039	0.053
<i>D4</i>	0.172	0.125	0.103	0.086	0.067	0.105
<i>D5</i> (high)	1.256	0.963	0.813	0.722	0.462	0.852

^{a,b} Statistically significant at the one and five percent levels, respectively.

Table II also presents two-way cuts on size and dispersion, to test if we are simply capturing a size effect. Each month, we assign stocks to one of five quintiles based on the level of market capitalization as of the third Thursday of the *previous* month. We rank stocks in each size quintile into five further quintiles based on dispersion in analyst earnings forecasts as of the *previous* month. Each of the 25 resulting portfolios contains an average of 113 stocks, if the Summary History file is used.⁷ Table II shows that the average monthly

⁷ We use conditional sorts here for consistency, since some of our later three-way sorts result in portfolios that are very thin if we use independent sorts. Our results are not significantly affected if we use independent sorts instead.

Table III
Mean Portfolio Returns by NYSE Market Capitalization Deciles

Each month stocks in each NYSE-based market capitalization decile (which is determined as of the third Thursday of the previous month) are sorted into five groups based on dispersion in analyst earnings forecasts for the previous month. Dispersion is defined as the ratio of the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts to the absolute value of the mean forecast, as reported in the I/B/E/S Summary History file. Stocks with a mean forecast of zero are assigned to the highest dispersion groups, and stocks with a price less than five dollars are excluded from the sample. Stocks are held for one month, and portfolio returns are equal-weighted. Only portfolios in the largest seven deciles are reported since there are fewer than eight stocks on average within each size/dispersion portfolio for size deciles 1 to 3. The time period considered is February 1983 through December 2000. The table reports average monthly portfolio returns; *t*-statistics in parentheses are adjusted for autocorrelation.

Mean Returns							
NYSE Market Capitalization Deciles (Average Number of Stocks in Each Portfolio)							
Dispersion Quintiles	4 (27)	5 (47)	6 (68)	7 (91)	8 (119)	9 (149)	10 (189)
<i>D1</i> (low)	1.21	1.22	1.34	1.42	1.50	1.47	1.45
<i>D2</i>	1.41	1.51	1.26	1.39	1.38	1.18	1.31
<i>D3</i>	1.38	1.31	1.41	1.04	1.30	1.16	1.20
<i>D4</i>	0.72	0.91	0.96	1.11	0.88	1.08	1.18
<i>D5</i> (high)	0.61	0.53	0.33	0.32	0.49	0.64	1.03
<i>D1–D5</i>	0.60	0.69	1.01 ^a	1.09 ^a	1.01 ^a	0.84 ^b	0.41
<i>t</i> -statistic	(1.48)	(1.86)	(3.08)	(3.61)	(3.93)	(2.33)	(1.49)

^{a,b} Statistically significant at the one and five percent levels, respectively.

return differential between low- and high-dispersion portfolios declines as the average size increases. While the return differential between the low- and high-dispersion stocks is positive and highly significant for the smaller stocks, it becomes insignificant for stocks in the two highest market-capitalization quintiles.⁸ In particular, the *D1 – D5* strategy for the smallest quintile earns an enormous 16.4 percent annual return on average. Thus, it does not appear that we are simply picking up a size effect, since the two-way sorts still produce a strong negative relation between average returns and dispersion for the first three quintiles.

However, I/B/E/S-based market-capitalization quintiles might be misleading because the universe of I/B/E/S firms is comprised of only relatively large firms—especially among the firms covered by at least two analysts. To be consistent with reporting convention, we also form portfolios using NYSE-based market capitalization deciles. Table III reports returns on these port-

⁸ Results do not change significantly when data are restricted to stocks covered by at least three analysts.

folios. There are not enough stocks in the three lowest deciles to produce representative average portfolio returns for these classes of stocks. The return differential for the fifth through ninth NYSE size deciles turns out to be statistically significant, indicating that the dispersion result holds for even relatively large stocks.

A. *Sorting by Size and Book-to-Market*

In this section, we triple-sort on size, book-to-market (BE/ME) ratio, and dispersion to test if we are merely picking up a book-to-market effect in returns. Since low-book-to-market stocks tend to have higher levels of market capitalization, we try to control for the strong size-related pattern observed in Table II (namely, that within small stocks, we see a huge return differential between low- and high-dispersion stocks) by sorting on size *and* book-to-market. Stocks are first sorted into three categories based on the level of market capitalization at the end of the previous month. Within each size category, the stocks are sorted into three groups based on the book-to-market ratio, and finally into three dispersion groups within each resulting group. Book equity (BE) is defined as the COMPUSTAT book value of stockholders' equity, plus balance sheet deferred taxes and investment tax credit (if available), minus the book value of preferred stock. Depending on availability, we use redemption, liquidation, or par value (in that order) to estimate the book value of preferred stock. To insure that the book equity figure is known to the market before the returns that it is used to explain, we match the yearly book equity figure for all fiscal years ending in calendar year $t - 1$ with returns starting in July of year t . This figure is then divided by market capitalization (ME) at month $t - 1$ to form the BE/ME ratio, so that the BE/ME ratio is updated each month. Finally, within each size and book-to-market group, stocks are sorted into three further subdivisions based on dispersion in earnings-per-share forecasts as of the previous month.

Table IV presents the returns on the resulting 27 portfolios. Note that each portfolio contains an average of 55 stocks. The return differential between low- and high-dispersion stocks is still significant in four of the nine categories, indicating that the dispersion effect is not simply capturing a book-to-market effect. Again we observe the strong size pattern first depicted in Table II, with small stocks exhibiting the largest return differential. However, no clear pattern emerges with respect to book-to-market ratios. While the return differential between low- and high-dispersion stocks is larger for value stocks (those with high BE/ME ratios), this result is not overwhelming. On the other hand, the lower half of the table is more revealing. In particular, value stocks have much larger average dispersion in forecasts than growth stocks, implying that book-to-market ratios are positively related to dispersion in analysts' forecasts.

Despite having lower average values of dispersion, growth stocks in the high-dispersion portfolios earn returns almost as low as value stocks, indicating that the same amount of dispersion in analysts' earnings per share

Table IV

Mean Portfolio Returns by Size, Book-to-Market, and Dispersion

Each month stocks are sorted into three groups based on the level of market capitalization at the end of the previous month. Each size group is then sorted into three book-to-market groups. The book-to-market ratio is computed by matching the yearly *BE* figure for all fiscal years ending in calendar year $t - 1$ to returns starting in July of year t ; this figure is then divided by market capitalization at month $t - 1$ to form the book-to-market ratio, so that the book-to-market ratio is updated each month. Each size and book-to-market group is further sorted into three dispersion groups. Dispersion is defined as the ratio of the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts to the absolute value of the mean forecast, as reported in the I/B/E/S Summary History file. Stocks with a mean forecast of zero are assigned to the highest dispersion groups, and stocks with a price less than five dollars are excluded from the sample. Stocks are held for one month, and portfolio returns are equal-weighted. The sample period is February 1983 through December 2000; t -statistics in parentheses are adjusted for autocorrelation.

Mean Returns									
Dispersion	Low Book-to-Market			Medium Book-to-Market			High Book-to-Market		
	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap
Low	1.35	1.42	1.40	1.64	1.50	1.32	1.50	1.58	1.42
Medium	1.27	1.40	1.20	1.29	1.11	1.27	1.13	1.46	1.37
High	0.73	1.28	1.38	0.67	0.86	1.12	0.70	1.10	1.34
Low-high t -statistic	0.63 ^b (2.10)	0.14 (0.47)	0.02 (0.06)	0.96 ^a (4.09)	0.64 ^b (2.37)	0.21 (0.96)	0.80 ^a (4.09)	0.48 (1.70)	0.08 (0.25)

Mean Dispersion									
Dispersion	Low Book-to-Market			Medium Book-to-Market			High Book-to-Market		
	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap
Low	0.04	0.03	0.02	0.05	0.03	0.02	0.10	0.04	0.03
Medium	0.08	0.05	0.03	0.09	0.05	0.04	0.21	0.11	0.07
High	0.49	0.34	0.20	0.56	0.37	0.26	1.04	0.86	0.59

^{a,b} Statistically significant at the one and five percent levels, respectively.

forecasts would cause lower returns for growth stocks than value stocks. This is not surprising, given that the same amount of disagreement about earnings per share should translate into a higher level of disagreement about the intrinsic value of a growth stock as compared to a value stock. This, in turn, would translate into a higher price run-up for growth stocks, and lower subsequent returns. However, low-dispersion growth stocks earn lower returns than low-dispersion value stocks because of the famous value premium, and that translates into lower spreads between low- and high-dispersion portfolios for growth stocks.

Table V
Mean Portfolio Returns by Size, Momentum, and Dispersion

Each month stocks are sorted into three groups based on the level of market capitalization at the end of the previous month. Each size group is then sorted into three momentum groups; momentum is computed based on past returns from $t - 12$ to $t - 2$, where "Losers" is an equal-weighted portfolio of stocks in the *worst*-performing 33 percent. Each size and momentum group is further sorted into three dispersion groups. Dispersion is defined as the ratio of the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts to the absolute value of the mean forecast, as reported in the I/B/E/S Summary History file. Stocks with a mean forecast of zero are assigned to the highest dispersion groups, and stocks with a price less than five dollars are excluded from the sample. Stocks are held for one month, and portfolio returns are equal-weighted. The sample period is February 1983 through December 2000; t -statistics in parentheses are adjusted for autocorrelation.

Mean Returns									
Dispersion	Losers			Winners					
	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap
Low	0.92	1.02	1.27	1.29	1.45	1.36	1.79	1.79	1.54
Medium	0.47	0.69	1.12	1.11	1.32	1.18	1.84	1.80	1.58
High	0.09	0.63	0.93	0.82	1.22	1.06	1.56	1.81	1.73
Low-high	0.82 ^a	0.40	0.35	0.48 ^a	0.23	0.30	0.23	-0.02	-0.19
t -statistic	(3.65)	(1.56)	(1.46)	(2.76)	(1.12)	(1.61)	(1.17)	(-0.10)	(-0.83)

Mean Dispersion									
Dispersion	Losers			Winners					
	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap	Small Cap	Mid Cap	Large Cap
Low	0.10	0.05	0.03	0.04	0.03	0.02	0.03	0.02	0.02
Medium	0.21	0.10	0.06	0.09	0.05	0.04	0.06	0.15	0.04
High	1.18	0.85	0.55	0.61	0.33	0.22	0.39	0.28	0.25

^{a,b} Statistically significant at the one and five percent levels, respectively.

B. Sorting by Size and Momentum

Our final portfolio strategy entails a three-way cut on size, momentum, and dispersion, to rule out the possibility that the momentum effect documented in Jegadeesh and Titman (1993) is driving our results. Stocks are first sorted into three categories based on the level of market capitalization at the end of the previous month. Within each size category, the stocks are sorted into three groups based on past returns from $t - 12$ to $t - 2$ (as in Fama and French (1996)). Lastly, within each size and momentum group, stocks are sorted into three further subdivisions based on dispersion in earnings-per-share forecasts as of the previous month.

Table V presents the returns on the resulting 27 portfolios. Again the return differential between low- and high-dispersion stocks is still strongly

significant in two of the nine categories, indicating that the dispersion effect is not simply capturing a momentum effect. However, a clear pattern emerges with respect to momentum, as the return differential between low- and high-dispersion stocks is strongest in stocks that have performed poorly over the past year (“losers”). For example, the $D1 - D3$ strategy for the smallest quintile of losers earns a 9.84 percent average annual return.

III. Regression Tests

Fama and French (1996) show that sorting stocks on variables such as BE/ME (the book-to-market ratio), E/P (the earnings-to-price ratio), or C/P (the cash-flow-to-price ratio) can produce a strong (in their case, positive) ordering of returns across deciles. They also argue, however, that estimates of three-factor time-series regressions indicate that the three-factor model captures these patterns in average returns. The regression intercepts are uniformly small, and the Gibbons, Ross, and Shanken ((1989), hereafter, GRS) tests never come close to rejecting the hypothesis that the three-factor model describes average returns. Along these lines, we employ similar tests to see if the three-factor model can capture the return patterns observed in Tables II to V.

A. Multifactor Time-Series Tests

Fama and French (1996) argue that many of the CAPM average-return anomalies are related, and that they are captured by the three-factor model in Fama and French (1993). In the model, $R_M - R_F$ is the excess return on a proxy for the market portfolio, SMB is the difference between the return on a portfolio of small stocks and the return on a portfolio of large stocks, and HML is the difference between the return on a portfolio comprised of high book-to-market stocks and the return on a portfolio comprised of low book-to-market stocks. The variable HML represents the value premium; high book-to-market stocks are value stocks, and low book-to-market stocks are growth stocks. Similarly, the variable SMB represents the size premium.⁹

Believers in the three-factor model suggest that it is an equilibrium pricing model in the spirit of Merton's (1973) intertemporal capital asset pricing model (ICAPM). In an ICAPM framework, investors care about market risk, but they also care about hedging against more specific aspects of their investment consumption decisions. For example, investors may wish to hedge against the relative prices of consumption goods. In this view of the three-factor model, SMB and HML are viewed as portfolios that proxy for factors that are of special hedging concern to investors, and β_i , s_i , and h_i are factor loadings or sensitivities, and $E(R_{Mt}) - R_{Ft}$, $E(SMB_t)$, $E(HML_t)$ are expected risk premiums. The main testable implication of the model in this setting is that $\alpha_i = 0$ for all assets. We test this hypothesis using a formulation of the Hotelling t^2 test (GRS (1989)). Extensions of the model (see,

⁹ See Fama and French (1996) for details on the construction of these factors.

e.g., Carhart (1997)) add a momentum factor to this model, included to capture the medium-term continuation of returns documented in Jegadeesh and Titman (1993). We include the variable *UMD* to capture the momentum premium; it equals the difference between the return on a portfolio comprised of stocks with high returns from $t - 12$ to $t - 2$ and the return on a portfolio comprised of stocks with low returns from $t - 12$ to $t - 2$.

Table VI reports estimates of three- and four-factor time-series regressions for monthly excess returns on five equal-weighted portfolios formed on dispersion in analysts' forecasts. The estimated intercepts indicate that the three-factor model leaves a large negative unexplained return for the portfolio of stocks in the highest dispersion quintile. The loadings indicate that low-dispersion stocks behave like big, value stocks (they load less on *SMB* but heavily on *HML*), while high-dispersion stocks behave like small, value stocks. An important point to mention, however, is that these loadings give a potentially misleading picture of the risk characteristics of these stocks, because they are highly variable. For example, in results not reported here, we find that in the sample period from 1983 to 1999, high-dispersion stocks load very heavily on *HML*, while low-dispersion stocks load more on growth. However, in the year 2000, these patterns were strongly reversed, as high-dispersion stocks behaved very much like growth stocks.

The GRS test rejects the hypothesis that the three-factor model explains the average returns on the five dispersion quintiles. The rejection of the three factor model is testimony to the explanatory power of the regressions (the average of the five regression R^2 is 0.932), so that even small intercepts are distinguishable from zero. Note, however, that the intercept on the high-dispersion quintile is a nontrivial 0.58 percent.

Table VI also reports estimates of the four-factor time-series regression for monthly excess returns on the dispersion quintiles. Again, the estimated intercepts indicate that the model leaves a large negative unexplained return for the portfolio of stocks in the highest dispersion quintile. Meanwhile, the estimated intercept on the lowest dispersion quintile is positive and marginally significant. The pattern of loadings is identical to the pattern found in the three-factor regressions, except that the stocks in the highest dispersion quintile now behave like small, distressed losers (they load more on *SMB*, *HML*, and less on *UMD*). This lowers the negative unexplained return for the portfolio of stocks in the highest decile, but the model still predicts higher postformation returns on these stocks. The GRS test again rejects the hypothesis that the four-factor model explains the average returns on the dispersion quintiles. Thus it appears that neither model can account for the return patterns displayed in Tables II to V.

B. Explaining Dispersion

In this section, we run a series of Fama and MacBeth (1973) cross-sectional regressions in order to illustrate the relation between dispersion in

Table VI
Time-Series Tests of Three- and Four-Factor Models
for Dispersion Equal-Weighted Quintiles

This table reports estimates of the Fama–French three-factor model, as well as a four-factor model, $E(R_{it} - R_{ft}) = b_{iM}(R_M - R_F) + s_i SML_t + h_i HML_t + m_i UMD_t$, for monthly excess returns on the equal-weighted dispersion quintiles. The market premium ($R_M - R_F$) uses the CRSP NYSE/AMEX/Nasdaq value-weighted index. The variables *HML* and *SMB* are created using the same methodology as Fama and French (1996). The momentum premium (*UMD*) is the difference between the return on a portfolio comprised of stocks with high returns from $t - 12$ to $t - 2$ and the return on a portfolio comprised of stocks with low returns from $t - 12$ to $t - 2$. The dispersion equal-weighted quintiles are formed as in Table II. The sample period is February 1983 to December 2000. Stocks with a price less than five dollars are excluded from the sample, and *t*-statistics (Newey-adjusted) are in parentheses.

Portfolio	Alpha (%)	Factor Sensitivities				Adj. R^2 (%)
		$R_M - R_F$	SMB	HML	UMD	
D1 (low-dispersion)	0.13 (0.88)	1.06 (39.78)	0.32 (3.17)	0.26 (2.83)		88.42
	0.27 (2.14)	1.06 (43.33)	0.33 (4.17)	0.21 (2.50)	-0.13 (-2.01)	89.18
D2	0.03 (0.30)	1.06 (44.13)	0.40 (5.76)	0.22 (2.95)		92.64
	0.17 (1.57)	1.06 (56.96)	0.41 (7.98)	0.17 (2.44)	-0.12 (-2.81)	93.35
D3	-0.07 (-0.82)	1.08 (42.71)	0.49 (6.59)	0.15 (2.91)		95.08
	0.09 (1.26)	1.07 (52.92)	0.51 (9.46)	0.09 (1.98)	-0.14 (-4.12)	95.94
D4	-0.18 (-2.56)	1.11 (52.79)	0.67 (12.74)	0.14 (3.31)		96.47
	-0.03 (-0.40)	1.10 (63.19)	0.68 (19.12)	0.08 (2.00)	-0.14 (-5.21)	97.17
D5 (high-dispersion)	-0.58 (-4.78)	1.15 (38.53)	0.88 (16.16)	0.13 (3.19)		93.74
	-0.35 (-2.62)	1.15 (45.66)	0.91 (26.12)	0.05 (1.11)	-0.20 (-5.18)	94.92
					GRS(3F) =	5.51
					GRS(4F) =	2.62

analysts' forecasts and various fundamentals. We also explore how dispersion in forecasts is distinct from other common measures of risk, such as standard deviation of returns, market β , earnings variability, and so forth. The goal is to determine what exactly makes dispersion in analysts' forecasts such an important valuation indicator.

Specifically, we attempt to disentangle the effect of dispersion in analysts' forecasts from stock fundamentals and additional measures of risk by performing a series of yearly cross-sectional regressions of dispersion (measured seven months before the end of the fiscal year) on various lagged firm characteristics. Aside from the fiscal year-end lagged values of market capitalization, book-to-market ratio, sales-to-assets ratio, and debt-to-book ratio, we include momentum (the return from $t - 12$ months to $t - 2$ months), market β (estimated using the past 36 to 60 months of data), *residual coverage*, the standard deviation of daily returns over the past 250 days (lagged one month), the average adjusted volume over the past 250 days (lagged one month), the average adjusted turnover over the past 250 days (lagged one month), the mean age of the forecasts, and a measure of earnings variability ($\sigma(\text{eps})$). *Residual coverage* is the residual from yearly regressions of $\ln(1 + \text{analyst coverage})$ on $\ln(M)$ and $\ln(B/M)$. A firm's adjusted volume is the daily volume from the last 250 days, divided by the mean volume of the firm's exchange; adjusted turnover is defined similarly. Finally, $\sigma(\text{eps})$ is the standard deviation of earnings per share, divided by the absolute value of mean earnings per share over the past five years; $\sigma(\text{eps})$ is lagged one year. We run the regressions yearly, rather than monthly, because the dispersion measures for each stock are highly autocorrelated.

The results, reported in Table VII, cast doubt on the interpretation of dispersion in analysts' forecasts as a proxy for risk. Dispersion is strongly positively related to market β , earnings variability, and standard deviation of past returns. Since these variables are commonly used measures of risk, we expect them to be positively related to returns. Therefore, the negative relation between dispersion and future returns is hard to explain within a traditional risk-based framework. Confirming the evidence presented in Tables II to IV, dispersion is positively related to *BE/ME*, and negatively related to size and momentum. Additionally, dispersion is positively related to leverage, turnover, and volume, but negatively related to sales and the mean age of the forecasts. Controlling for firm size, we find that dispersion is positively correlated with analyst coverage, indicating that there is a high demand for additional opinions in situations where earnings are difficult to forecast. Based on these results, we feel that the correct interpretation of dispersion in analyst earnings forecasts is as a proxy for differences in opinion among investors. For example, turnover is often used as a proxy for differences of opinion, and we find these two measures to be strongly positively related in our cross-sectional regressions.

IV. Robustness Checks

To ascertain that the persistently high returns we have documented thus far are not caused by a statistical fluke or an obvious explanation, we employ additional trading strategies and regression tests to demonstrate robustness.

Table VII
FM Regressions of Dispersion in Analysts' Forecasts
on Lagged Firm Characteristics

Fama and Macbeth (1973) cross-sectional regressions are run every fiscal year from 1983 to 2000. Dispersion in analysts' forecasts (measured seven months before the end of the fiscal year) is regressed on β (estimated using the past 36 to 60 months of data), M (the lagged fiscal year-end market capitalization), B/M (the lagged fiscal year-end book-to-market ratio), $R_{-12:-2}$ (the return from $t - 12$ months to $t - 2$ months), $RCov$ (the residual from yearly regressions of $\ln(1 + \text{analyst coverage})$ on $\ln(M)$ and $\ln(B/M)$), $\sigma(R)$ (the standard deviation of daily returns over the past 250 days, lagged one month), Vol (the average adjusted daily volume from the last 250 days, lagged one month), $Turn$ (the average adjusted turnover from the last 250 days, lagged one month), D/B (the lagged fiscal year-end debt-to-book ratio), S/A (the lagged fiscal year-end sales-to-assets ratio), $\sigma(eps)$ (the one-year lagged standard deviation of earnings per share divided by the absolute value of mean earnings per share over the past five years), and Age (the mean age of the forecasts). A firm's adjusted volume is its volume divided by the mean volume of the firm's exchange, while a firm's adjusted turnover is its turnover divided by the mean turnover of the firm's exchange. Newey–West t -statistics are in parentheses. Stocks with a price below five dollars are excluded.

β	$\ln M$	$\ln B/M$	$R_{-12:-2}$	$RCov$	$\sigma(R)$	Vol	$Turn$	$\ln D/B$	$\ln S/A$	$\sigma(eps)$	Age
0.131 (4.49)											
	-0.046 (-6.59)										
		0.075 (3.26)									
			-0.309 (-5.22)								
				0.135 (4.86)							
					0.141 (4.95)						
						0.005 (2.99)					
							0.044 (8.20)				
								0.039 (3.00)			
									-0.034 (-4.25)		
										0.005 (2.71)	
											-0.023 (-1.37)
0.112 (3.81)	-0.025 (-7.19)	0.058 (2.66)	-0.282 (-5.40)	0.114 (5.67)							
-0.030 (-1.41)	0.005 (0.81)	0.089 (3.29)	-0.221 (-6.04)	0.087 (5.07)	0.185 (3.66)						
0.080 (2.51)	-0.028 (-6.98)	0.059 (2.53)	-0.295 (-5.78)	0.089 (4.33)		0.002 (1.79)	0.039 (3.42)				
0.067 (1.90)	-0.030 (-7.85)	0.032 (1.55)	-0.288 (-6.23)	0.088 (3.86)		0.002 (1.24)	0.046 (3.23)	0.039 (2.23)	-0.031 (-3.30)	0.004 (2.89)	-0.020 (-0.80)




Figure 1. Robustness checks: Lags in portfolio formation. At the end of each month, all stocks in the I/B/E/S universe are ranked into deciles based on dispersion in analysts' earnings per share forecasts (defined as the ratio of the standard deviation to the absolute value of the mean forecast). Stocks in each dispersion decile are assigned into equal-weighted portfolios with a certain lag and held in the portfolio for one month. The difference in returns between the lowest- and highest-decile portfolios is calculated. The chart plots the mean monthly return differential, with the broken lines indicating the 95 percent confidence interval (adjusted for autocorrelation).

A. Lag in Portfolio Formation

We try waiting several months before assigning stocks to portfolios. The rationale for this is that even though I/B/E/S updates earnings forecasts frequently, the updates may not become available to the general public until much later. As shown in Figure 1, when the lag gets longer, returns clearly drop, as information becomes stale—prices on stocks in the portfolio may have already fallen due to selling by disappointed optimists. If the lag is longer than five months, the return difference is no longer statistically significant, even for the lowest market-capitalization quintile.

B. Different Holding Periods

We also try holding a stock in the portfolio for longer than one month. At the end of each month, stocks are ranked into deciles based on the dispersion in analyst earnings per share forecasts and assigned into portfolios without a lag. The stocks are then held in the portfolio for T months, with $1/T$ th of each portfolio reinvested monthly. Portfolio returns are equal weighted.

Figure 2 shows the results of the strategy for different holding periods. For relatively short holding periods, the results are similar to that of the one-month holding strategy. This is because essentially the same stocks are selected by both strategies every month due to persistence in analyst fore-




Figure 2. Robustness checks: Holding periods. At the end of each month, stocks are ranked into deciles based on dispersion in analysts' earnings per share forecasts and assigned into portfolios without a lag. The stocks are then held in the portfolio for T months, with $1/T$ th of each portfolio reinvested monthly. Portfolio returns are equal-weighted. The difference in returns between the lowest- and highest-decile portfolios is calculated. The chart plots the mean monthly return differential, with the broken lines indicating the 95 percent confidence interval (adjusted for autocorrelation).

cast dispersions. However, for longer holding periods, returns of this strategy are lower because stocks remaining in the portfolio too long may no longer satisfy the portfolio selection criteria.

C. Sorting Based on Dispersion in Two-Years-Ahead Earnings Forecasts

Here we sort stocks based on the dispersion in the following fiscal year's earnings forecasts (not reported here). There are fewer observations of the two-years-ahead earnings forecasts in I/B/E/S, and, therefore, slightly fewer stocks in each portfolio. Still, because the current and the next fiscal year's earnings-per-share forecasts are highly correlated, the forecast dispersions are also highly correlated, and the results are similar, although the return differential is somewhat lower.

D. Subperiod Analysis

We calculate mean monthly returns for the portfolio for the subperiods 1983 to 1991 and 1992 to 2000. As reported in Table VIII, the return differential is significant for all size quintiles for the period of 1983 to 1991. For the 1992 to 2000 time period, it is significant only for the smallest size quintile. As can be seen from the second part of the table, average dispersions of stocks in portfolios have also decreased in the later time period. This indicates that the overall magnitude of the disagreement has declined, possibly because investors have become better at valuing stocks. At the same

Table VIII
Subperiod Analysis

Each month stocks are sorted in five groups based on the level of market capitalization as of the third Thursday of the previous month. Each size group is then sorted into five subgroups based on dispersion in analyst forecasts. Dispersion is defined as the ratio of the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts to the absolute value of the mean forecast, as reported in the I/B/E/S Summary History file. Stocks with a mean forecast of zero are assigned to the highest dispersion groups, and stocks with a price less than five dollars are excluded from the sample. Stocks are held for one month, and portfolio returns are equal-weighted. The table reports the average monthly return differential between stocks in the lowest- and highest-dispersion quintile in each size group over the indicated time periods; *t*-statistics in parentheses are adjusted for autocorrelation.

Mean Returns						
Time Period	Small		Large			All Stocks
	S1	S2	S3	S4	S5	
1983–2000	1.37 ^a (5.98)	0.89 ^a (3.12)	0.67 ^b (2.41)	0.48 (1.55)	0.29 (0.93)	0.79 ^a (2.88)
1983–1991	1.54 ^a (6.11)	1.41 ^a (5.28)	0.94 ^a (3.56)	0.86 ^b (2.62)	0.66 (1.95)	1.16 ^a (4.63)
1992–2000	1.21 ^a (3.14)	0.37 (0.77)	0.40 (0.83)	0.11 (0.21)	−0.08 (−0.16)	0.41 (0.86)
Mean Dispersion						
Dispersion Quintiles	Size Quintiles					All Stocks
	S1	S2	S3	S4	S5	
Time period: 1983–1991						
D1 (low)	0.013	0.014	0.016	0.018	0.017	0.015
D5 (high)	1.434	1.148	0.986	0.938	0.487	1.008
Time period: 1992–2000						
D1 (low)	0.008	0.008	0.009	0.010	0.010	0.009
D5 (high)	1.081	0.780	0.642	0.509	0.437	0.698

^{a,b} Statistically significant at the one and five percent levels, respectively.

time, trading costs have come down, and firm-related information has become more readily available. All these factors may have contributed to the decline in the return differential over time.

E. Evaluation of Other Explanations

Hong et al. (2000) suggest that negative momentum in stock returns—the empirical finding that past losers continue to earn below normal returns in the future—is caused by the slow price adjustment to negative information. They empirically document that negative momentum is more pronounced for

stocks that have lower analyst coverage, after controlling for other factors which influence the amount of analyst coverage, such as size, covariance with the market, and the book-to-market ratio.

Though we have shown that momentum cannot explain away low returns of high-dispersion portfolios, we nonetheless check whether analyst coverage is a factor. As seen from Table VII, residual analyst coverage is positively correlated with dispersion in analysts' forecasts, indicating that there is a higher demand for expert opinions when existing information is difficult to interpret. This implies that the results described here cannot be explained by the slow information diffusion theory. Since high-dispersion stocks tend to have higher analyst coverage compared to other stocks, we expect that after controlling for effect of analyst coverage, our results will be even stronger.

Another concern is that the relative recency of the forecasts may be influencing the relation between forecast dispersion and future returns. In particular, since our tests form portfolios (at each date) that combine forecasts made immediately before the current date with those made up to a year before the current date, dispersion may be caused by the more recent forecasts that contain new information being combined with the old forecasts that contain no such information. In results not reported here, we calculate the mean age of the forecasts that we use in Tables II to V, and find that high-dispersion portfolios on average do indeed contain slightly more recent forecasts than low-dispersion portfolios. However, when we include the mean age of the forecasts in our regressions to explain dispersion (see Table VII), it is insignificant. Moreover, as mentioned earlier, our main argument as to why difference of opinions arise and influence prices does not depend on whether investors are asymmetrically informed or simply differ in the way they interpret the same information.

Our final robustness check addresses the possibility that by scaling the standard deviation of analysts' forecasts by the mean forecast, our results may simply be due to a denominator effect whereby the mean forecast is strongly positively related to future returns and thus drives our results. To address this possibility, we run Fama and MacBeth (1973) cross-sectional regressions each month on all securities in the intersection of CRSP, COMPUSTAT, and I/B/E/S from January 1980 to December 2000. The setup is analogous to that in Grinblatt and Moskowitz (1999), and provides a robustness check on our prior results by allowing us to avoid specifying breakpoints and to avoid weighting stocks by size. Specifically, we regress the cross section of individual stock returns at time t on a constant, market β ,¹⁰ size (log of market capitalization at $t - 1$), $\ln(BE/ME)$ (as defined in Table IV), a stock's past-year return ($ret_{-12:-2}$), the one-month past return ($ret_{-1:-1}$), a stock's long-run past return ($ret_{-36:-13}$), *residual coverage*, turnover, the one-month lag of *dispersion* in analysts' forecasts, the one-month lag of the standard deviation in analysts' forecasts, and the one-month lag of the mean analyst forecasts. *Residual coverage* is the residual from monthly regres-

¹⁰ We assign postranking β s according to the procedure in Fama and French (1992).

sions of $\ln(1 + \textit{analyst coverage})$ on $\ln(ME)$ and $\ln(BE/ME)$, and is included to address the concerns raised above that coverage may be driving our results. Turnover is the adjusted average turnover from the last 250 days (lagged one month); a firm's adjusted turnover is turnover divided by the mean turnover of the firm's exchange. The one-month past return variable is included to control for liquidity and microstructure effects documented by Jegadeesh (1990) that cause a short-term reversal in individual stock returns. The long-run past return variable is simply the return from $t - 36$ to $t - 13$, and is included to capture the three- to five-year reversal effect documented by DeBondt and Thaler (1985).

Since one can argue that turnover is another measure of differences of opinion (see, e.g., Harris and Raviv (1993) and Lee and Swaminathan (2000)), and since turnover has been shown to predict returns (Lee and Swaminathan), we include this variable in our regressions to test if turnover subsumes dispersion. Turnover is also the single most important explanatory variable for our measure of dispersion (see Table VII), making such a test an important robustness check on our prior results.

Table IX shows the results of the Fama and MacBeth (1973) cross-sectional regressions, including specifications with and without *residual coverage*, turnover, and the past-return variables. The coefficient on the dispersion variable is strongly negative in all of the specifications, in line with the patterns observed in Tables II to V. This result is robust to changes in the date at which dispersion is measured. Including size, book-to-market, and various momentum variables on the right-hand side of the regression controls for many of the previously documented anomalies in the literature, but does not capture the effect of dispersion in analysts' forecasts. Similarly, the inclusion of turnover does not drive out the dispersion effect; in fact, dispersion remains significant, while turnover is not.¹¹ The regression results confirm previous findings in this paper: There is a strong negative effect of dispersion in analysts' forecasts on future returns. In addition, Table IX demonstrates that coverage and scaling effects are not driving these results, since the standard deviation of analyst forecasts is negative and significant in all of the specifications, while the mean forecast is not.

V. Potential Explanations

As mentioned in Section III, we feel that the correct interpretation of dispersion in analysts' forecasts is as a proxy for differences of opinion among investors. Our results clearly reject the notion that dispersion can be viewed as a proxy for risk, since the relation between dispersion and future returns

¹¹ When we use unadjusted turnover in place of our measure of turnover, dispersion is still strongly negatively significant, but turnover is negative and significant as well. When we extend the horizon of our dependent variable to 12 months instead of 1, as in Lee and Swaminathan (2000), we find that both dispersion and unadjusted turnover are strongly negatively significant.

Table IX

Fama-McBeth Regressions: Explaining the Cross Section of Individual Stock Returns

Fama and Macbeth (1973) cross-sectional regressions are run every month from February 1980 to December 2000. The cross section of expected stock returns at time t is regressed on a constant (not reported), market β (estimated as in Fama and French (1992)), size (log of market capitalization at $t - 1$), book-to-market, residual coverage, turnover, the one-month lag of dispersion in analysts' forecasts, the one-month lag of the standard deviation in analysts' forecasts, and the one-month lag of the mean analyst forecast. The (log) book-to-market ratio ($\ln(BE/ME)$) is computed by matching the yearly BE figure for all fiscal years ending in calendar year $t - 1$ to returns starting in July of year t ; this figure is then divided by market capitalization (ME) at month $t - 1$ to form the book-to-market ratio, so that the book-to-market ratio is updated each month. Residual coverage (*Resid. Cov.*) is the residual from monthly regressions of $\ln(1 + \text{analyst coverage})$ on $\ln(ME)$ and $\ln(BE/ME)$. Turnover (*Turn*) is the adjusted average turnover from the last 250 days (lagged one month); a firm's adjusted turnover is its turnover divided by the mean turnover of the firm's exchange. Dispersion (*Disp*) is the ratio of the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts to the absolute value of the mean forecast (observations with a mean forecast of zero are deleted from the sample). Standard deviation (*Std. Dev.*) is the standard deviation of analysts' current-fiscal-year annual earnings per share forecasts, and *Mean* is the mean analyst earnings forecast. The return on each stock from $t - 12$ to $t - 2$ is designated $ret_{-12:-2}$. In addition, $ret_{-1:-1}$ and $ret_{-36:-13}$ are included to capture the short-term and long-run reversal effects in individual stock returns. Stocks with a price less than five dollars are excluded from the sample. t -statistics in parentheses are adjusted for autocorrelation.

β	$\ln(ME)$	$\ln(BE/ME)$	$ret_{-12:-2}$	$ret_{-1:-1}$	$ret_{-36:-13}$	Resid. Cov.	Turn	Disp	Std. Dev.	Mean
-0.156 (-0.40)	-0.039 (-0.81)	0.065 (0.80)	0.013 (5.62)	-0.037 (-6.25)	-0.002 (-1.96)					
-0.164 (-0.43)	-0.040 (-0.83)	0.069 (0.85)	0.013 (5.59)	-0.037 (-6.27)	-0.002 (-1.93)	0.008 (1.17)				
-0.061 (-0.18)	-0.033 (-0.69)	0.094 (1.10)	0.013 (5.84)	-0.037 (-6.36)	-0.002 (-1.51)	0.012 (1.70)	-0.096 (-1.41)			
-0.049 (-0.15)	-0.036 (-0.74)	0.094 (1.09)	0.013 (5.67)	-0.037 (-6.36)	-0.002 (-1.64)	0.012 (1.71)	-0.096 (-1.43)	-0.094 (-2.31)		
-0.201 (-0.46)	-0.031 (-0.63)	0.108 (1.22)				0.014 (1.74)		-0.121 (-2.70)		
-0.121 (-0.31)	-0.025 (-0.50)	0.116 (1.30)					-0.069 (-0.96)	-0.118 (-2.71)		
-0.192 (-0.43)	-0.029 (-0.61)	0.107 (1.20)						-0.119 (-2.62)		
-0.001 (-0.00)	-0.045 (-0.95)	0.109 (1.29)	0.012 (5.39)	-0.038 (-6.66)	-0.002 (-2.04)	0.011 (1.53)	-0.087 (-1.33)		-0.447 (-2.07)	0.042 (1.19)
-0.013 (-0.04)	-0.039 (-0.82)	0.153 (1.74)				0.016 (2.11)	-0.075 (-1.07)		-0.649 (-2.93)	0.060 (1.49)
-0.081 (-0.20)	-0.041 (-0.86)	0.147 (1.67)							-0.662 (-2.76)	0.056 (1.31)

is strongly negative. Returning to the hypotheses laid out in Section I, we feel that our results provide support for the predictions of the price-optimism model of Miller (1977). In this section, we explore this theory in greater depth.

First of all, some of the evidence presented in Sections II to IV can also be used to evaluate the Miller (1977) model. For example, the finding (in Table II) that the return differential between low- and high-dispersion stocks is strongest in small stocks is consistent with Miller's model, since these stocks are likely to be the ones that are hardest to short and least likely to have traded options. There is also an asymmetry in the return differential that is consistent with the Miller story; as noted earlier, in Table II we see that 68.4 percent of the overall spread comes from the short side of the trade. This is to be expected, since the Miller story stresses the fact that high-dispersion stocks will underperform; the converse, that low-dispersion stocks will outperform, need not be true. Indeed, the estimated intercepts (the "alphas") in Table VI support this claim. The three- and four-factor alphas are strongly negative for the high-dispersion quintiles; meanwhile, the alphas are positive and insignificant for low-dispersion stocks in the three-factor model, and positive and marginally significant for low-dispersion stocks in the four-factor model.

One of the implications of Miller's (1977) hypothesis is that convergence of prices down to fundamental values should coincide with the resolution of uncertainty. When uncertainty about the stock value is resolved, prices come down to reflect the unbiased view, leading to low returns on the high-dispersion stocks. Evidence in Scherbina (2001) supports this idea. She shows that 11 to 33 percent of the return differential between the low- and high-dispersion stocks falls on a three-day window around the earnings announcement dates for smaller stocks (for smaller stocks, more information gets revealed on the earnings announcement dates, while for larger stocks, information is available more steadily throughout the year).¹² She also shows that higher dispersion in analysts' forecasts leads to lower cumulative abnormal returns in the three-day window around the earnings announcements dates.

Our subperiod analysis in Table VIII also supports the Miller (1977) theory. In it, we see that the average return differential between the low- and high-dispersion stocks has declined in the later part of the sample, becoming insignificant for all but the smallest size quintile. This is to be expected in a Miller world for the following reasons. First of all, short-sale costs have come down over time, resulting in less binding short-sale constraints. At the same time, firm-related information processing and availability have improved, perhaps explaining the lower levels of disagreement we find today.

Additionally, Table IV indicates that future returns on growth stocks are more sensitive to the differences of opinion about the firms' expected earnings, consistent with Miller's (1977) theory. This is because a given degree of

¹² See Chari, Jagannathan, and Ofer (1988).

dispersion in an earnings forecast translates into a higher amount of disagreement about the underlying value for a growth stock than a value stock. For growth firms, equity value is based more on the business strategy and the hard-to-quantify value of research and development than on the book value of assets in place. Therefore, uncertainty about current earnings projected forward gives a further magnified uncertainty about the value of the growth stock. This, in turn, gives rise to the increased magnitude of the divergence in opinion.

Given that high short-sale costs are the main impediment to the revelation of negative information, Miller's (1977) effect should be more pronounced for stocks that are more difficult to short-sell. However, evidence relating returns to short-sale costs is not as strong. Scherbina (2001) argues that high-dispersion stocks in the S&P 500 index should not earn lower returns, because these stocks should be very cheap to short-sell. However, she still finds that stocks in the lowest dispersion quintile of the S&P 500 index outperform stocks in the highest dispersion quintile by 0.48 percent per month (a significant difference at the 10 percent level). Furthermore, Scherbina investigates whether various proxies for short-sale costs (such as float and turnover) predict future returns, and finds only a weak relation between proxies for short-sale costs and future returns. However, this result is likely due to the noisiness of the proxies used. Indeed, evidence of a link between short-sale costs and future stocks returns is mounting. For example, using more direct proxies for short-sale costs, D'Avolio (2001), Jones and Lamont (2001), and Lamont and Thaler (2001) all present evidence that high short-sale costs can lead to lower future returns, consistent with Miller's model of price optimism.

As mentioned earlier, given boundedly rational investors and limited arbitrage, any friction that prevents the revelation of negative opinions will produce the basic result that we find in the data. Analysts, who comprise an important channel of information to the market, are often reluctant to make low earnings-per-share forecasts or issue "sell" recommendations. McNichols and O'Brien (1977) argue that pessimistic analysts prefer to stop covering the stock. This self-censoring is due to analysts' incentive structure. Analysts rely on firm management for tips about future earnings, and risk being cut off when they voice pessimism about a firm's prospects. Additionally, analysts earn a percentage on commissions from stock sales and get rewarded whenever their employer wins lucrative investment banking deals, which furthers their incentives to follow stocks for which they have an optimistic outlook.

The self-censoring of pessimistic analysts creates an upward bias in consensus earnings-per-share forecasts. This bias is higher the larger the initial magnitude of disagreement about the firm's earnings. The following argument shows the causality involved. It is natural to assume that the more spread out the earnings estimates, the lower will be the cut-off point below which analysts will not issue forecasts. Suppose that the true underlying distribution of analyst earnings is normal with mean μ and standard deviation σ , and the cut-off point is $\mu - k\sigma$, where k is a constant identical for all analysts and firms. The observed distribution of the forecasts will be a

truncated normal distribution, with a mean higher than the true mean by $\sigma(\phi(k)/\Phi(k))$ where $\phi(\cdot)$ and $\Phi(\cdot)$ are the P.D.F. and the C.D.F. of a standard normal distribution, respectively.¹³ This shows that the bias in the mean reported forecast increases with the dispersion of the underlying distribution. If stocks are priced correctly with respect to the mean reported earnings forecast, they will be more overvalued the higher the amount of disagreement about the firm's earnings. This argument is similar in spirit to Miller's (1977) argument, except that the friction that prevents revelation of negative opinions is the incentive structure of analysts. Again, if investors who use earnings forecasts are fully rational, they should be able to infer the true mean of the distribution by taking into account the self-selection in analyst coverage.

We find strong support in the data for the conjecture that the upward bias in consensus earnings forecasts is positively related to forecast dispersion. Forecast error in a quarterly earnings-per-share forecast is defined as the realized minus the forecasted value. Historically, analysts have been too optimistic, so the errors are, on average, negative. We run a pooled time-series and cross-sectional regression of the quarterly forecast errors on standard deviation in quarterly earnings per share forecasts:¹⁴

$$FE_{it} = \alpha_0 + \alpha_1 disp_{it} + \epsilon_{it}. \quad (1)$$

The regression coefficients (with the t -statistics in parenthesis) are: $\alpha_0 = -0.019$ (-4.68) and $\alpha_1 = -0.727$ (-33.42). The adjusted R^2 for the regression is 4.5 percent. The negative intercept indicates that, on average, analysts have been too optimistic in their earnings forecasts. The highly significant negative t -statistic on forecast dispersion indicates that there is indeed a strong positive relationship between optimism in consensus forecasts and uncertainty surrounding the stock's earnings per share.

VI. Summary and Conclusions

In this paper we provide evidence that stocks with higher dispersion in analysts' earnings forecasts earn significantly lower future returns than otherwise similar stocks. Our results clearly reject the notion that dispersion in forecasts can be viewed as a proxy for risk, since the relation between dispersion and future returns is negative. Further, we present evidence that standard multifactor risk-based explanations cannot account for this relation. We feel that the correct interpretation of dispersion in analysts' forecasts is as a proxy for differences in opinion about a stock. We show that our results are consistent with the hypothesis that prices will reflect the optimistic view whenever investors with the lowest valuations do not trade.

¹³ Proof is available upon request.

¹⁴ Means and standard deviations in quarterly and annual earnings forecasts are calculated from the Detail History data set for the date immediately preceding the earnings announcement date.

Further, although Miller's (1977) theory focuses on short-sale costs as the reason that investors with the lowest valuations do not trade, we argue that any friction that prevents the revelation of negative opinions will produce the negative relation between dispersion and future returns that we document in this paper. To this end, we demonstrate that the incentive structure of analysts could serve as another such friction. Due to this incentive structure, analysts prefer not to issue forecasts when they are sufficiently pessimistic. This creates an upward bias in consensus forecasts. The bias increases with dispersion in the underlying forecasts. It would be interesting to isolate the importance of this effect on upward price bias. Recently, investors have started to rely more heavily on unofficial and largely anonymous "whisper" forecasts, which are available from several Internet sites.¹⁵ These forecasts should not be censored in the same way as analyst forecasts because people making them do not have similar incentives. By comparing the pre- and post-whisper samples, it would be possible to identify the relationship between forecast censoring and price bias. Perhaps, the increased reliance on non-analyst forecasts could also help to explain why the return differential between the low- and high-dispersion stocks has declined in later years.

REFERENCES

- Ackert, Lucy F., and George Athanassakos, 1997, Prior uncertainty, analyst bias, and subsequent abnormal returns, *The Journal of Financial Research* 20, 263–273.
- Bagnoli, Mark, Messod D. Beneish, and Susan G. Watts, 1999, Whisper forecasts of quarterly earnings per share, *Journal of Accounting and Economics* 28, 75–96.
- Carhart, Mark M., 1997, On persistence in mutual fund performance, *The Journal of Finance* 52, 57–82.
- Chari, V.V., Ravi Jagannathan, and Aharon R. Ofer, 1988, Seasonality in security returns: The case of earnings announcements, *Journal of Financial Economics* 21, 101–121.
- Chen, Joseph, Harrison Hong, and Jeremy C. Stein, 2001, Breadth of ownership and stock returns, *Journal of Financial Economics*, forthcoming.
- Cragg, John, and Burton Malkiel, 1982, *Expectations and the Structure of Share Prices* (University of Chicago Press, Chicago, IL).
- Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam, 2001, Overconfidence, arbitrage, and equilibrium asset pricing, *The Journal of Finance* 56, 921–966.
- D'Avolio, Gene, 2001, The market for borrowing stock, Working paper, Harvard University.
- DeBondt, Werner F.M., and Richard H. Thaler, 1985, Does the stock market overreact? *The Journal of Finance* 40, 793–808.
- Diamond, Douglas W., and Robert E. Verrecchia, 1987, Constraints on short-selling and asset price adjustment to private information, *Journal of Financial Economics* 18, 277–311.
- Fama, Eugene F., and Kenneth R. French, 1992, The cross section of expected stock returns, *The Journal of Finance* 46, 427–466.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns of stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *The Journal of Finance* 51, 55–84.
- Fama, Eugene F., and James MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

¹⁵ Bagnoli, Beneish, and Watts (1999) document that stock prices reflect "whisper" forecasts rather than analyst forecasts.

- Gebhardt, William R., Charles M.C. Lee, and Bhaskaran Swaminathan, 2001, Toward an implied cost of capital, *Journal of Accounting Research* 39, 135–176.
- Gervais, Simon, J.B. Heaton, and Terrance Odean, 2000, Capital budgeting in the presence of managerial overconfidence and optimism, Working paper, University of Pennsylvania.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of efficiency of a given portfolio, *Econometrica* 57, 1121–1152.
- Goetzmann, William N., and Massimo Massa, 2001, Heterogeneity of trade and stock returns: Evidence from index fund investors, ICF Working Paper No. 00-28, Yale University.
- Grinblatt, Mark, and Tobias Moskowitz, 1999, Do industries explain momentum? *The Journal of Finance* 54, 1249–1290.
- Gromb, Denis, and Dimitri Vayanos, 2001, Equilibrium and welfare in markets with financially constrained arbitrageurs, CEPR Discussion Paper No. 3049, Centre for Economic Policy Research.
- Harris, Milton, and Artur Raviv, 1993, Differences of opinion make a horse race, *Review of Financial Studies* 6, 475–506.
- Hong, Harrison, Terence Lim, and Jeremy C. Stein, 2000, Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies, *The Journal of Finance* 55, 265–295.
- Hong, Harrison, and Jeremy C. Stein, 2000, Differences of opinion, rational arbitrage and market crashes, Working paper, Stanford University.
- Jarrow, Robert, 1980, Heterogeneous expectations, restrictions on short sales and equilibrium asset prices, *The Journal of Finance* 35, 1105–1114.
- Jegadeesh, Narasimhan, 1990, Evidence of predictable behavior of security returns, *The Journal of Finance* 45, 881–898.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65–91.
- Jegadeesh, Narasimhan, and Sheridan Titman, 2001, Profitability of momentum strategies: An evaluation of alternative explanations, *The Journal of Finance* 56, 699–720.
- Jones, Charles M., and Owen A. Lamont, 2001, Short sale constraints and stock returns, *Journal of Financial Economics*, forthcoming.
- Lamont, Owen A., and Richard H. Thaler, 2001, Can the market add and subtract? Mispricing in tech stock carve-outs, Working paper, University of Chicago.
- LaPorta, Rafael, 1996, Expectations and the cross section of stock returns, *The Journal of Finance* 51, 1715–1742.
- Lee, Charles M.C., and Bhaskaran Swaminathan, 2000, Price momentum and trading volume, *The Journal of Finance* 55, 2017–2069.
- Mayshar, Joram, 1982, On divergence of opinion and imperfections in capital markets, *American Economic Review* 73, 114–128.
- McNichols, Maureen F., and Patricia O'Brien, 1997, Self-selection and analyst coverage, *Journal of Accounting Research* 35, 167–199.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- Merton, Robert C., 1987, A simple model of capital market equilibrium with incomplete information, *The Journal of Finance* 42, 483–510.
- Miller, Edward M., 1977, Risk, uncertainty, and divergence of opinion, *The Journal of Finance* 32, 1151–1168.
- Morris, Stephen, 1996, Speculative investor behavior and learning, *Journal of Financial Economics* 111, 1111–1133.
- Scherbina, Anna, 2001, Stock prices and differences of opinion: Empirical evidence that prices reflect optimism, Working paper, Northwestern University.
- Shleifer, Andrei, and Robert Vishny, 1997, The limits of arbitrage, *The Journal of Finance* 52, 35–55.
- Viswanathan, S., 2000, Strategic trading, heterogeneous beliefs and short constraints, Working paper, Duke University.
- Williams, Joseph T., 1977, Capital asset prices with heterogeneous opinion, *Journal of Financial Economics* 5, 219–241.