

Preliminary draft: please do not cite without permission. Comments welcome.

Analysts' Forecast Bias and the Mispricing of High Credit Risk Stocks

Mark Grinblatt¹, Gergana Jostova², Alexander Philipov³

This draft: July 30, 2014

ABSTRACT

This paper investigates whether financial analysts' power to move prices arises from investors' tendency to blindly follow analyst earnings estimates. Analyst forecasts are often overly optimistic. This optimism is predictable and may generate temporarily inflated stock prices. In addition, for high credit risk stocks, the quintile predicted to have the most optimistic forecasts outperforms the quintile with the least optimistic forecasts by about 19% per year. Certain types of firms attract significantly more analyst optimism than others—namely those with poor credit quality. For these firms, the price distortions caused by analyst optimism are so large and frequent that they account for the negative credit risk-return relation observed in the cross section of U.S. stocks.

¹UCLA Anderson School of Management, email: mark.grinblatt@anderson.ucla.edu.

²Department of Finance, School of Business, George Washington University, email: jostova@gwu.edu.

³Department of Finance, School of Management, George Mason University, email: aphilipo@gmu.edu.

Keen observers of the stock market recognize that analyst recommendations and forecasts move stock prices. What is unclear is whether such price movements arise because of the rational dissemination of the analyst’s superior information or because of “animal spirits”—the tendency of investors to slavishly follow experts’ opinions, irrespective of their merit. Distinguishing which of these two views of analyst opinions is correct advances our understanding of what determines stock prices. Researchers who subscribe to the view that markets are efficient believe that analysts move stock prices because market participants recognize that analysts have access to superior information or analysis, leading to a more accurate portrait of the fair pricing of a stock. The alternative view—that analysts’ power to move prices arises from the blind tendency to follow—represents a behavioral perspective.

If markets are efficient, the price movements generated by the news in analyst recommendations and forecasts should be permanent. That is, subsequent price movements will be random. Moreover, to the extent that the analyst’s view is predictable, and therefore not news, it should not move stock prices. If, alternatively, the behavioral perspective is correct, then both predictable and unpredictable analyst opinions may change stock prices. However, we would know that these analyst-driven price movements emerged from animal spirits and not news if they are temporary—that is, they subsequently tend to reverse. When these temporary price movements stem from a predictable analyst opinion, which cannot be news, the behavioral perspective on markets becomes particularly salient.

This paper finds strong evidence for the behavioral perspective: Firms with more optimistic consensus analyst forecasts subsequently earn lower risk-adjusted returns. The most likely explanation for this “analyst-bias anomaly” is that optimism temporarily inflates prices and inflated prices deflate as more accurate earnings and revenue information is disseminated to the market. This obvious behavioral explanation distinguishes the analyst-bias anomaly from other major anomalies, which do not as readily reveal their root causes. Supporting this behavioral view of the analyst-bias anomaly is that the lower returns of firms with up-

wardly biased forecasts are witnessed for a predictable component of the bias. A predictable component is not news, but it moves prices nonetheless.

Consensus analysts earnings forecasts in general tend to be overly optimistic, particularly long before earnings are announced.¹ Certain types of firms—e.g., those that have a history of optimism, those for which analyst opinions diverge, those with lower past returns, and those with negative past earnings surprises—tend to have more analyst optimism than others. These types of firms appear to be the most overpriced at the time of the earnings forecast; their stocks subsequently perform poorly compared to firms with less optimism.

Prior research links analyst earnings forecasts to future stock price changes, but seems inconclusive about whether these price changes are permanent responses to superior fundamental valuation or temporary consequences of analyst hunches. [Womack \(1996\)](#) identifies a positive link between analyst opinions and future stock returns. He shows that investors who mimic analysts' recommendations achieve superior performance. Womack's evidence portrays analysts as savants who identify when a stock's price differs from a fair value that the price ultimately gravitates towards. By contrast, [LaPorta \(1996\)](#) shows that stocks with the highest earnings growth forecasts tend to earn the lowest returns. At first glance, this is a view of analysts as "Pied Pipers." Investors who march to their forecast tunes are doomed to underperform the market. However, the link LaPorta finds between returns and prior earnings growth forecasts is strongly tied to the low returns of high growth stocks documented by [Fama and French \(1992\)](#). [Ackert and Athanassakos \(1997\)](#) and [Diether, Malloy, and Scherbina \(2002\)](#) show that analyst forecast dispersion predicts lower future returns. The theory here is that the market reacts only to the more optimistic half of the forecasts because of short-sale constraints, which inflates a stock's price the more dispersion there is.

¹Analyst bias towards optimism was first documented in [De Bondt and Thaler \(1990\)](#) and [Ali, Klein, and Rosenfeld \(1992\)](#). Since then, researchers have advanced several explanations for this analyst optimism. These include overconfidence ([Hilary and Menzly, 2006](#)), under/over-reaction to bad/good news ([Easterwood and Nutt, 1999](#)), underwriter affiliation and investment banking relationships ([Lin and McNichols, 1998](#)), bolstering trading commissions ([Hayes, 1998](#)), and access to firm managers ([Francis and Philbrick, 1993](#)).

As inflated share prices ultimately deflate, stocks with more dispersion earn lower future returns. However, [Avramov, Chordia, Jostova, and Philipov \(2009b\)](#) argue that firms with the greatest forecast dispersion tend to be firms with high credit risk, a prognosticator of subsequent poor performance. Once credit risk is controlled for, dispersion plays no role in predicting returns.

Our research design differs from others in its direct focus on the rationality of the consensus forecast. We study whether false consensus analyst optimism artificially inflates a stock's price, controlling for forecast dispersion, credit risk, momentum, value, firm size, and other factors. The trading strategies studied in our paper are tied to cross-sectional differences in the optimism of analyst earnings forecasts. For example, analysts tend to be most overly optimistic about the earnings of high credit risk firms. The average consensus forecast is 42% higher than actual earnings for firms in the worst (prior month) credit rating quintile, but is less than 9% above actual earnings for firms in the best credit rating quintile.

As noted above, the extant literature shows that credit risk correlates with a firm's risk premium, but not in the direction one might expect. [Dichev \(1998\)](#), [Campbell, Hilscher, and Szilagyi \(2008\)](#), and [Avramov, Chordia, Jostova, and Philipov \(2009a\)](#) document that high credit risk stocks earn lower returns than low credit risk stocks. This "credit risk effect" is considered anomalous because high credit risk stocks also have higher systematic risk than other stocks. We find that the negative risk-adjusted returns of high credit risk stocks disappear once we control for analyst optimism. The low returns for many of the stocks in the high credit risk sector are consistent with artificially inflated prices converging to their fundamental values as the extreme analyst optimism about their earnings prospects wanes.

This paper quantifies the influence of analysts' false earnings optimism on stock prices using either portfolio sorts or cross-sectional regressions of risk-adjusted returns on various firm characteristics and the predicted degree of false optimism in the consensus analyst earnings forecast. Trading against this false optimism leads to highly profitable market-neutral

strategies that cannot be explained by risk or any of the well-known stock return anomalies documented by finance researchers. This analyst bias anomaly is large, particularly for companies with the poorest credit ratings. The share valuations of these firms plausibly possess the greatest sensitivity to information, increasing the value of a rumor or false hunch perceived as information. Indeed, among stocks in the highest credit risk quintile, those with predicted analyst bias in the highest quintile underperform those with bias in the lowest quintile by 163 basis points per month.

There are three contributions to the literature here. First, our research documents that overly optimistic analyst forecasts are more prevalent among firms with high credit risk. Second, it shows that cross-sectional differences in analyst optimism lead to an efficient markets anomaly: An investor can earn abnormal risk-adjusted profits by selling high credit risk firms with the most overly optimistic consensus forecasts and/or buying those with the least optimistic forecasts. Third, we conclude that greater analyst optimism for high credit risk stocks is the likely explanation for their lower returns and negative alphas. Specifically, average risk-adjusted returns sorted by credit rating, as well as cross-sectional regressions of risk-adjusted returns on credit ratings, show that returns significantly decrease as the credit rating deteriorates (and credit risk increases). Including analyst bias in the sorts and regressions eliminates this inverse relation between a stock's credit risk and its average return. In fact, once we control for analyst bias, there is a positive (albeit weak) relationship between credit risk and average returns.

Our paper is organized as follows. Section I discusses the data and methodology. It also presents summary statistics relating analyst bias and credit risk to a variety of firm attributes. Section II presents our results, focusing on the joint effect of credit risk and analyst bias on risk-adjusted returns. Both portfolio sorts and cross-sectional regressions measure the profitability of trading strategies based on analyst bias. The last part of Section II presents a model that can be used to interpret the observed results. Section III concludes the paper.

I. Data and Methodology

This section first describes the filters used to create the sample of firms we study. Then, it discusses the methodology for computing risk-adjusted returns and analyst optimism bias, followed by a discussion of the specifications used to relate these two variables. Finally, it presents summary statistics for the data, conditional on credit risk and two rankings of predicted analyst bias.

Data Filters. Our analysis starts with all NYSE, AMEX, or NASDAQ-listed common stocks on the CRSP Monthly Returns File that trade from 1986 to 2012.² In each month t of the 27-year sample period, we include each stock i in a trading strategy based on equal-weighted portfolio sorts or regression analysis (with coefficients representing portfolio returns). The trades employ a signal computed from a month $t - 1$ forecast of stock i 's analyst optimism bias and control variables believed to be related to the cross-section of expected returns. We exclude stocks that lack a) share prices at or above \$1 at the end of month $t - 1$, or b) a month t CRSP return³ or c) a month $t - 1$ Standard & Poor's (S&P) long-term domestic issuer credit rating.⁴ These requirements generate a sample of 318,781 firm-month observations with the number of firms each month ranging from 776 to 1,234.

Risk-Adjusted Returns. We use risk-adjusted returns throughout the paper. In particular, following Brennan, Chordia, and Subrahmanyam (1998), we compute stock i 's month t risk-adjusted return as the difference between its realized month t excess return and its month t

²October 1985 represents the first month that the credit ratings of firms reliably appear on WRDS.

³We adjust for delisting months using the standard treatment for delisting returns, i.e. compounding delisting returns with standard returns (see Beaver, McNichols, and Price, 2007).

⁴We employ S&P's long-term issuer credit rating of each firm as listed in Compustat for each month t , or in S&P's RatingsXpress when the credit rating is missing from Compustat. As defined by S&P, the "long-term issuer credit rating is a current opinion of an issuer's overall creditworthiness, apart from its ability to repay individual obligations. This opinion focuses on the obligor's capacity and willingness to meet its long-term financial commitments (those with maturities of more than one year) as they come due." When reporting average credit ratings for groups of firms, we convert the 22 S&P letter ratings into numerical scores as follows: 1=AAA, 2=AA+, . . . , 10=BBB-, 11=BB+, . . . , 19=CCC-, 20=CC, 21=C, 22=D. Hence, higher scores indicate higher credit risk. Mapping letter ratings into numbers and vice versa serves the purpose of averaging ratings. Ratings AAA to BBB- are considered investment grade (denoted IG) and ratings BB+ to D are considered non-investment grade or high-yield (denoted NIG).

predicted excess return from the four-factor model of Carhart (1997):

$$r_{i,t}^* = (r_{i,t} - r_{f,t}) - \hat{\beta}_{i,MKT}MKT_t - \hat{\beta}_{i,SMB}SMB_t - \hat{\beta}_{i,HML}HML_t - \hat{\beta}_{i,UMD}UMD_t \quad (1)$$

where $\hat{\beta}_{ik}$ is beta estimated from a time-series regression of the firm's excess stock return on the four factors over the entire sample period. The regression is separately run for every stock that has at least 24 months of non-missing return data.⁵

Computing Analyst Optimism Bias. A firm's earnings forecast bias at a point in time is defined to be the percentage difference between its consensus earnings forecast and an unbiased estimate of its earnings given the information at that time. Unfortunately, this forecast bias is not directly measurable because we don't know what the unbiased forecast is. To estimate it, we rely on rational expectations. The unbiased forecast can be viewed as the realized future earnings of the firm plus mean zero noise. Hence, we use the future realized earnings in place of the unmeasurable unbiased earnings forecast. Specifically, firm i 's month t *ex-post* analyst bias is computed as its (end-of) month t consensus annual earnings per share (EPS) forecast for what I/B/E/S refers to as the fiscal year FY1 (or current year) forecast⁶ minus the actual EPS realized at the fiscal year end, standardized by the absolute value of the actual EPS.⁷ Formally,

$$AB_{i,t} = \frac{ConForecastEPS_{i,t}^T - EPS_i^T}{|EPS_i^T|} \quad (2)$$

⁵While this entails the use of future data in calculating factor loadings, Fama and French (1992) show that an in-sample approach does not bias coefficients and tends to only have a negligible influence on them compared to out-of-sample estimation. See also Avramov and Chordia (2006).

⁶Firm i 's end-of-month consensus forecast is an average of the most recent earnings per share forecasts collected by I/B/E/S from analysts following firm i . The FY1 forecast thus refers to the earliest fiscal year earnings that have yet to be announced by month end t .

⁷All calculations that use the *ex-post* bias exclude firm-month observations with obvious reporting errors. Among these are cases where month T 's earnings announcement precedes the firm's fiscal period-end date or is reported to be exactly zero. To prevent small positive and negative values of actual EPS from unduly affecting our inferences, analyses that employ the *ex-post* bias as an input only include bias observations between the 1st and 99th percentiles for the overall sample.

where EPS_i^T is firm i 's actual EPS for the fiscal year end (ultimately announced in month T) and $ConForecastEPS_{i,t}^T$ is the month t analyst consensus forecast of that annual EPS, made prior to month T .

Naturally, this analyst bias changes every month as analysts update their forecasts for the same upcoming fiscal year. We refer to this measure of analyst bias as “*ex-post*” because actual 10K earnings have yet to be announced in forecast month t . Indeed, the bias is generally not known until one to three months after the fiscal year ends.

To properly assess whether this bias influences stock prices, we need a measure of the firm’s analyst forecast bias at the time of the earnings forecast, not at the later date when true earnings are announced. An *ex-post* bias measure that looks ahead at future earnings to compute a stock’s degree of analyst optimism could inversely correlate with future returns for reasons that have nothing to do with the analysts’ tendency towards greater or less optimism for a stock—but rather, because future returns are leading indicators of future realized earnings.

We obtain our *ex-ante* bias measure as a prediction of the *ex-post* bias measure from instruments known at least one-month prior to the consensus forecast. Specifically, the prediction comes from a panel regression of a firm’s (end-of) month- t bias on control variables known at the end of month $t - 1$ and the prior fiscal year’s analyst bias, with the prior bias measured with the same delay as month t from the last 10K earnings announcement. The controls in the instrumental variable regression include:

- *dispersion*, as measured by the prior-month standard deviation of analyst EPS forecasts, standardized by the absolute value of the prior month consensus analyst forecast, subject to at least two analysts covering the firm;
- *coverage*, as measured by the prior month’s number of analysts covering a firm;
- *two regressors, one for positive (and one for negative) momentum*, as measured by the

maximum (or minimum) of zero and the firm’s cumulative past 6-month return (which excludes the return in the prior month);

- *two regressors, one for positive (and one for negative) earnings surprise*, as measured by the maximum (or minimum) of zero and the most recent year-to-year change in quarterly earnings known at the end of month $t - 1$, scaled by the standard deviation of the 8 most recent earnings changes for the same quarter;
- *dummies for small firms and value firms* that take values of one if the firm is below the sample’s prior-month median for size or above the prior-month median for book-to-market, respectively;
- *rating dummies* indicating prior-month membership in one of 17 notched S&P credit rating groups;⁸
- *industry dummies* for 19 of the 20 industries in Moskowitz and Grinblatt (1999).

Formally, our prediction, \widehat{AB}_i , of firm i ’s analyst bias during month t (dropping t for notational simplicity) is given by:

$$\begin{aligned}
 \widehat{AB}_i &= c_0 + c_1 PastAB_i + c_2 Dispersion_i + c_3 Coverage_i \\
 &+ c_4 PastRet_i^- + c_5 PastRet_i^+ + c_6 SUE_i^- + c_7 SUE_i^+ \\
 &+ c_8 D_{Small,i} + c_9 D_{Value,i} + \mathbf{d} \mathbf{D}_{Rating,i} + \mathbf{e} \mathbf{D}_{Industry,i}
 \end{aligned} \tag{3}$$

where $PastAB_i$ is the *actual* analyst bias for the prior fiscal year following the 10K earnings announcement by the same number of months as AB_i , and the coefficients $c_0, \dots, c_9, \mathbf{d}, \mathbf{e}$, are estimates from a full sample panel regression of AB_i on the regressors in the equation

⁸The five omitted credit ratings, CCC+, CCC, CCC−, CC, and C, embedded in the constant, are grouped together because they contain relatively few observations.

above.⁹

To illustrate, consider the prediction of firm i 's May 1995 analyst bias. Assume that firm i 's May 1995 consensus current-year earnings forecast is for a fiscal year ending in December 1995. Also, assume that firm i always reports its annual earnings in February. Then, firm i 's $PastAB_i$ regressor would be fiscal 1994's actual analyst bias for May 1994, three months past fiscal 1993's earnings announcement month. In some cases, a firm may report annual earnings in January one year and in March the next, in which case the $PastAB_i$ regressor would be 14 months rather than 12 months prior to the month in which AB_i is measured.¹⁰ Relating *ex-post* biases at the same point in two consecutive annual earnings forecast cycles—which usually, but not always, is 12 months apart—accounts for the fact that analyst optimism bias predictably diminishes over the cycle, as documented by Richardson, Teoh, and Wysocki (2004). In part, this is because an optimistic forecast bias cannot be maintained for the portion of annual earnings that ceases to be forecastable—a portion fully disclosed by the release of quarterly earnings in corporate 10Q statements.

To address the annual cyclicity in optimism bias, panel regression (3) is run separately for four groups of firms sorted by the number of 10Q earnings statements (denoted $q = 0, 1, 2,$ or 3) released by the firm for the relevant fiscal year.¹¹ In addition to constructing \widehat{AB} and \widehat{AB} quintile sorts within each month, we assign firms monthly to *cycle-specific* \widehat{AB} quintiles by comparing the \widehat{AB} s of firms sharing the same q . All firms in the same quintile for their cycle-specific cohort are then grouped together. We also create a cycle-adjusted measure of predicted analyst forecast bias, denoted \widehat{AB}^{CA} . For each subgroup q , we look at the

⁹The panel regression has no firm fixed effects. Although we use the same full-sample coefficients for each firm-month, the panel's in-sample regression coefficients are about the same for the first and second half of the sample period.

¹⁰We exclude observations for which the $PastAB_i$ precedes AB_i by more than 15 months (thus indicating a delay in the announcement preceding month t of more than 3 months or a major change in the fiscal year). We have verified that our results are robust to changing the maximum difference to 20 months.

¹¹Since 10Q earnings announcement dates in Compustat are more sparse than 10K announcement dates, we assume that the announcement month for each of the three 10Q earnings reports occurs at three-month intervals after the 10K announcement month.

average predicted analyst bias across all firm-months in that subgroup. The cycle-adjusted firm-month observations, \widehat{AB}^{CA} , in cycle q , are obtained by multiplying the corresponding firm-month observations of \widehat{AB} by the ratio of the sample average predicted bias for cycle 0 to that for cycle q .

Table I reports equation (3)'s coefficients for regressions using subgroups of firms, sorted by forecast cycle q . It shows that past analyst bias, analyst forecast dispersion, being a small firm, and being a value firm are positively related to future analyst bias; earnings surprises (particularly negative ones), coverage, and past returns (particularly negative ones) are inversely related to future analyst bias. These three inverse relationships could be accounted for by the fact that analysts update their earnings forecasts infrequently whereas true expectations of earnings update almost continuously—past returns and earnings surprises are two sources of information that would rationally update an earnings forecast, while a larger pool of analysts is more likely to witness some analyst within the pool updating, thus changing the consensus. The greater competition within larger pools of analysts may also deter procrastination by analysts when new information warrants an updated earnings forecast. Credit risk is also related to analyst bias, influencing variables on both the left and right sides of equation (3). We will discuss the role of credit risk in great detail later in the paper.

As this is a panel regression, with enormous amounts of data, all of these instruments are highly significant. Clustering of standard errors would not alter this conclusion. The regressions have R-squareds that range from 7% to about 8.5%, depending on the forecast cycle quarter. R-squareds of this magnitude are impressive in that the regressions are attempting to outperform the consensus forecast of supposed experts, estimating the degree to which the consensus is wrong. The R-squareds are not affected by the general tendency towards analyst optimism, only by predictable differences in the optimism bias across firms in a given month.

Lastly, we note that the coefficients reported in Table I are fairly consistent in sign and

magnitude across the four forecast cycles. In light of the high degree of cyclicity in forecast bias, it appears that forecasting the *ex-post* bias with cycle-specific regressions is the appropriate way to develop the instrument for analyst bias. As evidence of the degree of cyclicity in analyst bias, Figure IV Plot A documents forecast bias cyclicity for our sample. It graphs average *ex-post* bias in 48 event-time months centered around month 0, which corresponds to the most recent 10K announcement. The three lines correspond to average *ex-post* bias for all firms with credit ratings, as well as for investment-grade and non-investment-grade firms. All three lines in the graph show a similar 12-month pattern—analyst bias is largest at the beginning of the forecast cycle, just when earnings are announced. It diminishes monotonically as the next 10K earnings announcement approaches; then, the bias pops up again as the next earnings cycle begins. The cyclical pattern is more exaggerated for non-investment grade (NIG) firms.

Figure IV Plot B illustrates that *ex-ante* bias, \widehat{AB} , computed from equation (3), exhibits the same event-time cyclicity as the *ex-post* bias. The figure’s three lines graph average predicted bias in event time, with time measured relative to the most recent 10K announcement month. The figure shows there is 12-month cyclicity for all firms, but the zeniths of the lines are larger for NIG firms. The same cyclicity cannot be seen in Figure IV Plot C, which graphs the same three lines for cycle-adjusted analyst bias, \widehat{AB}^{CA} .

Relating Bias Level to the Cross-Section of Expected Returns. To assess the predictive power of analyst bias for future returns, we run monthly cross-sectional regressions of risk-adjusted stock returns, $r_{i,t}^*$, as defined by the 4-factor model in equation (1), on one-month lagged *ex-ante* analyst bias, $\widehat{AB}_{i,t-1}$, (or the analogous cycle-adjusted bias) and various combinations of lagged firm-level variables. The firm-level variables are chosen as controls because they may be related to the cross-section of 4-factor risk-adjusted returns. The full

specification is:

$$\begin{aligned}
r_{i,t}^* &= c_{0,t} + c_{1,t} \widehat{AB}_{i,t-1} + c_{2,t} (\widehat{AB}_{i,t-1} \times D_{CR5,t-1}) + c_{3,t} D_{CR5,t-1} \\
&+ c_{4,t} r_{i,t-7:t-2} + c_{5,t} Dispersion_{i,t-1} + c_{6,t} SUE_{i,t-1} + \epsilon_{i,t}
\end{aligned} \tag{4}$$

where $D_{CR5,t-1}$ is a credit risk dummy variable that takes the value of 1 if the credit rating indicates that the firm is one of the 20% most credit risky stocks, $r_{i,t-7:t-2}$ is the cumulative return over months $[t - 7 : t - 2]$, $Dispersion_{i,t-1}$ is the dispersion in analyst forecasts in month $t - 1$, and $SUE_{i,t-1}$ is the last reported quarterly earnings surprise as benchmarked by the random walk model.¹²

Lagging the predicted bias regressor, $\widehat{AB}_{i,t-1}$, by one month ensures that all of the instruments used to generate this regressor are known prior to the start of the return month, including the $PastAB_i$ instrument. Recall, from the illustration on page 9, that predicted bias for May 1995 is derived, in part, from a past *ex-post* bias computed for May 1994's forecast that becomes known only in February 1995. It is the June 1995 risk-adjusted return that we correlate with the May 1995 predicted bias, but May 1995's predicted bias uses past bias information known in February 1995, as well as other instruments, like coverage, known at the end of April 1995.

Summary Statistics. Table II presents summary statistics. For each of five (approximately equally populated) groups sorted at the beginning of each month, Table II reports the time-series average of the monthly cross-sectional means of the firm's credit rating (using the number mapping described earlier), non-investment-grade dummy (NIG=1 for a BB+ rating or below), highly speculative dummy (HS = 1 for a B+ rating or below), distressed dummy (DIS=1 implies a CCC+ rating or below), predicted analyst bias (\widehat{AB}), cycle-adjusted analyst bias, (\widehat{AB}^{CA}), actual analyst bias (AB), market capitalization, book-to-

¹²The random walk model's earnings surprise is the last reported quarterly earnings minus the earnings four quarters prior, standardized by the standard deviation of the last eight of these earnings changes.

market ratio, past 6-month cumulative return, current-month raw and 4-factor risk-adjusted return, CAPM and 4-factor Carhart betas, analyst dispersion, coverage, and earnings surprise (SUE). Panel A reports averages for five credit rating groups, Panel B for five predicted analyst bias groups, and Panel C for quintiles sorted on within-cycle predicted analyst bias.

In Panel A, credit group 1 (CR1) represents the highest-rated firms (averaging an A+ S&P rating), CR2 is the second highest (average rating of BBB+), CR3 is the third highest (average of BBB-), CR4 is the fourth (average of BB-), and CR5 is the lowest (average of B).¹³ The predicted and actual analyst biases tend to reflect an overall bias towards optimism, with more optimism the higher is the credit risk. The predicted biases are 10.17%, 15.28%, 21.71%, 34.55%, and 37.93% across the CR1 to CR5 groups, respectively. The cycle-adjusted bias equivalent shows a similar pattern. The actual ex-post biases are again monotonically increasing in credit risk: 8.68%, 15.08%, 24.07%, 38.19%, and 42.06%.

In addition, the higher credit risk groups in Panel A tend to be smaller firms with larger book-to-market ratios, greater analyst forecast dispersion, lower coverage, smaller past earnings surprises, and lower past and current monthly returns. At the same time, these groups tend to have higher systematic risk, with CAPM betas averaging 1.38 for CR5 firms and 0.89 for CR1 firms, while 4-factor market betas average 1.26 in CR5 and 0.98 in CR1. The tendency of high credit risk firms to have lower returns leads to the anomalous credit risk effect documented in the literature (Dichev, 1998). The low credit risk CR1 portfolio earns 104 basis points per month, while the high credit risk CR5 portfolio earns 46 basis points per month, even though CR5 firms have higher market betas. The 58 basis points per month spread between the CR1 and CR5 returns is about the same (49 basis points) with the four-factor risk adjustment, leaving a sizable alpha anomaly of about 7.5% per year.

¹³The credit groups are designed to have approximately 20% of all firms in each group in each month. Quintiles are formed each month, hence the quintile cutoffs may change due to the changing composition of rated firms in the sample.

Table II Panel B reports the same summary statistics for each of five predicted analyst-bias quintiles known at the end of the prior month. Predicted analyst bias correlates negatively with returns—past six month returns, as well as current-month raw and risk-adjusted returns—and negatively with same-date (signed) earnings surprises and firm size; it correlates positively with same-date forecast dispersion and book-to-market ratio. Predicted analyst bias also exhibits the same positive relationship between bias and credit risk observed in Panel A. Indeed, the correlation patterns observed in Panel B for predicted analyst bias are similar to the patterns observed for credit risk in Panel A. Panel C’s construction is like Panel B’s, except it sorts firms into quintiles after ranking a firm’s predicted analyst bias relative to other firms sharing the same quarter q ($q = 0, 1, 2, \text{ or } 3$) of their earnings forecast cycle. Panel C’s results are highly similar to those in Panel B.

II. Results and Discussion

This section analyzes the joint role of credit rating and analyst bias as determinants of the cross-section of expected returns. It first studies the issue using independent portfolio sorts, then using cross-sectional regressions with multiple specifications and additional controls. Finally, it develops an extraordinarily simple model of information revelation and inference that is consistent with the results.

Risk-Adjusted Returns Sorted by Credit Rating and Bias. Credit risk’s positive correlation with analyst bias and negative correlation with average return and risk-adjusted return complicates inferences about the role of credit risk and analyst bias in share price formation. To better understand how the former two characteristics influence average returns, Table III shows average risk-adjusted returns and t -statistics for 25 portfolios. Panel A’s portfolios are obtained from an independent 5×5 sort on credit rating and predicted analyst bias, \widehat{AB} ; Panel B’s sort assigns bias quintiles based on \widehat{AB} ranks among groups of stocks sharing the

same forecast cycle, q . In both panels, the risk-adjusted returns are benchmarked against Carhart's (1997) 4-factor model, as described earlier.

In addition to the time series average of the monthly cross-sectional means of the risk-adjusted returns, the border rows and columns of Table III's two panels difference the entries (category 5 less category 1) in the corresponding rows and columns and provide t -statistics. The t -statistics are generated from the time series of risk-adjusted return differences between two equal-weighted portfolios corresponding to categories 5 and 1.

Perhaps the highlight of the table is the bias-related spread in the panels' CR5 rows, consisting of the most credit-distressed firms. For these firms, the Panel A spread between the quintile of stocks predicted to have the greatest and least analyst optimism is 161 basis points per month, or about 19% per year. For Panel B's cycle-specific sorts, it is 163 basis points per month. In both panels, these spreads are accounted for more by firms with the most conservative forecasts (AB1) than by firms with the most optimistic forecasts (AB5). Thus, a long-short investment strategy that takes advantage of this spread would earn greater risk-adjusted returns from the long leg of the strategy than from the short leg, although both legs significantly contribute to the strategy's overall profitability.

The results from the double sorts in Panels A and B of Table III lead to three conclusions. First, except for firms with the most credit risk, risk-adjusted returns across analyst bias groups do not significantly differ from zero. In rows CR1-CR4, representing the more creditworthy firms, virtually no relationship exists between our *ex-ante* measure of analyst bias and the cross section of expected returns (as implied by their statistical nearness to zero). Second, only for the least creditworthy firms (row CR5 in Panels A and B), a monotonic relationship exists between predicted analyst bias and risk-adjusted returns. The relationship is remarkably strong. Third, the apparent inverse relationship between a stock's credit risk and its expected return, documented in Table II Panel A and by prior research, could be an artifact of the correlation between credit risk and analyst bias. Plots A and B of

Figure II are 3D bar graphs of the average number of firms in each of the 25 cells in Panels A and B of Table III, respectively. As the block heights in Figure II indicate, more firms appear in the top left and bottom right corners of Table III's two panels, reflecting credit risk's correlation with forecast bias. The risk-adjusted return difference between firms in this extreme pair of corner cells thus could be due to differences in credit risk or to differences in analyst bias. Given the abundance of firms in these two cells, it is obvious that the inference of a negative premium for credit risk from the negative correlation between credit risk and return is fraught with peril: the negative correlation could easily be due to a negative return "premium" for analyst bias.

When controlling for both credit risk and analyst bias, as in Table III, the data pattern of the entire matrix of risk-adjusted returns complicates conclusions about the separate roles played by credit risk and analyst bias. Increases in analyst bias reduce returns, *but only for the highest credit risk firms*. Alternatively, extreme credit risk has a depressing effect on the risk-adjusted returns of high-bias ($\widehat{AB5}$) firms and an enhancing effect for low-bias ($\widehat{AB1}$) firms. The more parsimonious of these two mathematically equivalent explanations is that analyst bias reduces returns for high credit risk firms. In short, Occam's razor suggests that the risk-adjusted returns in the bottom row of the two panels is the sum of a fixed (positive or negative) premium for credit risk and a negative premium for analyst bias that exists only among the least creditworthy firms.

The sign and magnitude of the constant credit risk premium is even trickier to flesh out from the data. The premium for extreme credit risk depends on the analyst bias column against which credit risk is benchmarked. The bottom-right risk-adjusted return of the two 5x5 panels is the sum of the top right risk-adjusted return plus a negative premium for credit risk. Moving leftwards along the bottom row of either panel enhances returns because analyst bias diminishes (but the credit risk premium is the same for each column). The bottom-left risk-adjusted return is the sum of the top left risk-adjusted return and a positive credit risk

premium. Moving rightward along the bottom row reduces returns because analyst bias increases. There are also credit risk premiums that exist in the middle categories of analyst bias that are close to zero.

The proper analyst bias category for benchmarking the sign and magnitude of a bias-independent credit risk premium depends on which degree of predicted analyst bias has no effect on returns. In part, this depends on one's view of markets. If investors understand that analysts tend to be optimistic overall, and shift their beliefs so that the average degree of analyst bias does not tend to generate inflated or deflated share prices, then column $\widehat{AB3}$, which contains stocks with average degrees of bias, is the benchmark for calculating the credit risk premium. With this benchmark, the $CR5 - CR1$ difference implies a (negative) credit risk premium of -13 to -16 basis points per month, depending on the panel. However, predicted analyst bias is closest to zero in the $\widehat{AB1}$ quintile of firms, which carries a huge and statistically significantly positive ($CR5 - CR1$) credit risk premium of 96 and 93 basis points per month in Panels A and B, respectively. Regardless of one's belief about whether there is a small insignificant credit risk premium or a large positive premium, it is the relatively greater concentration of firms in the $CR5 \widehat{AB5}$ and $CR1 \widehat{AB1}$ cells that appears to generate the odd negative correlation between credit risk and return.

The final insight from Table III comes from the similarity of the risk-adjusted returns in its two panels. Panel B distinguishes itself from A by combining firms that are at different points in their earnings forecast cycles for its quintile formation. This means that the spread in analyst bias across the columns of Panel B is narrower than in A because the former's comparison for quintile classifications is based on a narrower range of firms. (This is also confirmed from comparisons of Table II Panel C with Table II Panel B.) Yet, the less extreme optimism differences in Table III Panel B generate similar spreads in risk-adjusted returns.

One interpretation of Panel B's similar (indeed, negligibly greater) risk-adjusted return spread between $AB5$ and $AB1$ for the least creditworthy firms is that return-influencing bias

is a trait a firm possesses separate from its point in the forecast cycle. The stretch to this conclusion, however, is that a large portion of firms are on relatively synchronized forecast cycles due to their common fiscal years, which generally leads to January and February as the most popular months for announcing annual earnings. Nevertheless, there are two good reasons to believe that price-inflating optimism is not tied to forecast cyclicalities, even if one concluded that cyclicalities represent a deliberate “walk down” in optimism bias by analysts¹⁴ rather than a mechanical artifact of quarterly earnings disclosures and guidance by firms. First, valuations should be based on perpetual earnings streams, which are extrapolated from near-term earnings forecasts. A 20% overly optimistic fourth-quarter earnings forecast for a cycle $q = 3$ firm is only a 5% overly optimistic forecast of its annual earnings, but it is more likely to lead gullible investors to overestimate earnings for the next several years by 20% than by 5%. Put another way, the sum of the next four quarters of unrealized earnings have no walk down in their forecast. Second, analysts publish forecasts for earnings they expect to be generated in fiscal years other than the upcoming fiscal year. We (and probably most other researchers) just lack the ability to study the impact of these longer horizon forecasts due to data limitations. Historical observation of such forecasts (and even their existence) is far rarer than the data collected for FY1.

The Marginal Effect of Analyst Bias and Credit Risk on Risk-Adjusted Returns. Table III’s data pattern makes it difficult to pin down a credit risk premium because analyst bias only influences the least creditworthy quintile of firms. This makes 5x5 quintile sorts, a relatively agnostic approach to specification, unable to identify a credit risk premium—as noted above, the premium depends on the analyst bias quintile that represents the benchmark. However, it might be possible to identify a credit risk premium for a model with more structure to the specification. The identification here is generated by the functional form of the regressors, as explained later. This motivation, as well as the desire to verify that additional controls do not alter our key finding, lead us to run cross-sectional regressions of

¹⁴See, for example, Richardson, Teoh, and Wysocki (2004) for the first use of the term “walk-down.”

risk-adjusted returns on a more tightly structured specification of analyst bias, credit risk, interactions between the two, and control variables.

The controls include several instruments used as predictors in equation (3)'s *ex-ante* bias regression: earnings surprises, dispersion in analyst forecasts, the momentum characteristic, and industry fixed effects. These controls have a correlation pattern with both *ex-ante* optimism and risk-adjusted returns that might account for any observed negative relationship between analyst bias and future risk-adjusted returns. Negative past earnings surprises predict analyst optimism (and vice versa), and both Ball and Brown (1968) and Foster, Olsen, and Shevlin (1984), among others, document that earnings surprises are positively correlated with future returns.¹⁵ Dispersion in analyst forecasts is correlated both with analyst optimism and future returns.¹⁶ Finally, while our risk-adjusted returns control for a stock's exposure to the momentum factor, they do not control for the momentum characteristic; the Carhart 4-factor model may not be an adequate control for the momentum characteristic. Like our other controls, the momentum characteristic correlates both with bias (negatively) and future returns (positively).¹⁷

To better understand how all of these variables interact, Table IV reports the time-series average coefficients and their Fama and MacBeth (1973) *t*-statistics from monthly cross-sectional regressions of risk-adjusted returns on five specifications of regressors based on credit risk (proxied for by the CR5 dummy), predicted analyst bias (or cycle-adjusted analyst bias), and controls (including industry dummies, which appear in all regressions).¹⁸ Panel A measures bias as the predicted bias, \widehat{AB} , while Panel B measures bias with its

¹⁵In part, the latter correlation stems from the ability of past surprises to predict future "surprises," despite the unfortunate choice of the latter's name. See, for example, Freeman and Tse (1989) and Bernard and Thomas (1990).

¹⁶See for example, Table II, as well as Ackert and Athanassakos (1997) and Diether, Malloy, and Scherbina (2002)

¹⁷Doukas, Kim, and Pantzalis (2005) find that positive excessive analyst coverage (driven by banking incentives and self-interest) is associated with overvaluation and low future returns. Omitting this variable makes findings of an inverse correlation more conservative than they need to be.

¹⁸The results are highly similar without industry controls and are omitted from the table for brevity.

cycle-adjusted cousin, \widehat{AB}^{CA} .

Table IV's five specifications yield many noteworthy insights about the return effects of analyst bias, credit risk, and their interaction. Some of these insights are similar to those discussed for Tables II and III. Specifications 1 and 2's univariate regressions (plus industry fixed effects) indicate that risk-adjusted returns inversely relate to both credit risk and predicted analyst bias (and cycle-adjusted bias), even with industry controls. However, because analyst bias and its interaction with credit risk now possesses a specific functional form, we can isolate the effect of credit risk from the effect of analyst bias. The remaining specifications indicate that extreme credit risk has no significant independent effect on returns, while analyst bias adversely influences future risk-adjusted returns, albeit for the higher credit risk firms (specifications 4 and 5). Thus, credit risk's role is primarily to enhance the return-depressing effects of analyst optimism bias.

In contrast to Table III, the cross-sectional regression's functional form in Table IV, particularly the continuity of analyst bias and its distribution in a linear model, does pin down the credit risk premium as small and insignificant. Recall from Table III that one could just as easily view the CR5 row as being constructed from a negative or a positive credit risk premium (depending on the column one starts at) plus a return depressing analyst bias effect that only operates for firms in the CR5 category. In Table IV, we see aspects of this alternative perspective as well, since the coefficient on the interaction term (specifications 4 and 5) implies a credit risk premium that can be positive or negative, large or small, depending on the sign and magnitude of a continuous analyst bias variable. Here, however, the least squares criterion prevents us from fixing the credit risk premium at any level if we want the best fit for all the data. Each possible premium for the CR5 category implies a different fit for the distribution of analyst bias in the CR5 category. There are not enough degrees of freedom in the linear specification to have a constant credit risk premium per unit of analyst bias that perfectly offsets an arbitrary CR5 risk premium choice. Thus, the

linearity of the model and the continuity of analyst bias identifies the credit risk premium as the coefficient on the CR5 dummy, particularly in specification 5’s “all-variable” regression. This premium is small and positive but insignificant for the predicted analyst bias regression (0.16) and the cycle-adjusted analyst bias regression (0.11).

Despite the similarity of the interaction coefficients in specification 5 of Panels A and B, the Panel B coefficient is of larger economic magnitude because the cycle-adjusted bias metric, \widehat{AB}^{CA} , is scaled up to match the analyst bias observed at the beginning of the annual forecast cycle, i.e., cycle $q = 0$. The larger economic magnitude (coupled with greater statistical significance) buttresses our earlier conclusion: that stock returns are not influenced by a lessening of optimism due to the natural (or intentional) walk down of optimistic earnings forecasts that take place over the earnings forecast cycle.

Table IV reveals that the economic magnitude of the effect of predicted analyst bias is impressive. Cycle-adjusted or not, comparing a not too uncommon 20% optimistically biased CR5 firm with another CR5 firm that has a 10% bias reduces risk-adjusted returns by about 10 basis points per month (0.1 times the sum of the analyst bias and interaction coefficients). The cycle-adjusted analyst bias spread between the top- and bottom-bias quintiles is about 75% (see Panel B of Table II). Thus, even accounting for the effect of control variables like momentum and earnings surprises, a diversified long-short strategy of quintile extremes could earn abnormal returns of 10% per year. Moreover, since optimism tends to persist, the transaction cost of a long-short strategy focused on trading bias-sorted portfolios of high credit stocks could be relatively small compared to many other similarly profitable strategies from the literature, like momentum or industry momentum.¹⁹

Discussion: A Model of Market Beliefs. The results in Tables III and IV are consistent with the intuitive belief that when share prices are distorted by extreme analyst bias, they ultimately have a tendency to revert (partially or fully) to their fair values based on a more

¹⁹See, for example, Jegadeesh and Titman (1993) and Moskowitz and Grinblatt (1999).

rational assessment of fundamentals. For some firms with consensus forecasts in the extreme tails of the analyst bias distribution, the forecasters are “getting caught.” The market initially believes the analyst forecast, but eventually, some other unbiased information comes out that dispels the notion that the current analyst earnings forecast should be believed. For firms in the left tail of the analyst bias distribution, the alternative information is more likely to generate an upward revision in earnings estimates by *market participants*; for firms in the right tail, alternative information is more likely to generate downward revisions in these estimates. Compounding this effect is the possibility that analyst forecasts mean revert.

The simplest model consistent with the results in Tables III and IV is one where market participants value the firm at the beginning of each month based on the prevailing consensus analyst forecast. However, at the end of the month, with some probability p , the market receives a signal about earnings from another source. The signal is an unbiased forecast of earnings and, if it is received, generates a substantial revaluation of the firm by market participants. The signal could be earnings guidance by the firm, quarterly earnings releases, press commentary, or any other earnings-specific information from a source other than the analysts.

Without loss of generality, we assume that, if the signal is received, the market completely discards the consensus analyst forecast and, instead, employs the alternative information to form its earnings expectations. Section III later discusses how to generalize this to a model where the market’s reliance on the analyst forecast is merely diminished, rather than eliminated, when the alternative information is disclosed. For now, we prefer the pedagogy and simpler algebra of the more extreme assumption. If the firm’s analyst forecast bias remains constant over the month, this assumption implies that the expected difference in the market’s earnings forecast over a month is the probability that the alternative information is received times the degree to which the new information alters the market’s earnings forecast. The model’s transparency is also facilitated by avoiding the Jensen’s inequality effects of

convexity and concavity on the market's inferences. To abstract from this complexity, we assume that the market's valuation is a linear function of its earnings forecast, and in turn, the market's earnings forecast is a linear function of the information the market believes. Specifically, (and dropping the firm subscript i to simplify notation), at the end of month $t - 1$, the market's earnings per share forecast is a positive linear function of a consensus analyst EPS forecast that contains a bias, B :

$$MktForecastEPS_{t-1}^T = b_0 + b_1 ConForecastEPS_{t-1}^T \quad (5)$$

$$= b_0 + b_1(EPS^T + \tilde{u}_{t-1}^T + B_{t-1}), \quad (6)$$

where \tilde{u}_t^T is the end of month t analyst *ex-post* forecast error in the absence of a bias, and the market's valuation based on this forecast is:

$$P_{t-1} = k_0 + k_1 MktForecastEPS_{t-1}^T. \quad (7)$$

At the end of month t , no signal is revealed with probability $1 - p$. In this case,

$$MktForecastEPS_t^T = b_0 + b_1 ConForecastEPS_t^T \quad (8)$$

$$= b_0 + b_1(EPS^T + \tilde{u}_t^T + B_t), \quad (9)$$

the date t share price is

$$P_t = k_0 + k_1 MktForecastEPS_t^T, \quad (10)$$

and the change in the share price is

$$\Delta P_t = k_1 b_1 (B_t + \tilde{u}_t^T - B_{t-1} - \tilde{u}_{t-1}^T). \quad (11)$$

In the case where a signal is received,

$$MktForecastEPS_t^T = b_0 + b_1SignalEPS_t^T \quad (12)$$

$$= b_0 + b_1(EPST + \tilde{v}_t^T), \quad (13)$$

where v_t^T is the end of month t *ex-post* forecast error of the signal, and the change in the share price is:

$$\Delta P_t = k_1 b_1 (\tilde{v}_t^T - B_{t-1} - \tilde{u}_{t-1}^T). \quad (14)$$

Hence, the expected change in the price is

$$E[\Delta P_t] = k_1 b_1 [-pB_{t-1} + (1-p)(B_t - B_{t-1})]. \quad (15)$$

Because $k_1 b_1$ is positive and because the bias tends to mean revert, the expected price change of firms with low analyst bias at $t - 1$ exceeds the expected price change of firms with higher bias at $t - 1$. Mean reversion in analyst forecast bias is supported by Figure III. Figure III plots the event-time paths of both average predicted analyst bias (\widehat{AB} , two green lines) and average cycle-adjusted analyst bias (\widehat{AB}^{CA} , two red lines) for the two extreme quintile portfolios ranked at event date 0. The lines of the top bias quintile decline while the lines of the bottom bias quintile increase over the subsequent 18 months. For firms in the two extreme predicted bias quintiles, the greatest rate of mean reversion is witnessed over the first month. The bias decrease over the month for firms with large amounts of analyst bias and increase for firms at the opposite end of the spectrum enhances the effect on the price change contributed from the alternative information—an effect which relies on changes in investor gullibility. At the beginning of the month, investors are fully gullible about firms with extreme amounts of analyst bias. At the end of the month, they are completely ignoring the analyst forecast if the alternative information arrives. However, mean reversion in analyst bias would force prices to revert to fair value in the states where no alternative

information is revealed, which occurs with probability $1 - p$, would also contribute to share price reversion to fair value.

The model is easily parameterized to explain why bias only influences the risk-adjusted returns of firms with high credit risk. The market's sensitivity to analyst forecasts, and ultimately the bias inherent in such forecasts, might differ across the credit categories for a variety of reasons. Differences in p , k_1 , and b_1 across the credit rating categories could therefore explain why risk-adjusted return spreads, arising from cross-sectional differences in analyst bias, are only evident for high-credit-risk firms.

III. Conclusion

This paper demonstrates that analyst bias has profound effects on the risk-adjusted returns of stocks with high credit risk. High credit risk stocks with greater degrees of analyst optimism earn lower returns. This finding is consistent with the intuitive belief that analyst bias inflates (or deflates) share prices. Direct measurement of price inflation and deflation is a tall order, as it requires a proper model of fair value. Rather than commit to any particular model of fair value, we simply show that when analyst bias is extreme, risk-adjusted returns are consistent with share price inflation (or deflation) that reverts to zero. Equivalently, prices, distorted by analyst bias, gravitate to fair value. This conclusion follows from the plausible assumption that fair value, adjusted for risk and the time value of money, follows a random walk, a point elaborated on below.

While the return difference between stocks with extreme differences in analyst bias could potentially be explained by mean reversion in analyst optimism, market participants tend to find other less-biased sources of information that lead to less reliance on biased forecasts. Thus, changes in gullibility, even over a period as short as a month, could also lead to the abnormal risk-adjusted returns documented in this paper. Firms with high degrees

of analyst optimism bias lose their price inflation when gullibility with respect to analyst forecasts diminishes. Likewise, the deflated prices of firms with ultra-conservative analyst forecasts lose their deflation when that reliance diminishes as well. In short, this is an efficient markets anomaly with plausible root causes that derive from irrational behavior.

The theoretical contribution of the paper is a mathematical model of the influence of analyst bias on returns. The model is very stylized and is cast in a static two-period setting, but is easily generalized to a dynamic model where the market partially relies at all times on both the consensus analyst forecast and on the most recent realization of an alternative signal that is refreshed with a Poisson process. Only two additional assumptions are required. First, the reliance on the alternative signal at the end of the month must exceed the reliance at the beginning of the month if the alternative signal is refreshed. In other words, the market's gullibility must decline when the Poisson realization is achieved. For the model to be dynamically consistent and stationary, the market's gullibility must also increase in months where the Poisson realization does not take place. The rate of increase in gullibility in those months is determined by the rate of gullibility decrease in the Poisson realization months, the Poisson intensity of the event, and the degree of mean reversion in analyst forecast bias for states with no alternative information revelation.

The second critical assumption for generalizing the story is that changes in the model's alternative signal and the de-biased analyst forecast, v_t and u_t , be martingales. At the root of this assumption is a tautology: A stock's price must be the sum of a rational component—a fair value—and a (positive, zero, or negative) behavioral component sometimes referred to as “share price inflation.” Our model is one where the inflation mean reverts. To allow data to distinguish the behavioral component from the fair value component, we have to assume that changes in a stock's fair value, adjusted for risk and the time value of money, are purely random. The model's linearity and the Bayesian conclusion that rational current-year earnings estimates—wherever they appear in the model—follow a martingale, achieves

this goal.

The analyst-bias anomaly not only leads to highly profitable trading strategies, but explains the credit risk anomaly—the fact that high-beta high-credit-risk stocks earn far lower returns than low-beta low-credit-risk stocks. The inferior returns of high credit risk stocks disappears once analyst bias is controlled for. This leads to the conclusion that analyst bias may be an important omitted variable in many regressions used to assess the determinants of the cross-section of expected returns. We can only speculate, however, about why analyst bias has so much more of an effect on the returns of high credit risk stocks than on the returns of low credit risk stocks.

Analyst bias also has some relationship with momentum, a topic we intend to pursue as an extension of this paper. High-bias stocks are those with low past returns and low future returns. Controlling for analyst bias leads to a fairly negligible return effect from the momentum characteristic, but the four-factor risk-adjusted returns studied here already take out a momentum factor. Hence, further study is needed to assess the degree to which analyst bias accounts for the momentum anomaly.

There are many promising avenues for future research. We have studied effects on average risk-adjusted returns, but have yet to focus on how analyst bias influences the entire distribution of returns. We also have yet to study the causes of analyst bias. While we can predict analyst bias and perhaps improve upon the prediction accuracy with additional research, we do not yet have an economic model of why analysts forecast earnings with a bias, and what motivates cross-sectional differences in that bias. Finally, one of the most promising aspects of being able to predict analyst bias is that it may lead to better estimates of stock price inflation. Models that directly measure stock price inflation exist. However, the risk-adjusted returns that we predict are likely to stem from stock prices inflation and deflation. These risk-adjusted returns may offer a promising alternative for measuring inflation.

References

- Ackert, Lucy F., and George Athanassakos, 1997, Prior Uncertainty, Analyst Bias, and Subsequent Abnormal Returns, *Journal of Financial Research* 20, 263–273.
- Ali, Ashiq, April Klein, and James Rosenfeld, 1992, Analysts' Use of Information about Permanent and Transitory Earnings Components in Forecasting Annual EPS, *The Accounting Review* 67, pp. 183–198.
- Avramov, Doron, and Tarun Chordia, 2006, Asset Pricing Models and Financial Market Anomalies, *Review of Financial Studies* 19, 1001–1040.
- Avramov, Doron, Tarun Chordia, Gergana Jostova, and Alexander Philipov, 2009a, Credit Ratings and The Cross-Section of Stock Returns, *Journal of Financial Markets* 12, 469–499.
- Avramov, Doron, Tarun Chordia, Gergana Jostova, and Alexander Philipov, 2009b, Dispersion in Analysts' Earnings Forecasts and Credit Rating, *Journal of Financial Economics* 91, 83–101.
- Ball, Ray, and Philip Brown, 1968, An Empirical Evaluation of Accounting Income Numbers, *Journal of Accounting Research* 6, 159–178.
- Beaver, William, Maureen McNichols, and Richard Price, 2007, Delisting returns and their effect on accounting-based market anomalies, *Journal of Accounting and Economics* 43, 341 – 368.
- Bernard, Victor L., and Jacob K. Thomas, 1990, Evidence that stock prices do not fully reflect the implications of current earnings for future earnings, *Journal of Accounting and Economics* 13, 305–340.
- Brennan, Michael J., Tarun Chordia, and A. Subrahmanyam, 1998, Alternative factor specifications, security characteristics, and the cross-section of expected stock returns., *Journal of Financial Economics* 49, 345–373.
- Campbell, John Y., Jens Hilscher, and Jan Szilagyi, 2008, In Search of Distress Risk, *Journal of Finance* 63, 2899–2939.
- Carhart, Mark M., 1997, On Persistence in Mutual Fund Performance, *Journal of Finance* 52, 57–82.
- De Bondt, Werner F. M., and Richard H. Thaler, 1990, Do Security Analysts Overreact?, *The American Economic Review* 80, pp. 52–57.
- Dichev, Ilia D., 1998, Is the Risk of Bankruptcy a Systematic Risk?, *Journal of Finance* 53, 1131–1147.
- Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, Difference of Opinion and the Cross Section of Stock Returns, *Journal of Finance* 57, 2113–2141.

- Doukas, John A., Chansog(Francis) Kim, and Christos Pantzalis, 2005, The Two Faces of Analyst Coverage, *Financial Management* 34, pp. 99–125.
- Easterwood, John C., and Stacey R. Nutt, 1999, Inefficiency in Analysts' Earnings Forecasts: Systematic Misreaction or Systematic Optimism?, *Journal of Finance* 54, 1777–1797.
- Fama, Eugene F., and Kenneth R. French, 1992, The Cross-Section of Expected Stock Returns, *Journal of Finance* 47, 427–465.
- Fama, Eugene F., and James D. MacBeth, 1973, Risk, Return, and Equilibrium: Empirical Tests, *Journal of Political Economy* 81, 607–636.
- Foster, George, Chris Olsen, and Terry Shevlin, 1984, Earnings Releases, Anomalies, and the Behavior of Security Returns, *The Accounting Review* 59, 574–603.
- Francis, Jennifer, and Donna Philbrick, 1993, Analysts' Decisions As Products of a Multi-Task Environment, *Journal of Accounting Research* 31, pp. 216–230.
- Freeman, Robert N., and Senyo Tse, 1989, The Multiperiod Information Content of Accounting Earnings: Confirmations and Contradictions of Previous Earnings Reports., *Journal of Accounting Research* 27, 49 – 79.
- Hayes, Rachel M., 1998, The Impact of Trading Commission Incentives on Analysts' Stock Coverage Decisions and Earnings Forecasts, *Journal of Accounting Research* 36, pp. 299–320.
- Hilary, Gilles, and Lior Menzly, 2006, Does past Success Lead Analysts to Become Overconfident?, *Management Science* 52, pp. 489–500.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *Journal of Finance* 48, 65–91.
- LaPorta, Rafael, 1996, Expectations and the Cross-Section of Stock Returns, *Journal of Finance* 51, 1715–1742.
- Lin, Hsiou-wei, and Maureen F. McNichols, 1998, Underwriting relationships, analysts' earnings forecasts and investment recommendations, *Journal of Accounting and Economics* 25, 101 – 127.
- Moskowitz, Tobias J., and Mark Grinblatt, 1999, Do Industries Explain Momentum?, *Journal of Finance* 54, 1249–1289.
- Richardson, Scott, Siew Hong Teoh, and Peter D. Wysocki, 2004, The Walk-Down to Beatable Analyst Forecasts: The Role of Equity Issuance and Insider Trading Incentives, *Contemporary Accounting Research* 21, 885–924.
- Womack, Kent L., 1996, Do Brokerage Analysts' Recommendations Have Investment Value?, *Journal of Finance* 51, 137–167.

Table I. Prediction of analyst bias

This table, using monthly data on all U.S. exchange-listed non-penny stocks with credit ratings from 1986-2012, estimates the coefficients of (3). The estimation uses panel regressions run separately for groups of firms sorted by the number of 10Q earnings statements (denoted $q = 0, 1, 2,$ or 3) released by the firm for the relevant fiscal year. t -statistics, in parentheses, and the R-squared of each regression are reported as well.

	$q = 0$	$q = 1$	$q = 2$	$q = 3$
Constant	0.095 (1.95)	-0.029 (-0.74)	-0.091 (-3.02)	-0.051 (-1.91)
<i>PastAB</i>	0.123 (26.05)	0.100 (22.67)	0.082 (19.16)	0.058 (12.28)
<i>Dispersion</i>	0.248 (11.04)	0.367 (20.55)	0.511 (35.34)	0.453 (35.43)
<i>Coverage</i>	-0.001 (-1.35)	-0.002 (-3.67)	-0.002 (-4.10)	-0.002 (-5.17)
<i>PastRet</i> ⁻	-0.695 (-21.87)	-0.757 (-23.89)	-0.462 (-21.44)	-0.298 (-18.44)
<i>PastRet</i> ⁺	-0.345 (-16.61)	-0.249 (-14.94)	-0.183 (-13.67)	-0.067 (-5.22)
<i>SUE</i> ⁻	-0.040 (-9.34)	-0.063 (-16.47)	-0.035 (-13.52)	-0.035 (-14.41)
<i>SUE</i> ⁺	-0.027 (-9.39)	-0.029 (-11.10)	-0.018 (-9.88)	-0.008 (-4.72)
<i>D_{Small}</i>	0.141 (6.04)	0.112 (5.67)	0.099 (6.44)	0.093 (7.38)
<i>D_{Value}</i>	0.051 (6.17)	0.061 (8.64)	0.050 (8.95)	0.021 (4.41)
Adjusted R^2	0.085	0.086	0.080	0.070

Table II. Descriptive Statistics

Panel A ranks stocks each month, t , into quintiles based on their S&P credit rating. It reports the time-series average of monthly cross-sectional mean firm characteristics in month $t + 1$. Analyst bias is the analyst's EPS forecast for fiscal year 1 minus the actual EPS for the year standardized by the actual EPS. SUE is the last reported quarterly EPS minus the EPS four quarters prior, divided by the standard deviation of the last eight EPS changes. Dispersion is the standard deviation in analysts' EPS forecasts standardized by the absolute value of the consensus forecast. Analyst coverage is the number of analysts covering the firm. Panel B repeats this analysis, but sorting rated stocks into \widehat{AB} quintiles instead of rating quintiles. Panel C does the same across \widehat{AB} cycle-specific quintiles. Stocks priced below \$1 are removed. The sample period is January 1986 to December 2012.

Panel A. Characteristics across rating quintiles

Characteristic	Rating Quintile (CR1=Best, CR5=Worst)				
	CR1	CR2	CR3	CR4	CR5
S&P letter rating	A+	BBB+	BBB-	BB-	B
Fraction NIG firms ($\leq BB+$)	0.00	0.00	0.41	0.98	1.00
Fraction rated B+ or worse	0.00	0.00	0.00	0.24	0.99
Fraction rated CCC+ or worse	0.00	0.00	0.00	0.00	0.16
\widehat{AB} (%)	8.68	15.08	24.07	38.19	42.06
\widehat{AB} (%)	10.17	15.28	21.71	34.55	37.93
\widehat{AB}^{CA} (%)	15.42	23.18	32.75	52.78	58.86
Market capitalization (\$bln)	18.76	5.07	2.78	1.17	0.71
Book-to-market ratio	0.64	0.82	0.86	0.87	1.17
Past six-month return (%)	6.58	6.84	7.25	6.40	2.97
Current month return (%)	1.04	1.08	1.15	0.96	0.46
Risk-adjusted return (%)	0.10	0.11	0.09	-0.07	-0.39
CAPM beta	0.89	0.93	1.07	1.24	1.38
MKT beta	0.98	1.01	1.12	1.20	1.26
SMB beta	-0.05	0.21	0.47	0.81	0.98
HML beta	0.41	0.53	0.56	0.50	0.42
UMD beta	-0.06	-0.10	-0.16	-0.23	-0.32
Dispersion in analyst forecasts	0.05	0.08	0.13	0.23	0.41
Coverage (# of analysts)	17.68	13.23	11.42	8.01	6.21
SUE	0.84	0.49	0.43	0.25	-0.03

Table II. Descriptive Statistics (continued)Panel B. Characteristics across \widehat{AB} quintiles

Characteristic	Bias Quintile ($\widehat{AB}1$ =Low, $\widehat{AB}5$ =High)				
	$\widehat{AB}1$	$\widehat{AB}2$	$\widehat{AB}3$	$\widehat{AB}4$	$\widehat{AB}5$
S&P letter rating	A–	BBB+	BBB	BBB–	BB
Fraction NIG firms ($\leq BB+$)	0.16	0.19	0.26	0.41	0.67
Fraction rated B+ or worse	0.05	0.05	0.07	0.13	0.31
Fraction rated CCC+ or worse	0.01	0.00	0.00	0.01	0.02
AB (%)	1.36	4.10	9.39	20.35	48.81
\widehat{AB} (%)	0.78	8.42	14.45	23.24	49.63
\widehat{AB}^{CA} (%)	–0.39	13.41	23.74	38.03	75.39
Market capitalization (\$bln)	15.20	9.83	7.08	4.88	2.46
Book-to-market ratio	0.53	0.70	0.78	0.85	0.94
Past six-month return (%)	18.56	10.50	6.38	2.30	–5.77
Current month return (%)	1.08	1.02	1.02	1.01	0.81
Risk-adjusted return (%)	0.16	0.03	0.02	–0.00	–0.22
CAPM beta	0.98	0.94	1.00	1.11	1.26
MKT beta	1.01	1.01	1.06	1.14	1.23
SMB beta	0.11	0.18	0.27	0.43	0.69
HML beta	0.30	0.47	0.52	0.53	0.52
UMD beta	–0.08	–0.09	–0.13	–0.17	–0.25
Dispersion in analyst forecasts	0.03	0.04	0.06	0.10	0.32
Coverage (# of analysts)	16.99	14.44	13.23	11.95	9.99
SUE	1.91	0.89	0.44	0.01	–0.52

Panel C. Characteristics across cycle-specific \widehat{AB} quintiles

Characteristic	Cycle-Specific Bias Quintile ($\widehat{AB}1$ =Low, $\widehat{AB}5$ =High)				
	$\widehat{AB}1$	$\widehat{AB}2$	$\widehat{AB}3$	$\widehat{AB}4$	$\widehat{AB}5$
S&P letter rating	A–	BBB+	BBB	BBB–	BB
Fraction NIG firms ($\leq BB+$)	0.15	0.20	0.28	0.46	0.67
Fraction rated B+ or worse	0.05	0.05	0.08	0.16	0.32
Fraction rated CCC+ or worse	0.01	0.00	0.00	0.01	0.02
AB (%)	1.62	5.08	11.64	24.16	50.48
\widehat{AB} (%)	1.67	9.65	16.26	27.89	49.49
\widehat{AB}^{CA} (%)	0.76	14.15	24.72	44.10	80.34
Market capitalization (\$bln)	15.14	9.57	6.72	4.36	2.39
Book-to-market ratio	0.55	0.69	0.79	0.84	0.98
Past six-month return (%)	18.60	10.34	7.11	–0.07	–7.30
Current month return (%)	1.06	1.01	1.02	1.03	0.81
Risk-adjusted return (%)	0.16	0.00	0.10	–0.02	–0.30
CAPM beta	0.96	0.95	1.02	1.14	1.25
MKT beta	1.00	1.01	1.08	1.16	1.23
SMB beta	0.10	0.19	0.31	0.48	0.69
HML beta	0.31	0.46	0.51	0.54	0.54
UMD beta	–0.08	–0.10	–0.13	–0.19	–0.25
Dispersion in analyst forecasts	0.03	0.04	0.06	0.12	0.33
Coverage (# of analysts)	16.97	14.32	13.00	11.34	10.13
SUE	1.85	0.91	0.41	–0.08	–0.63

Table III. Risk-Adjusted Returns Sorted on Analyst Bias

Panel A shows average Carhart risk-adjusted returns for 25 categories of firms constructed from the intersection of independently sorted quintiles of prior month credit ratings and predicted analyst bias. Panel B is like Panel A except that there are only two credit ratings groups: investment grade (IG, rated *BBB-* or better) and non-investment grade (NIG, *BB+* or worse). Panel C is like Panels B, except that a firm's predicted analyst bias quintile is determined by its rank relative to other firms within the same quarterly forecast cycle. Panel D is like Panel B except that the division is into firms rated non highly speculative (NHS, AAA to BB-) and highly speculative (HS, B+ to D). The last column (row) in Panels B, C, and D display the return spreads $\widehat{AB5} - \widehat{AB1}$ (NIG-IG) after controlling for momentum, earnings momentum, and analyst dispersion as well as industry. In particular, we run monthly Fama-MacBeth cross-sectional regressions of Carhart risk-adjusted returns on dummies indicating the independently sorted IG/NIG and $\widehat{AB1}$ to $\widehat{AB5}$ groups, while also including in the regressions $r_{t-7:t-2}$, SUE_{t-1} , and $Disp_{t-1}$ characteristics as well as industry dummies. To compute each return spread and its t-statistic, we omit the dummy with respect to which the corresponding spread is computed and report the dummy and its associated t-statistic.

Panel A: 5×5 independent sort on predicted analyst bias, \widehat{AB}_{t-1} , and credit rating, CR .

Rating Groups	Bias Quintiles ($\widehat{AB1}$ =Low, $\widehat{AB5}$ =High)					$\widehat{AB5} - \widehat{AB1}$
	$\widehat{AB1}$	$\widehat{AB2}$	$\widehat{AB3}$	$\widehat{AB4}$	$\widehat{AB5}$	
CR1	0.09 (1.05)	0.02 (0.32)	0.02 (0.24)	0.23 (1.92)	0.06 (0.26)	-0.03 (-0.11)
CR2	0.21 (1.98)	0.09 (1.01)	0.00 (0.01)	0.04 (0.30)	0.19 (1.02)	-0.02 (-0.10)
CR3	-0.07 (-0.52)	0.03 (0.24)	-0.03 (-0.26)	0.09 (0.70)	-0.12 (-0.76)	-0.05 (-0.21)
CR4	0.20 (0.99)	-0.05 (-0.32)	0.18 (1.15)	-0.25 (-1.87)	-0.07 (-0.50)	-0.27 (-1.17)
CR5	1.02 (3.02)	0.01 (0.05)	-0.14 (-0.65)	-0.13 (-0.72)	-0.58 (-3.15)	-1.61 (-4.26)
CR5-CR1	0.95 (2.80)	-0.00 (-0.00)	-0.16 (-0.69)	-0.36 (-1.66)	-0.62 (-2.16)	

Panel B. 5x5 earnings-cycle-specific sort on \widehat{AB}_{t-1} and CR

Rating Groups	Cycle-Specific Bias Quintiles ($\widehat{AB1}$ =Low, $\widehat{AB5}$ =High)					$\widehat{AB5} - \widehat{AB1}$
	$\widehat{AB1}$	$\widehat{AB2}$	$\widehat{AB3}$	$\widehat{AB4}$	$\widehat{AB5}$	
CR1	0.08 (0.94)	0.06 (0.78)	0.16 (1.73)	0.07 (0.56)	-0.02 (-0.07)	-0.10 (-0.44)
CR2	0.16 (1.73)	0.02 (0.23)	0.14 (1.53)	0.15 (1.30)	0.26 (1.50)	0.10 (0.48)
CR3	0.09 (0.79)	-0.02 (-0.22)	0.07 (0.68)	0.13 (1.19)	-0.06 (-0.38)	-0.15 (-0.77)
CR4	0.15 (0.91)	-0.10 (-0.72)	0.09 (0.75)	0.00 (0.01)	-0.18 (-1.30)	-0.33 (-1.54)
CR5	1.04 (3.29)	0.31 (1.34)	0.03 (0.18)	-0.09 (-0.62)	-0.59 (-3.85)	-1.63 (-4.60)
CR5-CR1	0.96 (3.02)	0.25 (1.05)	-0.13 (-0.69)	-0.16 (-0.86)	-0.58 (-2.34)	

Table IV. Cross-Sectional Regressions of Returns on Firm Characteristics

We run Fama-MacBeth monthly cross-sectional regressions of Carhart risk-adjusted returns, $r_{i,t}^*$, on combinations of lagged firm-level variables:

$$r_{i,t}^* = c_{0,t} + c_{1,t} \widehat{AB}_{i,t-1} + c_{2,t} (\widehat{AB}_{i,t-1} \times D_{CR5,t-1}) + c_{3,t} D_{CR5,t-1} + c_{4,t} r_{i,t-7:t-2} + c_{5,t} SUE_{i,t-1} + c_{5,t} Disp_{i,t-1} \quad (16)$$

where $r_{i,t-7:t-2}$ is the past six month cumulative return, $SUE_{i,t-1}$ is the quarterly earnings surprise, $Disp_{i,t-1}$ is the dispersion in analyst earnings forecasts, and $D_{CR5,t-1}$ is a dummy variable indicating whether the firm belongs to the quintile of worst-rated firms in month $t-1$. All specifications include dummies for the twenty industries of Moskowitz and Grinblatt (1999). Panel A uses predicted analyst bias \widehat{AB} , while Panel B uses cycle-adjusted analyst bias, \widehat{AB}^{CA} .

Panel A: Using predicted analyst bias, \widehat{AB}

Characteristic	Regression Specification				
	1	2	3	4	5
Constant	-0.07 (-0.69)	-0.01 (-0.09)	-0.06 (-0.57)	-0.10 (-0.96)	-0.30 (-2.72)
$\widehat{AB}_{i,t-1}$	-0.58 (-2.57)		-0.55 (-2.52)	-0.36 (-1.66)	0.16 (0.73)
$\widehat{AB}_{i,t-1} \times D_{NIG,t-1}$				-0.87 (-1.78)	-1.12 (-2.09)
$D_{NIG,t-1}$		-0.57 (-4.08)	-0.22 (-1.26)	0.09 (0.39)	0.13 (0.52)
$r_{t-7:t-2}$					0.37 (1.35)
SUE_{t-1}					0.05 (3.04)
$Disp_{t-1}$					-0.18 (-0.82)
Industry fixed effects	Yes	Yes	Yes	Yes	Yes
Adjusted R^2 (%)	6.69	5.01	7.14	7.46	8.53

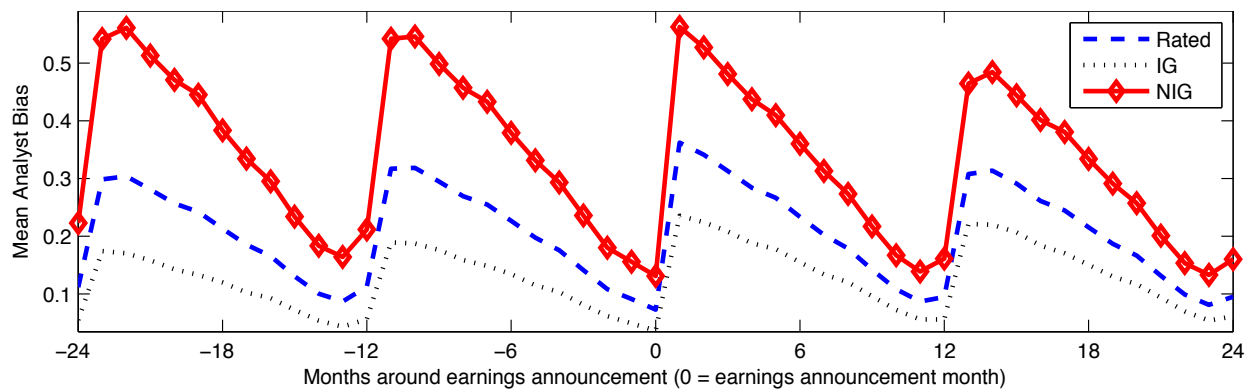
Panel B: Using cycle-adjusted analyst bias, \widehat{AB}^{CA}

Constant	-0.07 (-0.70)	-0.01 (-0.09)	-0.06 (-0.59)	-0.11 (-1.00)	-0.29 (-2.54)
$\widehat{AB}_{i,t-1}$	-0.41 (-2.77)		-0.40 (-2.76)	-0.25 (-1.69)	0.11 (0.67)
$\widehat{AB}_{i,t-1} \times D_{NIG,t-1}$				-0.76 (-2.46)	-1.06 (-3.02)
$D_{NIG,t-1}$		-0.57 (-4.08)	-0.20 (-1.15)	0.18 (0.84)	0.30 (1.25)
$r_{t-7:t-2}$					0.29 (1.01)
SUE_{t-1}					0.04 (2.69)
$Disp_{t-1}$					-0.23 (-0.97)
Industry fixed effects	Yes	Yes	Yes	Yes	Yes
Adjusted R^2 (%)	6.77	5.01	7.22	7.52	8.56


Figure I. Analyst bias around 10K earnings announcements.

The figure presents the observation-weighted time-series average of the mean analyst bias around 10K earnings announcements. Results are presented for all rated firms, as well as firms rated investment grade (IG) and non-investment grade (NIG) during month 0. The observation-weighted average weights each monthly mean by the number of observations (earnings announcements) during that month, giving more weight to months with more earnings announcements.

Plot A: *Ex-post* analyst bias: AB



Plot B: Predicted analyst bias: \widehat{AB}



Plot C: Cycle-adjusted analyst bias: \widehat{AB}^{CA}

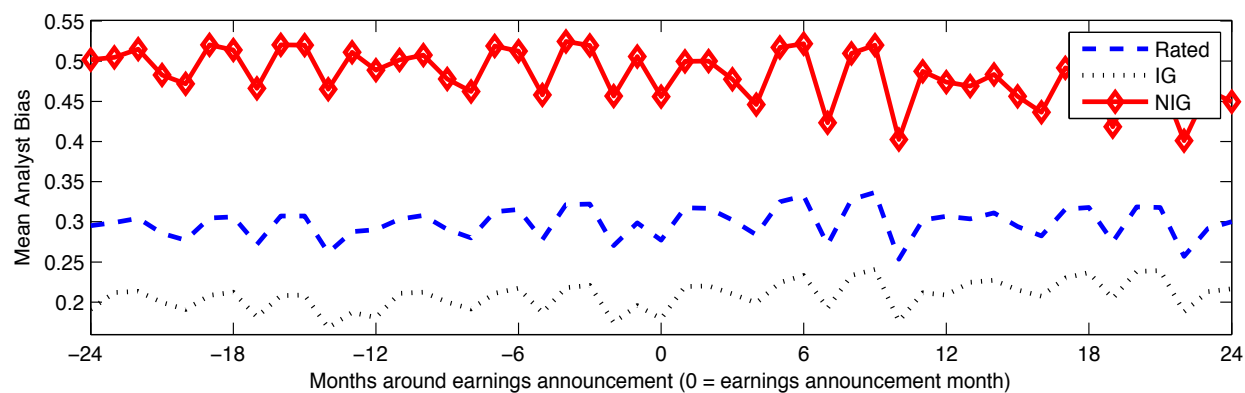



Figure II. Firm counts by predicted analyst bias and credit rating.

Each month, t , stocks are ranked into quintiles based on S&P credit rating and (independently) into quintiles based on predicted analyst bias, \widehat{AB} . The 3-D bar plot depicts the average number of firms populating the 25 categoris based on the intersection of the two sets of quintiles.

Plot A. Average observation counts using \widehat{AB} .



Plot B. Average observation counts using cycle-adjusted \widehat{AB} .


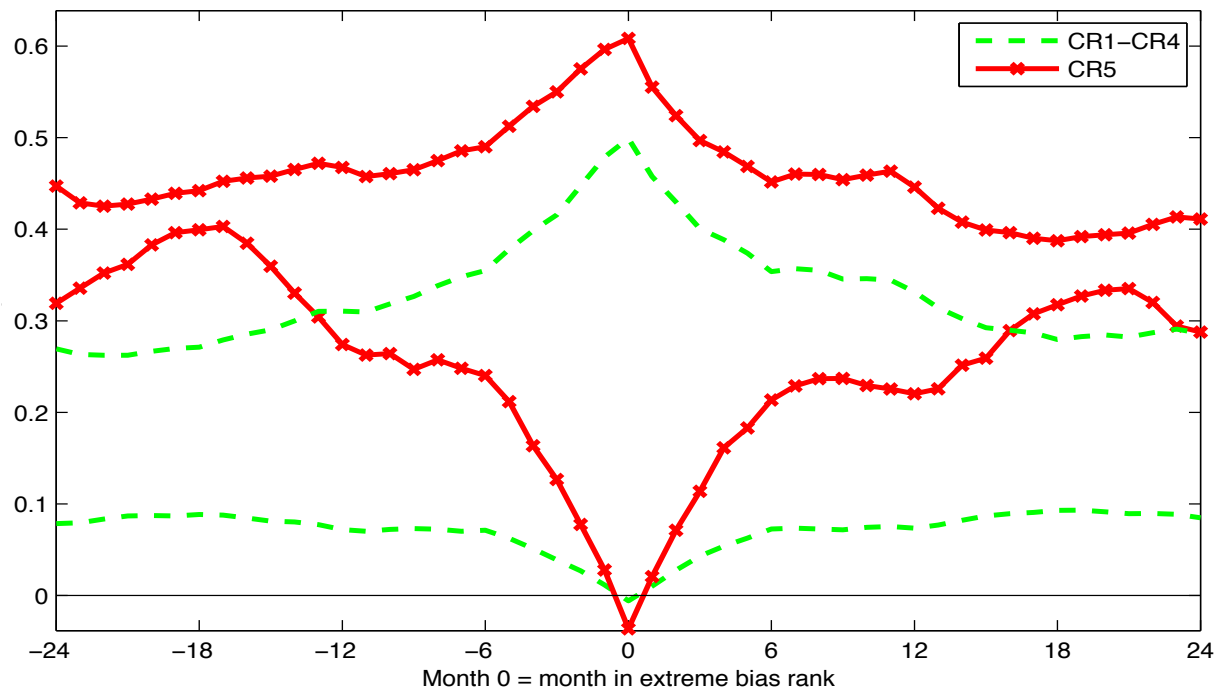


Figure III. Analyst bias around extreme bias ranks.

The figure presents the time-series average of the monthly cross-sectional mean predicted analyst bias (Plot A), or cycle-adjusted analyst bias (Plot B), around month 0 in which a firm is in the lowest (two bottom lines) or highest (two top lines) analyst bias quintile.

Plot A: Predicted analyst bias (\widehat{AB})



Plot B: Cycle-adjusted analyst bias (cycle-specific \widehat{AB})

