Forecasting styles

Michael S. Drake*
Brigham Young University
mikedrake@byu.edu

James R. Moon, Jr.
Georgia Institute of Technology
robbie.moon@scheller.gatech.edu

James D. Warren
University of Arkansas
jwarren@walton.uark.edu

November 2022

We thank Kimball Chapman, Jake Thornock, Ken Merkley, Joe Pacelli, Quinn Swanquist, Ke Wang (discussant), and workshop participants at BYU, Georgia Tech, the University of Kentucky, and the 2022 Conference on Financial Economics and Accounting (CFEA) for helpful comments and suggestions.

* Corresponding Author

Forecasting styles

ABSTRACT

We employ a novel machine-learning technique ("clusterwise linear regression") to identify five distinct forecasting styles employed by equity research analysts. We first document significant variation in how each of the forecasting styles contributes to the consensus analyst forecast accuracy and in how each incorporates public information from four sources (firm fundamentals, valuation multiples, momentum signals, and indicators from other intermediaries). We next find that, incremental to the number of analysts following the firm ("analyst coverage"), the number of unique forecasting styles ("style coverage") relates positively to both consensus forecast dispersion and accuracy. Further, our tests reveal that greater style coverage, but not analyst coverage, improves the information environment of firms as reflected in smaller earnings announcement surprises, higher earnings response coefficients, and reduced information asymmetry. Finally, we provide insight into factors that appear to drive forecasting style, such as overall market conditions, volatility, and the earnings announcement environment. Overall, our study sheds light on how unique forecasting styles contribute to the information environment of the firm.

1. Introduction

An extensive literature examines properties of analyst forecasts, such as accuracy, dispersion, and bias (e.g., Lys and Sohn 1990; Das et al. 1998; Duru and Reeb 2002). This literature primarily focuses on how specific analyst characteristics (e.g., experience, specialization), resources availability (e.g., brokerage size), and incentives (e.g., investment banking spillover) influence the forecasts. Far less attention has been given to *how* analysts construct their forecasts. The dispersion observed in analyst forecasts is striking given that analysts have access to the same set of public information and suggests the existence of various forecasting styles that perhaps consider and weight different types of public and private information. Our objective is to explore whether we can empirically identify different forecasting styles based on how the forecasts appear to incorporate publicly available information and then examine the implications of such styles for the market. Specifically, we address two empirical questions: (1) Can we identify different forecasting styles? And if so, (2) How does the presence of more unique styles impact the information environment of a firm?

Our focus on forecasting styles is motivated by extant research that provides some evidence that analysts may approach the research task in different ways. For example, Bradshaw et al. (2004) provides evidence that analysts use a variety of valuation models as part of their research, and both Ertimur et al. (2011) and Mauler (2018) suggest that analysts engage in different levels of disaggregation in developing their forecasts. Survey responses in Brown et al. (2015) indicate that analysts place different weights on common research inputs in developing their forecasts. Finally, survey respondents and interviewees in Brown et al. (2015) also disclose different levels of access to private information from managers and highlight that this private information is a key input into their forecasting process. This heterogeneity raises the possibility that different styles of forecasting likely exist *across forecasts*, both within an analyst (i.e., analyst *i* exhibits style *x* for firm *a*, but style *y* for firm or *b*) as well as *across analysts* (i.e., analyst *i* typically relies upon style *x* and analyst *j* typically employs style *y*). If such variation in forecasting style exists, then the extent to which a particular forecast contributes to a firm's information environment is unlikely to be uniform. It may also then be possible that certain styles produce more useful information than others, or that a diversity of styles provides richer overall information (Da and Huang 2019).

The central challenge with addressing this question relates to the identification of forecasting styles because we only observe the ultimate output of the forecasting process and not the manner in which the forecast was generated. To address this challenge, we use a machine-learning technique referred to as "clusterwise linear regression" (Späth 1979; hereafter CR) that empirically shuffles forecast revisions into different groups depending on how the revision appears to incorporate and weight different information.[1] It is important to note that the technique sorts forecast *revisions*, not necessarily analysts, into groups. This focus on forecasts, and not on forecasters, allows for the likely possibility that analysts vary their styles of forecasting across covered firms and over time.

To model individual analyst forecast revisions, we identify four categories of public information that analysts potentially consider when revising their forecasts: (1) firm fundamentals, (2) valuation multiples, (3) momentum-based signals, and (4) forecast herding indicators. Within each category, we identify three representative variables, and regress individual analyst revisions on these 12 variables. While we expect these factors reasonably capture public information, we also expect that forecasting styles reflect differences in analysts' access to private information

[1] We develop our own implementation of CR, which is very similar to the model outlined in Späth (1979). In different settings, Larcker and Richardson (2003) and Allen, Larson, and Sloan (2013) both employ Latent Class Mixture Models, which are similar in spirit to CR.

relevant for forecasting. Consistent with this view, in a pooled regression, the modeled parameters capturing public information explain less than one-half (roughly 40 percent) of the overall variation in forecast revisions.

Like many unsupervised machine learning techniques, CR requires a researcher to specify the number of unique groups (*k*), or in our setting, unique styles of forecasting. CR then estimates *k* separate weighted-least squares regressions, where each observation is assigned a weight for each regression such that observation weights sum to one. These regressions are iteratively estimated, and weights are re-assigned based on the *k* regression residuals; larger (smaller) residuals equate to smaller (larger) weights in the next estimation, and the process stops once weights converge. Forecasting style is then assigned based on the largest regression weight for each observation. Based on our diagnostic procedures, we focus on the five groups, or styles, that represent more than 99 percent of forecasts.

We estimate the CR using a pooled sample of observations between 2010 and 2017. For model training, we use one-quarter-ahead forecasts issued at the earnings announcement (e.g., the Q3 analyst forecast issued at the Q2 earnings announcement), which roughly holds constant forecast horizon and public information available to the analysts. The R^2 values within the five predominant styles identified by the CR average nearly 80 percent, which is more than double the explanatory power of the pooled regression model.

To address our first empirical question (i.e., can we identify different forecasting styles?), we perform two sets of analyses. First, we examine whether the CR procedure identifies forecasting styles that vary in terms of forecast accuracy. We conduct this test at the firm-quarter level and utilize firm fixed effects. Our estimates suggest that the addition of a forecast from any one of the forecasting styles within a quarter corresponds to a significant increase in consensus accuracy. However, the magnitude of the consensus accuracy forecast improvements varies considerably depending on which style is added. Adding a forecast from the most accurate style in a given quarter is four times more impactful than increasing coverage in the least accurate style. This test provides initial evidence that the CR procedure groups forecasts into meaningful styles that contribute in different ways to the quality of the consensus forecast.

Our second set of analyses examines differences in how the five forecasting styles incorporate public information about firm fundamentals, valuation multiples, momentum, and outstanding forecasts and press coverage. We plot the relative importance of these 12 factors in explaining forecast revisions within each forecasting style. The plots reveal substantial variation across styles in how public information is used. For instance, compared to the most accurate style, the least accurate forecasting style places much greater weight on information related to fundamentals and price momentum. Perhaps the most interesting observation in this plot is that the most accurate forecast style appears to rely *least* on the modeled public information. This style exhibits the lowest R^2, and coefficient estimates for this style hover around zero (though they are still statistically significant). This implies that forecasts in this group rely relatively more on unmodeled, and thus presumably private, information (consistent with Brown et al. 2015 who note that private information from managers is important to analysts in forming their forecasts). Interestingly, the second most accurate style has the highest R^2, and this is largely driven by herding behavior. Overall, we conclude from these first two sets of analyses that the CR technique is identifying unique forecasting styles based on how information is used to revise forecasts.

Next, we turn to our second empirical question (i.e., how does the presence of more unique styles impact the information environment of a firm?). To address this question, we measure "style coverage" as the number of unique forecasting styles present for a particular firm in a given quarter and "analyst coverage" as the number of analysts following the firm in the same quarter. We then test whether style coverage is incremental to analyst coverage in explaining the consensus forecast dispersion, consensus forecast accuracy, and market activity (i.e., return magnitude and spreads) at the earnings announcement. With respect to consensus dispersion, we find that analyst (style) coverage is negatively (positively) associated with dispersion. These results are expected, as it suggests that greater style coverage corresponds to a greater diversity of forecasts. More interestingly, our evidence also suggests that both analyst coverage and style coverage relate *positively* to forecast accuracy. For the median firm in our sample, adding a revision from a new style corresponds to an increase in consensus accuracy that is 44 percent larger than adding an analyst from a style already included in the consensus. This finding is consistent with prior evidence that independence among social media forecasters on Estimize yields more accurate consensus forecasts (Da and Huang 2019) and that greater cultural differences among analysts yields higher quality consensus forecasts (Merkley et al. 2020).

We next consider potential capital market effects of forecasting style with a focus on earnings announcement reactions and information asymmetry for the period being forecasted. We first examine the magnitude of earnings announcement returns. If greater style coverage enhances the information environment of the firm *during the quarter* as the forecasts are published, then we expect the information content of the earnings announcement to be lower, as the market is less surprised by the quarterly performance. Alternatively, it is possible that greater style coverage corresponds to a noisier information environment, as suggested by our results for forecast dispersion, with greater resolution at the earnings announcement and larger market reactions. Consistent with the first argument, we find that the magnitude of earnings announcement returns declines with style coverage (i.e., the number of unique styles of forecasting in a quarter). Our estimates suggest that increasing style coverage from two to three (three to four) is associated with a ten (seven) basis point reduction in absolute returns. While this suggests that the amount of new information conveyed at the earnings announcement declines with style coverage, it is also possible that the per-dollar response to earnings increases. We also find results that this is the case—adding one additional style to the consensus increases the earnings response coefficient by approximately 10 percent.

We next examine information asymmetry around earnings announcements. Prior research suggests information asymmetry temporarily, but substantially increases at earnings announcements because certain investors can more quickly process and respond to earnings news (Kim and Verrechia 1994; Lee et al. 1993; Amiram, Owens, and Rozenbaum 2016). On the one hand, if greater style coverage before the earnings announcement creates a richer information environment for all investors, then this could mitigate the information advantage of sophisticated investors. This reduced advantage would, in turn, reduce the spike in information asymmetry at the earnings announcement. On the other hand, greater style coverage could create a noisier information environment that is more difficult for less sophisticated investors to process, which could increase the spike in spreads at the earnings announcement. Our evidence is more consistent with the first argument; we find that abnormal bid-ask spreads at the earnings announcement are negatively associated with style coverage. Surprisingly, when both analyst coverage and style coverage are included in the model, we find that analyst coverage has a positive association with announcement spreads. For the median firm, adding one style reduces spreads by approximately 0.29, or 9 percent of the mean of abnormal spread. Overall, the results of these analyses suggest that there are capital markets benefits associated with a firm having a greater style coverage.

Our final set of tests are aimed at understanding whether forecasting style is driven by the specific analyst, the forecasting environment, or a combination of the two. First, we consider both the distribution of styles within an analyst (i.e., do analysts use multiple styles in a quarter?), and the persistence of forecast style for a particular analyst for a given covered firm. Overall, we find little evidence that forecasting style is an analyst specific attribute. We find that analysts typically employ multiple forecasting styles; for instance, of analysts covering three firms in a quarter, only 6.6 percent use a single style for their three revisions. More surprisingly, we find that forecasting style is not very persistent over time for a given analyst-firm combination. That is, for a given covered firm, the majority of analysts tend to adopt a variety of forecasting styles over time.[2] One implication of this result is that individual analyst accuracy for a given firm is unlikely to be persistent over time, and subsequent analyses suggest that this is indeed the case.[3]

Second, we examine the extent to which cross-sectional, time-varying factors explain forecasting style assignment. To do this, we explore the extent to which macroeconomic factors, firm-specific characteristics, analyst properties, and earnings-announcement environments predict the likelihood of being assigned in more or less accurate style. We find some evidence that more positive market sentiment and higher VIX negatively predict the likelihood of an accurate forecasting style. Firm-specific volatility also predicts style assignment—more (less) accurate styles are more likely for less (more volatile) firms. We find little evidence that analyst (All-Star) or brokerage (Brokerage Size) characteristics explain style assignment. As for properties of the

[2] We also confirm that forecasting styles are not simply capturing time period or industry. Specifically, the distribution of forecasting styles across both time and industry is relatively stable.

[3] This conclusion is based on two tests. First, we rank analysts covering a firm by accuracy in a given quarter, and regress this rank on analyst-firm fixed effects. Those fixed effects explain only 5.6% of variation in this variable. Second, we compare a regression of forecast accuracy on firm fixed effects to one using crossed firm-analyst fixed effects. The increase in R^2 from adding the fixed effects is only marginal.

earnings announcement our evidence is mixed. Earnings announcement busyness and the sign of the prior earnings announcement both exhibit non-linear patterns with style assignment.

We contribute to the growing literature within accounting that examines the benefits of diversity among groups of experts when drawing consensus. Most relevant to our setting, Merkley et al. (2020) find that cultural diversity among the analyst coverage of a firm is associated with higher forecast quality, in part, because such diversity leads to opinions that are more diverse. The focus in Merkley et al. (2020) is squarely on diversity observed at the analyst (human) level, which cannot change over time. We extend this line of research by identifying and exploring the market outcomes of differences in style observed at the forecast revision (forecast) level, which can (and does) change over time or across covered firms.

We also contribute to research examining the determinants of analyst forecast accuracy. Consistent with conventional wisdom, this literature suggests the quality of a firm's information environment is increasing in the number of analysts covering the stock. Our evidence indicates that this characterization is incomplete. Our tests reveal a more nuanced view that the relation between a firm's information environment and following depends heavily on forecasting style.

We also contribute to the literature examining various outcomes related to the "style" of different stakeholders, like investors (Cronqvist and Siegal 2015), management (Bertrand and Schoar 2003; Bamber et al. 2010; Dyreng et al. 2010; Ge et al. 2011), board members (Qu 2020), and auditors (Francis et al. 2014; Johnston and Zhang 2020). We extend this work by examining the style of analyst forecasting and finding that most analysts do not employ a single forecasting style, but rather employ multiple forecasting styles within their coverage. While beyond the scope of our study, our results open the door to research exploring *why* analysts use different forecasting styles for different firms or at different times.

Finally, we suggest a new method of identifying "style" using CR. Prior studies identify style using individual fixed effects or observable analyst specific characteristics. Existing research uses Latent Class Mixture models to explore variations in audit pricing (Larcker and Richardson 2003) and accruals (Allen, Larson, and Sloan 2013). This research focuses more on differences at the firm level. In contrast, we believe CR could be used in other settings to assign individuals, like analysts, into distinct group.

2. Prior Literature & Empirical Questions

The literature that examines sell-side analyst earnings forecasts is vast and continues to grow. Much of this literature focuses on identifying determinants associated with individual analyst forecast attributes such as its accuracy or bias. These determinants can largely be categorized into three groups, including analyst characteristics (e.g., experience, specialization), resources availability (e.g., brokerage size), and incentives (e.g., investment banking spillover). Representative papers that examine these groups of determinants include O'Brien and Bhushan (1990), Clement (1999), Jacob et al. (1999), Irvine (2004), and Corwin, Larocque, and Stegemoller (2017). This literature, however, has placed less focus on *how* analysts construct their forecasts. That is, what is the information they use to determine their forecasts of earnings? The lack of research in this area does not suggest a lack of interest in the topic or a signal of unimportance. Rather, it likely stems from data unavailability, as it is difficult for researchers to observe the forecasting process of analysts. Brown et al. (2015) attempt to pierce this black box by asking analysts how useful different types of information sources are for determining their earnings forecasts. Survey results suggest that analysts' industry knowledge, private communication, earnings calls, management guidance, and accounting reports are most useful, though analysts are not asked to identify specific pieces of information used to forecast.

While much of this literature focuses on individual analyst forecasts, a relatively smaller sub-stream focuses on the *consensus* earnings forecast (for reviews of this literature, see Schipper (1991), Ramnath et al. (2008), and Bradshaw (2011)).[4] Since Brown et al. (1987) revealed that the consensus analyst forecast is superior to the time-series forecast, the consensus forecast has been commonly used as a proxy by practitioners and researchers for general market expectations of earnings performance. Indeed, some argue that the consensus forecast is among the most widely used financial metrics in capital markets (Graham et al. 2005; Chang et al. 2021; Merkley et al. 2020).

Given the importance of the consensus analyst forecast to markets, this sub-stream of research seeks to identify the key determinants of consensus forecast accuracy. For instance, Lys and Soo (1995) argue that the number of analysts covering a stock is an indicator of competition among analysts and that increased competition will motivate analysts to increase the precision of their forecasts by increasing spending and effort on research. Consistent with their prediction, they find that analyst earnings forecast accuracy is increasing in the number of analysts covering the stock. Duru and Reeb (2002) provide similar evidence of a positive association between consensus accuracy and analyst coverage.

While there is considerable focus on accuracy in the literature, other research considers the extent to which forecasts in the consensus disagree with one another, which is typically referred to as forecast dispersion. Theory shows that dispersion is a product of uncertainty and divergence in beliefs (Barry and Jennings 1992; Abarbanell, Lanen, and Verrecchia 1995; Barron, Kim, Lim,

[4] Some prior research suggests that analyst forecast accuracy is not of primary importance to analysts because they are not incentivized to forecast accurately. Groysberg et al. (2011) fail to provide any evidence from one investment bank that analyst forecast accuracy is significantly related to the analyst's compensation. Brown et al. (2015) provide survey evidence that analysts rank the accuracy of their forecasts as the least important factor (of nine potential factors) of their compensation. These findings notwithstanding, the consensus analyst forecast continues to be one of most cited financial metrics of a company. Further, our empirical tests also examine other market effects, in addition to consensus forecast accuracy.

and Stevens 1998), and research commonly uses dispersion as a proxy for these constructs (Ajinkya et al. 1991). High dispersion has been linked to capital market consequences. For instance, Diether et al. (2002) find that stocks with higher dispersion earn lower future returns and interpret their evidence as suggesting that a high level of disagreement can result in overvaluation. Johnson (2004) suggests this result is driven by unpriced information risk.

Dispersion arises from differences in forecasts, of which one possible source is the *uniqueness* of the analysts contributing to the consensus. Merkley et al. (2020) examine the cultural diversity within the sell-side analyst following and find that it is positively related to the accuracy of the consensus. Additional tests suggest the consensus improvement stems, in part, from diversity in forecast errors. This diversity reduces forecast error covariance, which improves forecast quality. Similarly, Da and Huang (2019) find that the consensus earnings forecast crowdsourced from social media analysts on Estimize.com is more accurate when the forecasts are more independent, exhibiting less "herding" behavior. The authors conclude that the wisdom of crowds is best harnessed by the presence of more independent voices. Somewhat surprisingly, the literature has largely stopped exploring other determinants of consensus forecast accuracy in favor of researching factors associated with individual analysts' forecast accuracy and bias, such as analyst experience, analyst effort, resource availability, and brokerage incentives (Clement 1999; Jacob, Lys, and Neale 1999). The focus on individual forecasts makes sense given the richness of the underlying variation available to researchers. However, we argue that the consensus earnings forecast remains a very important metric in capital markets. It is generally more accurate than the underlying individual forecasts (Clement 1999; Zarnowitz and Braun 1993) and is commonly reported by data aggregators (e.g., Yahoo Finance) "for free" to investors. Thus, it is a particularly important metric for smaller retail investors who do not subscribe to data providers.

The intersection of the literatures discussed above motivates our primary research objective which is to first identify forecasting styles and then to examine its market implications. We distill this research objective down by addressing the following specific empirical questions:

EQ1: Can we identify unique forecasting styles?

EQ2: Does style coverage improve the information environment of the firm?

3. Sample, Data and Research Design

3.1 Sample Information & Revision Regression

We begin with a sample of quarterly earnings forecasts in IBES issued by analysts in the days following earnings announcements. We focus on the days following earnings announcements because the majority of analysts following a firm revise their forecasts immediately following firms' earnings announcements, and earnings announcements are important information events that should be relevant to all forecasters. We examine one-quarter-ahead earnings forecasts (FPI=6 in IBES) to hold the forecast horizon relatively constant. In other words, following the earnings announcement for a given firm, nearly all analysts will revise forecasts for the next quarter, and the horizon of those forecasts is roughly similar at approximately one quarter. We consider a forecast to be issued at the earnings announcement if it occurs within 5 days of the earnings announcement date (days 0 to +5, inclusive). In total, we identify 370,771 forecasts issued by 3,399 unique analysts at earnings announcements for 2,666 unique firms between 2010 and 2017.[5]

Our measure of forecasting style requires us to model observable information an analyst potentially uses when revising their forecast. While we cannot include factors in the model that

[5] We end our sample in 2017 because, as of October 2018, Thomson-Reuters re-assigned a significant portion of the broker and analyst identifiers in IBES (see Thompson-Reuters "Product Change Notification" documentation on WRDS). WRDS explains that individual analyst identifiers have been and will continue to be subject to reshuffle without warning, and that some analyst identifiers will be anonymized. This makes it difficult to track all individual analysts over time and link them across IBES datasets.

capture unobservable private information, we include a broad set of public factors to increase the likelihood that modeled variation captures a significant portion of information common to analysts. We rely on prior research (e.g., Stickel 1990; Drake, Rees, and Swanson 2011) to identify publicly available signals that are likely considered by analysts. Similar in spirit to Drake et al. (2011), we organize these factors in categories of similar constructs in the following regression model:

$$Revision_{i,j,t} = a_{0,t} + \beta_t Fundamentals_{i,t} + \gamma_t Momentum_{i,t} + \delta_t Valuation\ Multiples_{i,t} + \zeta_t Herding_{i,t} + e \quad (1)$$

Revision refers to analyst *j*'s revision for firm *i*'s forecast at time *t*, defined as the analyst's current forecast minus prior forecast scaled by price, and then multiplied by 100 for exposition. *Fundamentals, Momentum, Valuation Multiples,* and *Herding* refer to groups of proxies related to each category, which largely follow Drake et al. (2011) and Stickel (1990). For each category, we identify three proxies.[6] *Fundamentals* includes factors pertaining to changes in key financial statement variables. Specifically, we include the change in sales (*ΔSales*), change in operating cash flows (*ΔOCF*), and change in capital expenditures (*ΔCapEx*).[7] *Momentum* includes variables capturing the trajectory of firm performance. We measure momentum based on investor responses to recent earnings surprises (*EAReturns_Prior4*), buy-and-hold returns since the analyst's last forecast (*ReturnSinceLast*), and the number of recent earnings increases (*IncomeInc*). *Valuation Multiples* captures metrics typically used by analysts and investors to assess relative valuations. Specifically, we include the book-to-market ratio (*BM*), the cash-flow-to-price ratio (*CFP*), and the sales-to-price ratio *(SP).*[8] Finally, *Herding* includes indicators of possible "herding behavior", or the tendency of an analyst to follow news from other intermediaries in constructing his or her

[6] Appendix A includes detailed variable definitions.
[7] We do not include change in earnings performance in *Fundamentals* for two reasons. First, the factors we identify are relevant for predicting future performance for nearly all firms whereas earnings is less relevant for growth firms or in periods with losses. Second, we include the trajectory of earnings-based performance in *Momentum*.
[8] We use *SP* instead of the price-to-earnings ratio to avoid issues with zero-earnings or loss firms.

revision. We include two measures derived from analysts (Stickel 1990): the change in consensus (*ΔConsensus*) and the change in long-term growth forecasts (*ΔLTG*) since the analyst's prior forecast. We also include one measure derived from the business press—the average sentiment of business press articles published during the quarter about the firm from RavenPack (*BusPress*).

3.2 Cluster Regression (CR)

To model forecast style, we use CR to identify discrete groups that maximize the fit of equation (1). CR essentially fits a least squares regression where the sample is partitioned into *k* groups, but group assignment is unknown *a priori* (Späth 1979). This empirical problem can be represented by the following equation:

$$Y_i = \sum_{k=1}^{m}(\alpha_k + \beta_{1,k}x_{1,i} + \beta_{2,k}x_{2,i} + \ldots + e_i) \tag{2}$$

where *k* represents group (or cluster) membership (i.e., observation *i* is assigned to one group, *k*, where $k \in \{1,\text{m}\}$). In Späth's original algorithm, observations are iteratively re-assigned values for *k* until the model converges.

Canned implementations of CR are not available in more recent programming languages, so we develop our own implementation in Python.[9] Our procedure emulates the spirit of equation (2), but with a few modifications to improve performance. We first describe the basic, iterative procedure and then provide details on the specific process we use. First, in each iteration, we estimate *k* separate regressions and examine each observation's fit within each regression. Second, to aid with convergence, we use weighted least squares regression where each observation is assigned *k* different weights that must sum to one. In machine learning parlance, this is referred to as "softmax" or "soft responsibility" (e.g., Kawano et al. 2013), which contrasts with discrete

[9] We largely base our implementation on code provided here: https://stackoverflow.com/questions/39208679/libraries-for-regression-clustering-in-python. We modify the code to stabilize initialization, provide additional diagnostic information, and allow for out-of-sample prediction.

responsibilities where each observation receives a weight of zero or one for each cluster (as in equation (2)). Third, we employ Ridge regression instead of OLS. Ridge regression introduces a regularization factor that will down-weight overly influential features and is useful when regressors are highly correlated. While not a concern in our overall sample, it could become a concern within individual clusters. Fourth, rather than randomly assigning initial weights in our CR procedure, we initialize *k* weights for each observation using K-means clustering. Our specific implementation follows these steps:

(1) We estimate a pooled version of equation (1) and apply k-means clustering to regression residuals.[10] K-means clustering identifies natural centroids (or cluster centers) based on the magnitude of the regression residuals, and we assign initial weights based on distances to the *k* centroids identified in the K-means procedure. We then re-scale the weights such that they scale to one.[11]

(2) We estimate equation (1) *k* times using a weighted Ridge regression with the initial K-Means derived weights. We compute the overall mean-squared-error (MSE) from this estimation (adjusted for regression weights) and compute the *k* residuals per observation.

(3) We compute new weights (or responsibilities) using the following formula:

$$Resp_{i,k} = exp(-e_{i,k}^2 / MSE) \tag{3}$$

$Resp_{i,k}$ is observation *i*'s regression weight for cluster *k*, $e_{i,k}$ is observation *i*'s residual from the cluster *k* regression, and MSE is described previously. This formula assigns greater weights to smaller residuals. As before, these responsibilities are re-weighted such that they sum to 1.

Steps (2) and (3) are repeated until there is little change in responsibilities.[12] Upon convergence, we assign observations to the cluster corresponding to its largest responsibility.

---

[10] We find that using randomly assigned initial weights leads to variability in final cluster assignment. Using k-means to initialize weights greatly reduces this variability.
[11] To illustrate, suppose k = 3, and the three cluster centroids were -1, 0, and 1. If the residual from the pooled regression for a given observation were 0.25, we would compute three distances to centroids of 1.25, 0.25, and 0.75, respectively. These distances are then rescaled such that they sum to 1 (i.e., 0.56, 0.11, and 0.33).
[12] We use a tolerance of 0.00001 to define "little change".

Like many unsupervised machine-learning methods, CR requires the researcher to specify the number of clusters (*k*). We evaluate model fit for between 2 and 20 clusters and employ a hold-out sample to avoid overfitting. We consider two metrics to evaluate fit: the R^2 value of estimation and the average "confidence" of cluster assignment (or the average weights for assigned clusters) relative to unconditional assignments.[13] We plot these diagnostics in Figure 1. We note that, in a pooled regression, equation (1) has an R^2 of approximately 40 percent. This number nearly doubles around *k=5*, suggesting substantial variation across forecasting styles. Confidence appears to peak between four and six clusters and then declines at higher values. Based on these diagnostics, we choose to use six clusters.[14]

We note that the CR procedure does not impose balanced clusters, and the final cluster assignments are imbalanced. One cluster in particular contains less than one percent of all revisions. We find that this cluster is largely comprised of outlier revisions; the mean revision in this style is nearly six times stock price. Therefore, we exclude it from our tests and focus on the five other forecasting styles, each of which contains at least 10 percent of all revisions. We label the smallest (largest) remaining forecasting style *Style 1 (Style 5*).

One potential issue with the unsupervised nature of CR is that it simply groups observations based on other, unmodeled factors that affect the relevance of information, such as industry trend. To assess whether this is the case, we present descriptive statistics that examine style assignment by industry in Table 1, where we present a breakdown of forecasting style by Fama-French 12-industry classifications. As noted by the bottom row, approximately 10 percent of the sample falls in the least populated style, and 31 percent in the largest. In each cell in the table, we first present

[13] We define confidence as the mean cluster fit minus 1/k. For instance, if, for *k=4*, the average cluster fit was 75%, then confidence equals 50% (75% - 25%).
[14] When choosing *k*, our intent is to balance model fit with generating reasonable variability in the number of forecast styles. Our main inferences are robust to using values of 4 or 5 for *k*.

the number of observations in that industry-style combination. In parentheses, we also report the percentage of each style corresponding to each industry (i.e., the column percentages sum to 100 percent). While we observe some small fluctuations, these data suggest that CR does not merely partition by industry membership. We conduct additional tests to better understand the potential drivers of forecasting styles when we address EQ1.

3.3 Descriptive Statistics

In Table 2, we present descriptive statistics for variables discussed previously, as well as several defined later. In Panel A of Table 2, we provide descriptive statistics for the variables used to identify forecasting styles, and in Panel B, we present statistics for variables used to examine the capital market consequences of forecasting style. Consistent with prior research, we find that *Revision* and *ΔConsensus* are both slightly negative, consistent with analyst forecast "walk-down" (Richardson et al. 2004). We observe small average declines (increases) in operating cash flows (capital expenditures). Both returns measures (*EAReturns_Prior4* and *ReturnSinceLast*) have means close to zero. Median media sentiment (*BusPress*) is close to neutral (Ravenpack codes sentiment on a 0 to 100 scale, where 50 is neutral), though the mean is slightly negative. Finally, valuation multiples (*BM, CFP, SP*) are in line with prior research. For example, we find a mean book-to-market ratio is 0.437, which is comparable to that reported in Drake et al. (2011).

4. Empirical Design & Results

4.1 EQ1: Can we identify forecast style?

To address EQ1, we perform two sets of analyses. First, we consider whether the relations between analyst following and consensus forecast accuracy vary by style. Second, we examine variation in how different forecasting styles use publicly available information.

4.1.1 Unique forecasting styles and consensus analyst forecast accuracy

Our first analysis focuses on whether consensus analyst forecast accuracy relates to the earnings revisions we use to identify forecasting styles with CR. Research generally assumes that a larger following corresponds to a more accurate consensus forecast (Lys and Soo 1995; Duru and Reeb 2002; Merkley et al. 2020). The objective of this test is to examine whether forecasting style also helps explain consensus forecast accuracy. Because there is only one "right" forecast, we argue that different styles should exhibit different levels of accuracy, which will impact the quality of the overall consensus. More specifically, to answer EQ1, we first test whether the styles we identify group forecasts in a way that leads to different associations with consensus forecast accuracy using the following model:

$$Accuracy_{i,t} = \alpha + \Sigma\beta_k StyleFollowing_{i,t,k} + \Sigma Controls_{i,t} + \Sigma FirmFE + e_{i,t} \quad (4)$$

We estimate equation (4) at the firm-quarter level. *Accuracy* captures the consensus forecast accuracy. $StyleFollowing_{i,t,k}$ refers to a vector of variables, indexed with k, each of which equals the natural log of one plus the number of forecasts for firm i's one-quarter ahead earnings identified as belonging to style k. In essence, we take the traditional measure of analyst following—the number of unique forecasters for a given firm in a given quarter—and decompose it by style. We estimate equation (4) with and without a vector of firm controls (*Controls*), which are based on prior research (Merkley et al. 2020). Specifically, we control for firm size (*Size*), the book-to-market ratio (*BM*), return on assets (*ROA*), volatility of performance (*StdROA*) returns *(Ret*), and volatility of returns (*StdRet*). We include firm and time (year-quarter) fixed effects and calculate t-statistics based on standard errors clustered by firm.

We present the equation (4) results in Table 3. In column 1 (2) of the table, we provide the results without (with) control variables. Our evidence indicates that important differences exist in how following by style contributes to consensus accuracy. Recall that the label *StyleFollowing* uses the same convention described previously, such that $StyleFollowing_1$ ($StyleFollowing_5$) corresponds to *Style 1 (Style 5*). Interestingly, this assignment, which was based on the size of each style cluster, also ranks styles by accuracy. Estimates in column 1 suggest that all but the least accurate style (i.e., style 1) contribute to consensus accuracy. In column 2, we introduce controls, and find that all estimates in all five styles correspond to increased forecast accuracy. However, the magnitude of these effects varies greatly. The coefficient on $StyleFollowing_5$ is more than four times as large as that on $StyleFollowing_1$ (difference significant at p-value < 0.01). In sum, the evidence presented in Table 3 suggests that the CR procedure identifies important differences in forecasting style, which translate to different levels of consensus accuracy.

4.1.2 Use of Publicly Available Information

We further examine whether we can identify differences in forecasting style by considering variation in how different styles use publicly available information. We begin by revisiting equation (1) and, in Figure 2, examining plots of the 12 coefficients included in the model by forecasting style. Consistent with earlier tests, the *x*-axis corresponds to style assignments and, based on results in Table 3, is increasing with style accuracy. To facilitate magnitude and trend comparisons, all variables (dependent and independent) in these regressions are standardized to have mean 0 and standard deviation 1. Note that the y-axis varies with the scale of the coefficients, so we include a red horizontal line at a standardized coefficient magnitude of zero.

Overall, in Figure 2 we observe significant variation in coefficients across styles. This suggests that different forecasting styles place different weights on observable signals. Beginning with *Fundamentals* (top row), the least accurate style (far left in each sub-figure) weights these factors heavily. Moving from left to right, the magnitude of coefficients rapidly drop to close to zero, indicating that the more accurate styles place little weight on fundamental signals. More specifically, the most accurate style exhibits small, positive coefficients for *ΔSales* and *ΔOCF* and a slightly negative coefficient for *ΔCapEx*. We observe similar patterns with the *Momentum* signals. The least accurate style appears to weight these measures at levels that are considerably higher than more accurate forecasting styles. For *Valuation Multiples* we observe less consistent patterns with respect to the least accurate style. Specifically, coefficients on *BM* and *SP* are both strongly negative, while the coefficient on *CFP* is significantly positive. This pattern reverses for the second least accurate style and then largely stabilizes for the remaining styles (*CFP* is significantly positive for the third style and then drops close to zero). Once again, the magnitude of coefficients for the most accurate style are generally the smallest. Finally, for *Herding*, the least accurate style again seems to heavily weight all three factors. Interestingly, the third style appears contrarian, as each of the three herding measures exhibit negative associations with the revision.

As discussed, the most accurate style appears to weight the modeled information the least, as evidenced by the smallest coefficient magnitudes. This suggests that the most accurate forecasting style is primarily relying on unmodeled, and potentially private, information. This inference is further supported by the R^2 for the most accurate style, approximately 71 percent (untabulated), which is the lowest across the five clusters. Interestingly, the second most accurate cluster has the highest R^2 (96 percent). Much of this explanatory power is driven by *ΔConsensus*, which has a coefficient nearly equal to one and a t-statistic of over 300. In other words, this forecasting style appears to mimic other analyst revisions with little adjustment. In sum, this figure provides additional evidence of differences across forecasting styles. All 12 variables we model exhibit significant variation. In fact, all 12 have at least one positive and one negative coefficient.

The evidence presented in Table 3 and Figure 2 suggests the answer to *EQ1* is clearly "yes." Based only on differences in how revisions appear to weight the information we model in equation (1), our CR procedure identifies meaningful differences in forecasting styles. Note that this evidence is exploratory in that we made no *ex ante* prediction as to how each style following would relate to these outcomes. In the next sections, we consider whether a diversity of forecasting styles for a given firm provides incremental benefits to the information environment.

4.2 EQ2: Does style coverage improve the information environment of the firm?

Our second empirical question asks whether diversity of styles enhances or weakens the information environment. Specifically, we examine how analyst coverage from a larger *number of unique forecasting styles*, a construct we refer to as "style coverage", relates to two aspects of the analyst forecast consensus: consensus dispersion and accuracy. Then, we examine how it relates to two capital market outcomes: subsequent earnings announcement returns and subsequent earnings announcement information asymmetry.

4.2.1 Unique Forecasting Styles and Consensus Forecast Dispersion and Accuracy

A key motivation for our study is the notion that analyst forecasts are derived from different styles of forecasting, suggesting they may influence the consensus forecast in unique ways. To examine this, our next tests consider how a diversity of forecasting styles relate and contribute to the overall attributes of the consensus forecast, as measured by its dispersion and accuracy.

With respect to consensus forecast dispersion, on the one hand, prior work suggests higher levels of analyst following is associated with lower dispersion (e.g., Ajinkya et al. 1991). On the other hand, our earlier evidence suggests different forecasting styles exhibit different associations with accuracy. If adding styles to a consensus adds less accurate forecasts, then we will likely find higher levels of analyst forecast dispersion when more unique styles contribute to the consensus.

With respect to consensus forecast accuracy, in Table 3 we observe that there are significant differences in how forecasts from various styles relate to forecast accuracy. It could be that when more than one style contributes to the consensus, then the consensus incorporates a greater diversity of views or more unique perspectives of the firm. Prior work suggests more independent thought is associated with higher quality forecasts. For example, Da and Huang (2019) provide evidence that requiring "blind" forecasts on Estimize (asking users to submit an estimate before viewing the current consensus) promotes independence and enhances forecast accuracy. If different styles similarly reflect independence across analysts, then forecast accuracy may improve. However, it could also be the case that adding more forecasting styles to the consensus simply increases (decreases) consensus forecast accuracy based on whether any new forecasting style that is introduced is generally more (less) accurate than the existing forecasting styles including in the consensus. Thus, it is an open empirical question whether the addition of more unique styles to the consensus will lead to a higher or lower quality consensus forecast.

To test these, we estimate the following model:

$$Dispersion_{i,t}/Accuracy_{i,t} = \alpha + \beta_1 Analyst\ Coverage + \beta_2 Style\ Coverage + \Sigma Controls_{i,t} + \Sigma FirmFE + e_i \quad (5)$$

We estimate equation (5) at the firm level. The dependent variable and *Controls* are the same as in equation (4). *Style Coverage* equals the natural log of one plus the number of unique styles contributing to a given consensus, and *Analyst Coverage* equals the natural log of one plus the number of estimates. In essence, this model horseraces style following (or the number of distinct styles contributing to a consensus) with the traditional measure of analyst following (the number of estimates included in a consensus).

We present the equation (5) estimation results using *Dispersion* as the dependent variable in Table 4 Panel A. Note that we define *Analyst Coverage* in a manner consistent with prior research (the natural log of one plus the number of estimates in the consensus), which assumes a concave association with accuracy (i.e., moving from 3 to 4 analysts is more impactful than moving from 13 to 14). To facilitate comparisons, we define *Style Coverage* similarly, though it is unclear whether concavity should be assumed. Therefore, we repeat estimation of equation (5) using unlogged versions of *Analyst Coverage* and *Style Coverage* and present the results in columns 3 and 4 of Table 4. Consistent with prior research, we find a negative and significant coefficient on *Analyst Coverage,* suggesting higher levels of analyst coverage is associated with lower forecast dispersion. However, in columns 2 and 4 we find a *positive* and significant coefficient on *Style Coverage,* suggesting that having more unique forecasting styles in the consensus increases the dispersion of the consensus analyst forecast. This result is consistent with analyst forecasting styles truly offering uniquely different perspectives rather than simply herding together to a single forecast number. This result also suggests that we are capturing unique forecasting styles rather than simply higher levels of analyst coverage.

We present the equation (5) estimation results using *Accuracy* as the dependent variable in Table 4 Panel B. Consistent with multiple viewpoints enhancing accuracy, our evidence suggests that style coverage is incremental to analyst coverage in explaining consensus forecast quality. For the median firm in our sample (following = 7, styles = 3), adding a forecast from a new style corresponds to an increase in accuracy of 0.055, approximately 12 percent of the mean value of *Accuracy*. This effect is 41 percent larger than the effect of adding another analyst from an already included forecasting style.[15] While forcing *Analyst Coverage* to have a linear association with

[15] Moving from 7 to 8 analysts corresponds to an increase in *Analyst Coverage* of ln(8/7), or 0.13. Multiplying this value by the coefficient on *Analyst Coverage* (0.28) equals 0.037. Moving from 3 to 4 styles corresponds to an increase in *Style Coverage* of

Accuracy likely understates its true effect, we believe this evidence further suggests that the diversity of forecasting styles is an important determinant of overall accuracy.

4.2.2 Unique Forecasting Styles and the Information Content of Earnings

To further gauge the capital market consequences associated with style coverage, we examine the information content of the earnings announcement. If firms with more forecasting styles have better information environments, then this should sharpen investors' expectations for the upcoming earnings announcement, thereby limiting the likelihood of large earnings announcement surprises. Accordingly, we test whether the information content of the earnings announcement (|*EAReturns*|) varies with the quantity of unique forecasting styles contributing to the consensus for that earnings announcement. We do this by replacing *Accuracy* with |*EAReturn*| in equation (5). Note that we do not include fixed effects in this regression because the dependent variable is a short-window return reaction to a news announcement. We also cluster standard errors by date instead of firm since cross-sectional correlation in error terms is a bigger concern than serial correlation when returns is the dependent variable (Petersen 2009).

We present the equation (5) estimation results using |*EAReturn*| as the dependent variable in Table 5, Panel A. In column 1 of Panel A, we report the results excluding *Style Coverage*. Perhaps surprisingly, we observe a positive association between |*EAReturn*| and *Analyst Coverage*. In column 2, we include both variables, and we find that *Style Coverage* has a negative and significant coefficient (*t*-statistic = -3.32), opposite that of *Analyst Coverage.* In other words, the strength of the pre-earnings announcement information environment increases with the diversity in analyst styles, but not overall analyst coverage. To facilitate economic interpretation, we repeat these tests in column 3 and 4 after including only *Style Coverage* as a discrete count (instead of a

ln(4/3), or 0.29. Multiplying this by the coefficient on *Style Coverage* (0.053) equals 0.015. Summing 0.037 and 0.015 equals 0.052, which is 41 percent higher than 0.037.

logged measure). The results suggest that adding an additional style to a firm's coverage corresponds to a 9.5% reduction in |*EAReturn*| as a percentage of the mean.

While we find that the magnitude of the overall earnings price response declines with greater style coverage, it is still possible to find a stronger response *per dollar* of earnings news for firms with greater style coverage. To test this possibility, we evaluate whether the ERC varies with the number of unique forecasting styles included in the consensus forecast. If firms with more unique styles have a higher quality earnings expectation metric, then this should increase investors' ability to interpret deviations from expectations, thereby increasing the firm's ERC. To test this conjecture, we adjust our approach in Panel A of Table 5 in two ways. First, we change the dependent variable to signed abnormal returns around the earnings announcement (*EAReturn*). Second, we include the firm's earnings surprise (*EarnSurp*) and interact this variable with all other variables in the model. Our coefficient of interest is on *EarnSurp * Style Coverage*.

We present the ERC results in Panel B of Table 5. In column 1, we again report the results excluding *Style Coverage*. We find an insignificant coefficient on *EarnSurp * Analyst Coverage,* suggesting higher levels of analyst coverage do not impact ERCs. In column 2, we include both variables, and we find a positive and significant coefficient on *EarnSurp * Style Coverage* (t-statistic = 2.09). This result suggests that investors are more sensitive to earnings announcement news surprises when the consensus is made up of more unique analyst forecasting styles. Similar to Panel A, we repeat our tests in columns 3 and 4 using a discrete count of analyst and style coverage (instead of logged measures). Focusing on column 4, this result suggests that adding an additional style to a firm's coverage corresponds to a 10.4 percent increase in the earnings response coefficient (i.e., 0.033/0.316).

The results from Table 5 suggest that having more unique forecasting styles in the consensus sharpens investors' expectations of upcoming earnings releases such that investors are less surprised at earnings announcements and better able to process earnings news.

4.2.3 Unique Forecasting Styles and Earnings Announcement Information Asymmetry

In addition to the amount of information released at the earnings announcement, we consider the spike in information asymmetry that generally occurs at the announcement. Earnings announcements precipitate significant information flows that allow more sophisticated investors to gain an information advantage (Kim and Verrecchia 1994; Lee et al. 1993). Amiram et al. (2016) suggest that analyst forecasts reduce information asymmetry upon their issuance, which they interpret to mean that the forecasts are more useful to otherwise less informed investors. However, extant research finds little evidence that analyst coverage mitigates *earnings announcement* information asymmetry (e.g., Yohn 1998; Gomez et al. 2022).[16] One potential reason for this lack of evidence is that not all forecasts are equal, as suggested by our earlier tests.

We present the equation (5) estimation results using *AbSpread* as the dependent variable in Table 6, where we use the same table structure as Tables 4 and 5. In column 1 of Table 6, we observe a marginally significant and positive association between *Analyst Coverage* and *AbSpread*, which is somewhat consistent with findings in Yohn (1998). In column 2, we include both variables and continue to observe a significantly positive significant coefficient on *Analyst Coverage* that increases substantially in magnitude. Conversely, we find that the coefficient on *Style Coverage* is *negative* and significant, suggesting more unique styles decreases the spike in earnings announcement information asymmetry. We emphasize that it is difficult to measure the economic significance of *Style Coverage* holding all else constant (i.e., increasing *Style Coverage*

[16] Yohn (1998) actually estimates a positive, though insignificant (*t*-statistic=1.50) relation between analyst following and the spike in information asymmetry at the earnings announcement.

requires an increase in *Analyst Coverage*). Therefore, we again present unlogged versions of our following variables in columns 3 and 4. The estimation results suggest that adding an additional style to a firm's coverage corresponds to a 0.640 basis point reduction in *AbSpread*, or 21 percent of the mean. Additionally, the reduction in spreads attributable to increasing style coverage (0.640) more than offsets the increase in spreads corresponding to the addition of an analyst (0.156).

4.3 Summary

Tables 4 through 6 suggest CR identifies meaningful differences in styles that are associated with capital market ramifications. Further, the results support two broad inferences. First, not all forecasts are equal. While this inference may not appear surprising or new to the literature, we highlight two important considerations. Namely, our evidence in Table 3 suggests substantial differences in terms of both quality (accuracy), and our identification of style relies only on how revisions relate to factors likely relevant for forecasting earnings. Second, despite the differences in forecast quality across styles, the evidence in Tables 4 through 6 indicates that a greater diversity of styles enhances a firm's information environment. In quarters with greater style coverage, firms experience more accurate consensus forecasts, smaller earnings announcement returns, and reduced spikes in information asymmetry at the earnings announcement.

5. Additional Analysis - Properties of forecasting style

In this section, we consider the extent to which forecasting style is attributable to the analyst. We then examine other time-varying factors that potentially contribute to forecasting style.

5.1 Individual Analysts and Forecasting Style

Prior research suggests that certain analyst attributes can drive the quality of their forecasts (e.g., Merkley et al. 2020). However, the procedure we use to identify forecasting styles allows analysts to use different styles across their portfolio of covered firms. In other words, the CR procedure does not force an analyst into one style across coverage or time. We believe this is an important empirical choice because, in practice, the same analyst can change their forecast approach across firms and over time. In this section, we evaluate the extent to which this happens.

We begin by considering whether analysts tend to use the same, or different styles, for the firms they cover. In Table 7, we cross-tabulate the number of firms followed by an analyst (rows) with the number of styles employed (columns). We collapse these data at the analyst-calendar quarter level, yielding 56,890 unique analyst-firm-quarter observations. In each cell, we present two values. The first value represents the proportion of the sample in that cell. The second value in parentheses is the percentage of that row falling in that column. To illustrate, 4.3 percent of the analysts in our sample follow 3 firms and use two different styles. For analysts following 3 firms, 53.1 percent employ two different styles. While subjective, we view the evidence in Table 7 as inconsistent with analysts exhibiting a single style. For instance, for analysts covering 2 firms (9.4 percent of the sample), 77.1 percent use two styles. For analysts covering more than 3 firms, only 3.0 percent employ a single style for all firms. In sum, forecasting style does not appear to reflect an analyst-specific attribute.

While this evidence suggests substantial variation *within* an analyst's coverage, analysts may use the same style for the same firm over time. To evaluate the extent to which analysts change forecasting styles for a given firm, we tabulate the proportion of quarter-over-quarter changes by style in Table 8. We tabulate these by style since the number of observations in each style is not balanced. We also include the proportion of the sample in each style and the unconditional likelihood of changing (i.e., 1 minus the sample proportion). For all five styles, change rates are less than unconditional rates, suggesting some persistence. However, we note that the difference is not that large; between 65 and 83 percent of analysts change styles for a given firm quarter-over-quarter, which is only slightly lower than unconditional rates. Overall, this evidence suggests that forecasting style is unlikely to be a manifestation of specific analyst attributes (i.e., experience or specialization), as analysts employ different styles across their portfolio of firms and exhibit little persistence with respect to style for a given firm.

While the results in Tables 7 and 8 may be somewhat surprising, we also find that individual analyst accuracy for a given firm is also not very persistent. This inference is supported by two tests (both untabulated). In the first test, we rank analysts covering a firm by accuracy, and regress this rank on analyst-firm fixed effects. We find that the analyst-firm fixed effects explain only 5.6% of the variation in this ranked accuracy variable. In the second test, we compare a regression of forecast accuracy on firm fixed effects to one using crossed firm-analyst fixed effects. The increase in R^2 from crossing the analyst and firm fixed effects is marginal (it increases the R^2 by only 1.4%). This lack of persistence in analyst forecast accuracy could be explained by the lack of incentives to be accurate as described in Groysberg et al. (2011) and Brown et al. (2015).

5.2 Determinants of Forecasting Style

We next consider potential determinants of forecasting styles. We consider factors related to macroeconomics, volatility and performance of the firm, the forecasting analyst, and the earnings announcement of the firm. To perform this analysis, we estimate a multinomial logistic regression that attempts to predict classification into a given style. Multinomial logistic regressions produce "*k*-1" coefficients for each regressor, where *k* is the number of unique classifications; in our case, we have five styles, so 4 sets of coefficients. To facilitate discussion, we refer to styles in terms of relative accuracy, and we set the middle style, Style 3, as our baseline comparison group, so each set of coefficients captures the likelihood of moving away from this baseline style. Recall that Style 1 (Style 5) corresponds to analysts with the least (greatest) forecast accuracy. We present the results in Table 9. In column 1, we report coefficients that predict the likelihood of a forecast being classified as *Style 1* (the least accurate forecast) relative to *Style 3*. Similarly, in columns 2, 3, and 4, we report coefficients pertaining to *Style 2*, *4*, and *5*, respectively.

The first set of variables we consider relate to macroeconomic conditions. Namely, we include expected volatility (VIX) in the 30 days before the analyst forecast (*VIX_prior30days*) and a measure of market sentiment in the month before the analyst forecast (*MktSentiment*, using the data from Baker and Wurgler 2006). For VIX, we observe 3 negative coefficients, implying higher volatility reduces the likelihood of classification into *Style 2*, *Style 4*, and *Style 5*. We find that analysts are more likely to be classified in the least accurate style when sentiment is more positive (column 1) and a more accurate style (columns 3 and 4) when market sentiment is more negative. These results suggest that negative sentiment and lower volatility both predict a more accurate style classification.

Next, we consider factors related to the firm covered by the analyst, including return volatility over the prior 12 months (*StdRet*), prior accounting performance volatility (*StdROA*), and stock performance over the prior year (*Ret*). High volatility likely captures events that increase processing costs related to assimilating information, as events that are surprising or difficult to interpret correspond to more volatile and/or larger magnitude returns. For both measures of volatility, we find that Style 1 (Style 5) is more likely when volatility is greater (smaller). Specifically, the coefficients on both *StdRet* and *StdROA* decrease monotonically as we move from Style 1 to Style 5. With respect to stock performance (*Ret*), our results suggest that analysts are less (more) likely to be classified in the least (a more) accurate style when prior stock performance is better. These results are intuitive, as volatility makes forecasting difficult, and good news is generally easier to process than bad news.

Next, we turn to factors related to the forecasting analyst. We include brokerage size (*(Ln)BrokerageSize*) and an indicator equal to one if the analyst was classified as an All-star in the prior year (*AllStar*). We find no consistent patterns of results for these measures. We find some evidence that more accurate clusters are associated with smaller brokerage houses (i.e., a negative coefficient in column 4) and less likely to be All-Stars (i.e., a negative coefficient in column 3), but estimates in other columns are insignificant. These results are consistent with our classification technique identifying factors related to analyst forecast style that are not easily observable.

We also examine three factors related to firms' earnings announcements. First, we consider the earnings announcement busy season (*Busyness*). Our rationale is that analysts may revert to simple heuristics or herding on busy earnings announcement days, which prior research suggests is associated with increased processing costs (Hirshliefer et al. 2008; Blankespoor et al. 2020). *Busyness* is an indicator variable equal to one if the number of earnings announcements on the day of the revision issued by the analyst is above median, and zero otherwise. We observe relatively stable patterns across the forecasting styles in that all styles present a negative coefficient on *Busyness*. Because Style 3 serves as our baseline, this suggests that analysts are most likely to be assigned to Style 3 when they face more busyness and distraction. Second, we consider the sign of the earnings announcement news issued by the firm prior to the analyst forecast revision. Prior research suggests that negative earnings news yields smaller per-dollar market reactions than positive earnings news (e.g., Bartov, Givoly, and Hayn 2002; Lopez and Rees 2002). Accordingly, we consider whether forecasting style varies by the sign of the earnings surprise immediately preceding the revision. Interestingly, we find that Styles 1 and 5 are less likely when prior earnings surprise is positive, while Styles 2 and 4 are more likely. Third, we examine whether the fiscal quarter that is forecasted is an important consideration. Specifically, we include *FourthQuarterEA*, which is an indicator variable equal to one if the quarter for which the analyst provides a forecast is the fourth quarter, and zero otherwise. The coefficient is positive for Style 1, while it is negative for Styles 4 and 5, suggesting that more accurate styles are likely when the analyst is forecasting for a quarter earlier in the year.

Overall, our determinants analysis provides some interesting insights, but also suggests that forecasting style assignment is not fully explained by the factors we consider. As a final test, we consider the overall ability to predict style classification. To maximize predictive power, we employ a Random Forest classifier that uses the same data as Table 9. We tune the model using 5-fold cross-validation with 80 percent of the data and evaluate fit on the remaining hold-out sample.[17] Table 10 tabulates predicted classification vs. actual style. To facilitate discussion, we scale rows by the number of observations in each style (so rows sum to 100 percent).

A few interesting patterns emerge. First, the largest percentages of "accurate" assignments occur in Style 1 (55%) and Style 5 (40%), the least and most accurate styles, respectively. However, 25 percent of observations in Style 5 have a predicted classification of Style 1. Moreover, far more observations were predicted to be in Style 1 than actually are (20,382 vs. 7,364, untabulated). We interpret this to suggest that there are many instances in which observable factors suggest a forecast should be poor, but analysts perform much better than expected.

6. Conclusion

We employ a novel machine-learning technique ("regression-clustering") to identify five distinct analyst forecasting styles. We first document significant variation in how each of the forecasting styles contributes to the consensus analyst forecast accuracy and in how each

[17] We use the RandomForecastClassifier object from python's scikit learn. Our tuning procedure suggested balanced class-weights, a maximum depth of 100, maximum features equal to the natural log of total features, and 500 estimators. Final model accuracy was 38%, or nearly twice the unconditional expectation of 20%.

incorporates public information from four sources (firm fundamentals, valuation multiples, momentum signals, and herding indicators). We next find that, incremental to the number of analysts following the firm ("analyst coverage"), the number of unique forecasting styles ("style coverage") relates positively to both consensus forecast dispersion and accuracy. Further, our tests reveal that greater style coverage, but not analyst coverage, improves the information environment of firms as reflected in smaller earnings announcement surprises and reduced information asymmetry. Finally, we provide evidence of cross-sectional variation in analyst forecasting style based on the fiscal quarter forecasted, analyst busyness, firm volatility, and the sign of the most recent earnings news. Overall, our study sheds light on how analysts develop their earnings forecasts and documents significant capital market benefits associated with different analyst forecasting styles.

References

Abarbanell, J., W. Lanen, and R. Verrecchia. 1995. “Analysts’ Forecasts as Proxies for Investor Beliefs in Empirical Research.” *Journal of Accounting and Economics* 20, no. 1: 31–60.

Allen, E. J., Larson, C. R., & Sloan, R. G. (2013). “Accrual reversals, earnings and stock returns. *Journal of Accounting and Economics*.” *56*(1), 113–129.

Ajinkya, B., R. Atiase, and M. Gift. 1991. “Volume of Trading and the Dispersion in Financial Analysts’ Earnings Forecasts.” *The Accounting Review* 66, no. 2: 389–401.

Amiram, D., E. Owens, and O. Rozenbaum. 2016. “Do Information Releases Increase or Decrease Information Asymmetry? New Evidence from Analyst Forecast Announcements.” Journal of Accounting and Economics 62, no. 1: 121–38.

Baker, M. and J. Wurgler. 2006. “Investor Sentiment and the Cross-Section of Stock Returns.” The Journal of Finance 61, no. 4.

Bamber, L., J. Jiang, and I. Wang. 2010. “What’s My Style? The Influence of Top Managers on Voluntary Corporate Financial Disclosure.” *The Accounting Review* 85, no. 4: 1131–62.

Barron, O., O. Kim, S. Lim, and D. Stevens. 1998. “Using Analysts’ Forecasts to Measure Properties of Analysts’ Information Environment.” *The Accounting Review* 73, no. 4: 421–33.

Barry, C., and R. Jennings. 1992. “Information and Diversity of Analyst Opinion.” *Journal of Financial and Quantitative Analysis* 27, no. 2: 169.

Bartov, E., D. Givoly, and C. Hayn. 2002. “The Rewards to Meeting or Beating Earnings Expectations.” *Journal of Accounting and Economics* 33, no. 2: 173–204.

Bertrand, M., and A. Schoar. 2003. “Managing with Style: The Effect of Managers on Firm Policies.” *The Quarterly Journal of Economics* 118, no. 4: 1169–1208.

Blankespoor, E., E. deHaan, and I. Marinovic. 2020. “Disclosure Processing Costs, Investors’ Information Choice, and Equity Market Outcomes: A Review.” *Journal of Accounting and Economics* 70, no. 2: 101344.

Bradshaw, M. 2011. “Analysts’ Forecasts: What Do We Know after Decades of Work?” SSRN Scholarly Paper.

———. 2004. “How Do Analysts Use Their Earnings Forecasts in Generating Stock Recommendations?” *The Accounting Review* 79, no. 1: 25–50.

Brown, L., A. Call, M. Clement, and N. Sharp. 2015. “Inside the ‘Black Box’ of Sell-Side Financial Analysts.” *Journal of Accounting Research* 53, no. 1: 1–47.

Brown, L., G. Richardson, and S. Schwager. 1987. “An Information Interpretation of Financial Analyst Superiority in Forecasting Earnings.” *Journal of Accounting Research* 25, no.1:49–67.

Chang, Y., P. Hsiao, A. Ljungqvist, and K. Tseng. 2021. “Testing Disagreement Models.” *Journal of Finance,* forthcoming.

Clement, M. 1999. “Analyst Forecast Accuracy: Do Ability, Resources, and Portfolio Complexity Matter?” *Journal of Accounting and Economics* 27, no. 3: 285–303.

Corwin, S., S. Larocque, and M. Stegemoller. 2017. “Investment Banking Relationships and Analyst Affiliation Bias: The Impact of the Global Settlement on Sanctioned and Non-Sanctioned Banks.” *Journal of Financial Economics* 124, no. 3: 614–31.

Cronqvist, H., and S. Siegel. 2015. “The Origins of Savings Behavior.” *Journal of Political Economy* 123, no. 1: 123–69.

Da, Z., and X. Huang. 2019.“Harnessing the Wisdom of Crowds.” *Management Science* 66, no. 5: 1847–67.

Das, S., C. Levine, and K. Sivaramakrishnan. 1998. “Earnings Predictability and Bias in Analysts’ Earnings Forecasts.” *The Accounting Review* 73, no. 2 (1998): 277–94.

Daniel, K., M. Grinblatt, S. Titman, and R. Wermers. 1997. Measuring Mutual Fund Performance with Characteristic-Based Benchmarks. *The Journal of Finance* 52 (3): 1035–58.

Diether, K., C. Malloy, and A. Scherbina. “Differences of Opinion and the Cross Section of Stock Returns.” *The Journal of Finance* 57, no. 5 (2002): 2113–41.

Drake, M., L. Rees, and E. Swanson. 2011. “Should Investors Follow the Prophets or the Bears? Evidence on the Use of Public Information by Analysts and Short Sellers.” *The Accounting Review* 86, no. 1: 101–30.

Duru, A., and D. Reeb. 2002. “International Diversification and Analysts’ Forecast Accuracy and Bias.” *The Accounting Review* 77, no. 2: 415–33.

Dyreng, S., M. Hanlon, and E. Maydew. 2010. “The Effects of Executives on Corporate Tax Avoidance.” *The Accounting Review* 85, no. 4: 1163–89.

Ertimur, Y., W. Mayew, and S. Stubben. 2011. “Analyst Reputation and the Issuance of Disaggregated Earnings Forecasts to I/B/E/S” *Review of Accounting Studies.*

Francis, J., M. Pinnuck, and O. Watanabe. 2014. “Auditor Style and Financial Statement Comparability.” *The Accounting Review* 89, no. 2: 605–33.

Green, J., J. Hand, and F. Zhang. 2017. "The Characteristics That Provide Independent Information about Average U.S. Monthly Stock Returns." *The Review of Financial Studies* 30, no. 12: 4389–4436.
Gomez, E., F. Heflin, J. Moon, and J. Warren. 2022. "Crowdsourced Financial Analysis and Information Asymmetry at Earnings Announcements." Social Science Research Network.
Graham, J., C. Harvey, and S. Rajgopal. 2005."The Economic Implications of Corporate Financial Reporting." *Journal of Accounting and Economics* 40, no. 1–3: 3–73.
Groysberg, B., P. Healy, and D. Maber. 2011. "What Drives Sell-Side Analyst Compensation at High-Status Investment Banks?" *Journal of Accounting Research* 49, no. 4: 969–1000.
Hirshleifer, D., J. Myers, L. Myers, and S. Teoh. 2008. "Do Individual Investors Cause Post-Earnings Announcement Drift? Direct Evidence from Personal Trades." *The Accounting Review* 83, no. 6: 1521–50.
Irvine, P. 2004. "Analysts' Forecasts and Brokerage-Firm Trading." *The Accounting Review* 79, no. 1: 125–49.
Jacob, J., T. Lys, and M. Neale. 1999. "Expertise in Forecasting Performance of Security Analysts." *Journal of Accounting and Economics* 28, no. 1: 51–82.
Johnson, T. 2004. "Forecast Dispersion and the Cross Section of Expected Returns." *The Journal of Finance* 59, no. 5: 1957–78.
Johnston, J., and J. Zhang. 2020. "Auditor Style and Financial Reporting Similarity." *Journal of Information Systems* 35, no. 1: 79–99.
Kawano, A., K. Honda, H. Kasugai, and A. Notsu. 2013. "A Greedy Algorithm for K-Member Co-Clustering and Its Applicability to Collaborative Filtering." *Procedia Computer Science* 22 (2013):
Kim, O., and R. Verrecchia. 1994. "Market Liquidity and Volume around Earnings Announcements." *Journal of Accounting and Economics* 17, no. 1: 41–67.
Larcker, D. F., & Richardson, S. A. (2004). Fees Paid to Audit Firms, Accrual Choices, and Corporate Governance. *Journal of Accounting Research*, *42*(3), 625–658.
Lee, C., B. Mucklow, and M. Ready. 1993. "Spreads, Depths, and the Impact of Earnings Information: An Intraday Analysis." *Review of Financial Studies* 6, no. 2: 345.
Lopez, T, and L. Rees. 2002. "The effect of meeting analyst forecasts and systematic positive forecast errors on the information content of unexpected earnings." *Journal of Accounting, Auditing, and Finance.* 17 (Spring): 155-184.
Lys, T., and S. Sohn. 1990. "The Association between Revisions of Financial Analysts' Earnings Forecasts and Security-Price Changes." Journal of Accounting and Economics 13, no. 4: 341–63.
Lys, T., and L. Soo. 1995. "Analysts' Forecast Precision as a Response to Competition." *Journal of Accounting, Auditing & Finance* 10, no. 4: 751–65.
Mauler, L. 2018. "The Effect of Analysts' Disaggregated Forecasts on Investors and Managers: Evidence Using Pre-Tax Forecasts." *The Accounting Review* 94, no. 3: 279–302.
Merkley, K., R. Michaely, and J. Pacelli. 2020. "Cultural Diversity on Wall Street: Evidence from Consensus Earnings Forecasts." *Journal of Accounting and Economics* 70, no. 1: 101330.
O'Brien, P., and R. Bhushan. 2009. "Analyst Following and Institutional Ownership." *Journal of Accounting Research* 28 (1990): 55–76.
Petersen, M. 2009. "Estimating standard errors in finance panel data sets: Comparing approaches." *Review of Financial Studies* 22(1): 435–480.
Qu, C. 2020. "Board Members with Style: The Effect of Audit Committee Members and Their Personal Styles on Financial Reporting Choices." *Journal of Accounting, Auditing & Finance* 35, no. 3: 530–57.
Ramnath, S., S. Rock, and P. Shane. 2008. "A Review of Research Related to Financial Analysts' Forecasts and Stock Recommendations." *Foundations and Trends in Finance* 2
Richardson, S., S. Teoh, and P. Wysocki. 2004. "The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives." *Contemporary Accounting Research*, 21(4), pp.885-924.
Schipper, K. 1991. Commentary on analysts' forecasts. *Accounting Horizons*. 5 (December): 105-121.
Späth, H. 1979. "Algorithm 39 Clusterwise Linear Regression." *Computing* 22, no. 4: 367–73.
Stickel, S. 1990. "Predicting Individual Analyst Earnings Forecasts." *Journal of Accounting Research* 28, no. 2: 409–17.
Yohn, T. 1998. "Information Asymmetry Around Earnings Announcements." *Review of Quantitative Finance & Accounting* 11, no. 2: 165–82.
Zarnowitz, V., and P. Braun. 1993. "Twenty-Two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance: Studies in Business Cycles, Volume 28." *In Business Cycles, Indicators, and Forecasting,* 11–84. Chicago: University of Chicago Press, 1993.

APPENDIX A

Variable Definitions

Variable	Definition
Variables Used to Construct Analyst Forecast Style	
Revision	Sell-side analyst forecast revision, measured as the EPS forecast of the individual analyst minus the most recent previous EPS forecast of that same analyst, scaled by prior period stock price.
d_Sales	The four-quarter rolling change in quarterly sales revenue.
d_OCF	The year-over-year change in operating cash flow per share.
d_CapEx	The year-over-year change in capital expenditures per share.
EAReturns_Prior4	The average of the firm's past four earnings announcement abnormal returns, where abnormal returns are measures as buy and hold abnormal returns (using portfolio returns calculated from Daniel, Grinblatt, Titman, and Wermers (1997), and if missing, the value-weighted return from CRSP) over day 0 and +1 relative to the earnings announcement.
ReturnSinceLast	Following Stickel (1990), the buy and hold abnormal return from the day after the analyst's prior forecast to the three days before the analyst's current forecast. Abnormal returns are measures as buy and hold abnormal returns (using portfolio returns calculated from Daniel, Grinblatt, Titman, and Wermers (1997), and if missing, the value-weighted return from CRSP).
IncomeInc	Following Green, Hand, and Zhang (2017), the number of consecutive quarters in which the firm had an earnings increase over the same quarter last year.
BM	The natural logarithm of the firm's book-to-market ratio at the beginning of the period.
CFP	Following Green, Hand, and Zhang (2017), Operating cash flows divided by year-end market capitalization.
SP	Following Green, Hand, and Zhang (2017), annual sales divided by year-end market capitalization.
d_Consensus	Following Stickel (1990), the change in the consensus from after the previous analyst forecast to right before the current analyst forecast.
d_LTG	Change in the consensus long-term growth forecast from after the previous analyst forecast to right before the current analyst forecast.
BusPress	The average RavenPack event sentiment score of all business press articles during the quarter.

Variables used to Evaluate Outcomes of Analyst Forecast Style

Accuracy	Consensus analyst forecast accuracy, defined as the absolute difference between the median consensus analyst earnings forecast and actual earnings, multiplied by -100, scaled by beginning-of-the-period price.
Dispersion	Consensus analyst forecast dispersion, defined as the standard deviation of all analyst forecasts that make up the consensus.

|*EAReturns*|	The absolute value of buy and hold abnormal returns (using portfolio returns calculated from Daniel, Grinblatt, Titman, and Wermers (1997), and if missing, the value-weighted return from CRSP) over day 0 and +1 relative to the earnings announcement for which the analyst forecasted.
EAReturns	Buy and hold abnormal returns (using portfolio returns calculated from Daniel, Grinblatt, Titman, and Wermers (1997), and if missing, the value-weighted return from CRSP) over day 0 and +1 relative to the earnings announcement for which the analyst forecasted.
AbSpread	Average abnormal bid-ask spread on the day of and day following the earnings announcement for which the analyst provided a forecast. Abnormal spread is measured as the average percentage effective bid-ask spread on the day of and following the earnings announcement, subtracted by the average percentage effective bid-ask spread during the quarter, and multiplied by 10,000 for scaling purposes.
$StyleFollowing_X$	The number of analysts following the firm that issue a forecast with x forecast style.
Style Coverage	The number of unique forecast styles covering the firm.
Analyst Coverage	The number of analysts following the firm.
Size	The natural logarithm of the firm's total assets at the beginning of the period.
BM	The natural logarithm of the firm's book-to-market ratio at the beginning of the period.
ROA	The firm's return on assets at the beginning of the period, measured as income before extraordinary items scaled by total assets.
StdROA	The standard deviation of the firm's ROA over the last five years.
Ret	The firm's average monthly return over the last 12 months.
StdRet	The standard deviation of daily stock returns over the last 12 months.
EarnSurp	Earnings surprise of the EA for which the analyst is forecasting, calculated using the most recent analyst consensus prior to the EA and scaled by beginning-of-the-period stock price.
VIX_prior30days	The average VIX level in the 30 days prior to the analyst forecast, as measuring using the Chicago Board Options Exchange (CBOE) data.
MktSentiment	Investor sentiment in the month of the analyst forecast, as measured using the data from Baker and Wurgler (2006).
Ln(BrokerageSize)	The natural logarithm of one plus the number of analysts employed at the analyst's brokerage.
AllStar	An indicator variable equal to one if the analyst was rated an All-Star in the current year, and zero otherwise.
Busyness	An indicator variable equal to one if the analyst forecast occurs on a day with an above median number of firm earnings announcements, and zero otherwise.
PriorGoodNewsEA	An indicator variable equal to one if the most recent earnings announcement before the analyst forecast was good news, and zero otherwise.
FourthQuarterEA	An indicator variable equal to one if the earnings announcement for which the analyst is forecasting is the 4th quarter, and zero otherwise.

Fig. 1. Regression Diagnostics

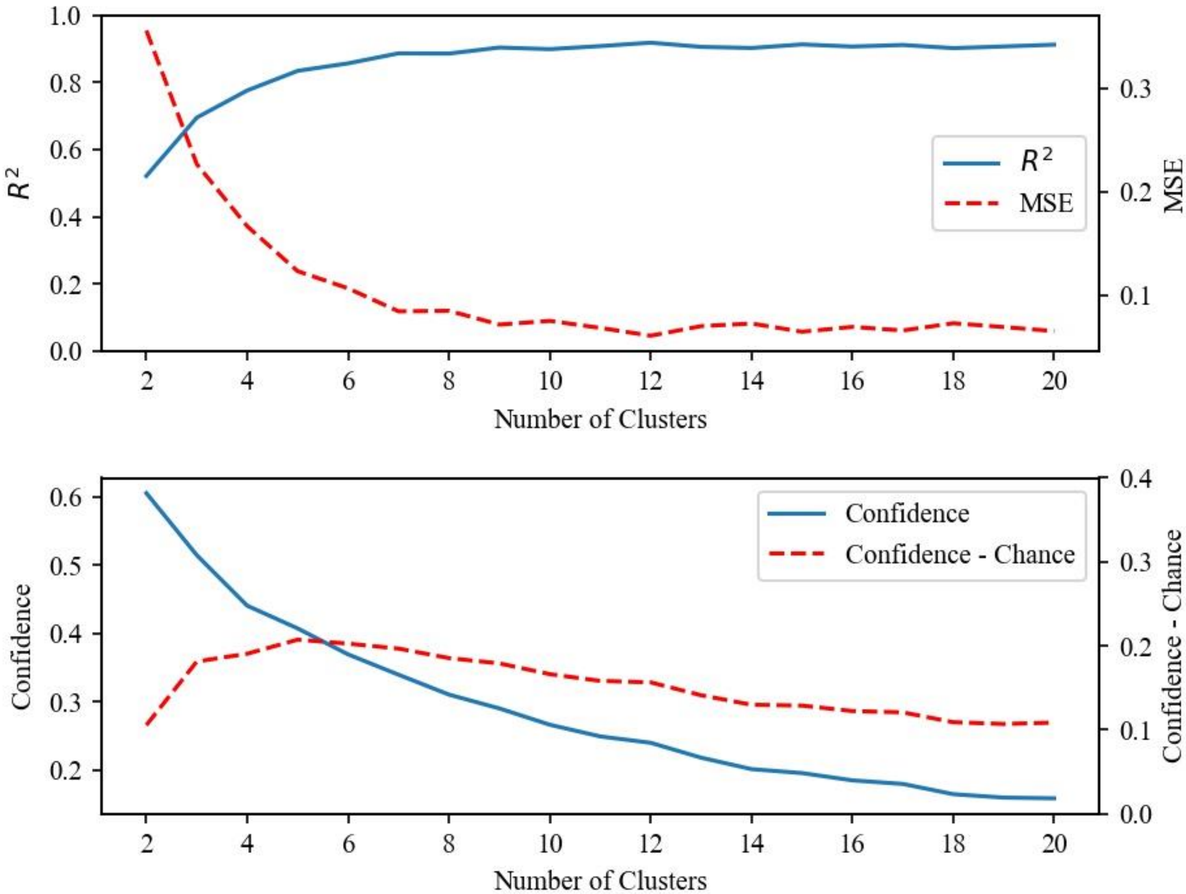

Figure 1 presents results from evaluating the number of appropriate analyst forecast clusters. We evaluate model fit for between 2 and 20 clusters and employ a hold-out sample to avoid overfitting. We consider two metrics to evaluate fit. (1) The R^2 value of estimation and (2) the average "confidence" of cluster assignment (or the average weights for assigned clusters) relatively to unconditional assignments.

Fig. 2. Use of Public Information by Analyst Forecast Style

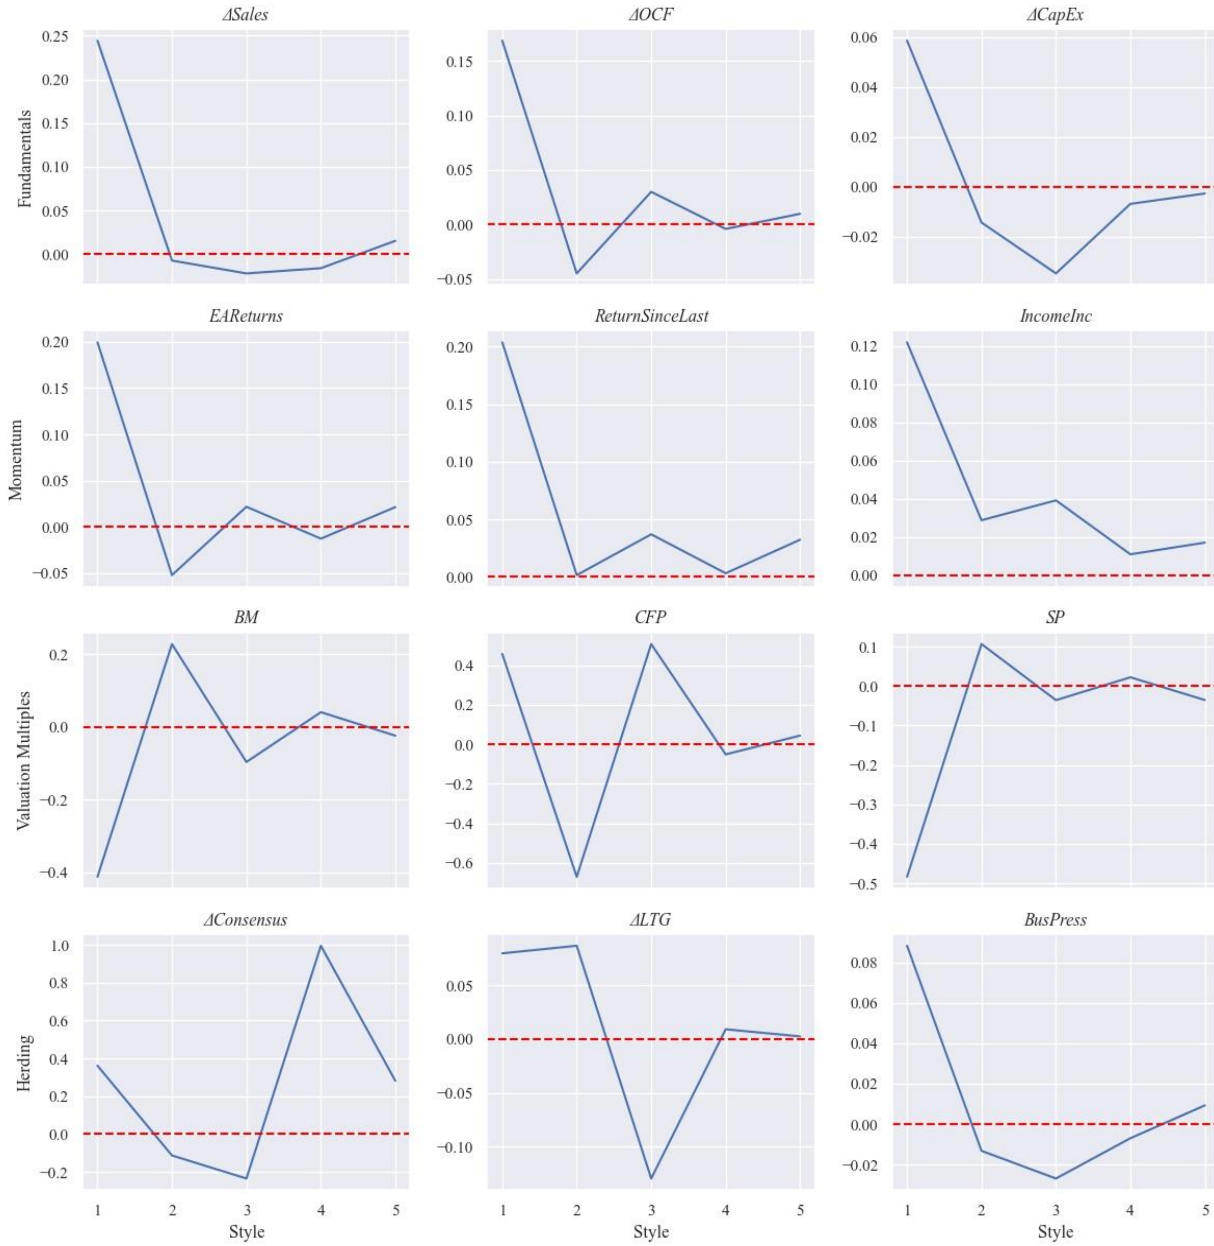

Figure 2 presents plots of the 12 coefficients included in equation 1. The *x*-axis corresponds to style assignments and is increasing with style accuracy. To facilitate magnitude and trend comparisons, all variables in these regressions are standardized to have mean 0, standard deviation 1. The y-axis varies with the scale of the coefficients, so we include a red horizontal line at a standardized coefficient magnitude of zero.

Table 1

Sample Breakdown by Fama-French Industry Classification

Fama-French Industry	Style 1	Style 2	Style 3	Style 4	Style 5	Total
1	1,580 (4.2%)	2,741 (5.0%)	3,445 (5.5%)	4,542 (4.6%)	6,162 (5.3%)	18,470 (5.0%)
2	968 (2.6%)	1,373 (2.5%)	1,758 (2.8%)	2,051 (2.1%)	2,734 (2.4%)	8,884 (2.4%)
3	4,006 (10.7%)	5,915 (10.9%)	7,196 (11.5%)	9,879 (10.0%)	12,589 (10.9%)	39,585 (10.8%)
4	3,286 (8.8%)	4,350 (8.0%)	4,669 (7.5%)	6,758 (6.8%)	7,893 (6.8%)	26,956 (7.3%)
5	940 (2.5%)	1,565 (2.9%)	1,868 (3.0%)	2,494 (2.5%)	3,537 (3.1%)	10,404 (2.8%)
6	11,128 (29.8%)	13,594 (25.0%)	14,722 (23.6%)	23,202 (23.5%)	23,680 (20.5%)	86,326 (23.5%)
7	1,401 (3.7%)	1,501 (2.8%)	1,875 (3.0%)	2,866 (2.9%)	3,494 (3.0%)	11,137 (3.0%)
8	389 (1.0%)	508 (0.9%)	626 (1.0%)	877 (0.9%)	1,577 (1.4%)	3,977 (1.1%)
9	3,530 (9.4%)	6,578 (12.1%)	8,125 (13.0%)	14,323 (14.5%)	16,227 (14.1%)	48,783 (13.3%)
10	3,486 (9.3%)	5,598 (10.3%)	5,964 (9.6%)	9,835 (10.0%)	10,907 (9.5%)	35,790 (9.7%)
11	2,277 (6.1%)	3,936 (7.2%)	4,717 (7.6%)	9,599 (9.7%)	11,636 (10.1%)	32,165 (8.7%)
12	4,391 (11.7%)	6,730 (12.4%)	7,345 (11.8%)	12,258 (12.4%)	14,843 (12.9%)	45,567 (12.4%)
Total	**37,382 (10.2%)**	**54,389 (14.8%)**	**62,310 (16.9%)**	**98,684 (26.8%)**	**115,279 (31.3%)**	**368,044**

Table 1 presents a breakdown of the sample of forecast styles by Fama-French 12-industry classifications. Each cell in the table first presents the number of observations in that industry-style combination. In parentheses, we also report the percentage of each style corresponding to the row-denoted industry (i.e., the column percentages sum to 100 percent).

TABLE 2

Descriptive Statistics

Panel A: *Variables used to construct analyst forecast style*

Variable	N	Mean	Lower Quartile	Median	Upper Quartile	Std Dev
Dependent Variable:						
Revision	368,044	-0.154	-0.182	-0.039	0.053	0.833
Independent Variables:						
Fundamentals						
Δ*Sales*	368,044	1.018	0.990	1.013	1.040	0.062
Δ*OCF*	368,044	-0.010	-0.380	0.024	0.392	1.614
Δ*CapEx*	368,044	0.022	-0.040	0.009	0.081	0.393
Momentum						
EAReturns_Prior4	368,044	0.005	-0.074	0.007	0.087	0.148
ReturnSinceLast	368,044	0.000	-0.055	0.000	0.053	0.116
IncomeInc	368,044	0.891	0.000	1.000	1.000	1.113
Valuation Multiples						
BM	368,044	0.437	0.208	0.361	0.586	0.343
CFP	368,044	0.100	0.053	0.083	0.131	0.104
SP	368,044	1.041	0.324	0.624	1.199	1.278
Herding						
Δ*Consensus*	368,044	-0.063	-0.048	0.000	0.007	0.329
Δ*LTG*	368,044	-0.532	-0.596	0.000	0.228	26.386
BusPress	368,044	39.337	43.7	49.6	52.8	21.9

Panel B: *Variables used to evaluate outcomes of analyst forecast style*

Variable	N	Mean	Lower Quartile	Median	Upper Quartile	Std Dev
Dependent Variables						
Accuracy	43,645	-0.543	-0.453	-0.182	-0.076	1.166
Dispersion	43,645	0.270	0.055	0.115	0.263	0.449
|*EAReturns*|	43,645	0.544	0.017	0.039	0.075	0.050
EAReturns	43,645	0.001	-0.038	0.000	0.040	0.074
AbSpread	43,645	3.069	-0.195	0.934	3.419	9.445
Style Variables						
$StyleFollowing_1$	43,645	0.819	0.000	0.000	1.000	1.790
$StyleFollowing_2$	43,645	1.191	0.000	0.000	2.000	2.009
$StyleFollowing_3$	43,645	1.366	0.000	1.000	2.000	2.129
$StyleFollowing_4$	43,645	2.161	0.000	1.000	3.000	2.650
$StyleFollowing_5$	43,645	2.523	1.000	2.000	3.000	2.708
Style Coverage	43,645	2.817	2.000	3.000	3.000	0.999
Other Variables						
Analyst Coverage	43,645	8.061	4.000	7.000	10.000	5.093
Size	43,645	7.819	6.594	7.699	8.885	1.693
BM	43,645	0.356	0.203	0.331	0.489	0.227
ROA	43,645	0.007	0.001	0.011	0.022	0.034
StdROA	43,645	0.017	0.004	0.009	0.019	0.023
Ret	43,645	0.015	-0.001	0.015	0.031	0.030
StdRet	43,645	0.024	0.016	0.022	0.030	0.011

Table 2 presents descriptive statistics. Panel A presents descriptive statistics for variables used in generating analyst forecast cluster style assignments. The data is at the analyst forecast level. Panel B presents descriptive statistics for variables used in evaluating outcomes of analyst forecast styles. The data is at the firm-quarter level.

TABLE 3

Can We Identify Analyst Forecast Style?

Dependent Variable: *Accuracy*

	[1]	[2]
StyleFollowing$_1$	**0.010**	**0.020****
	(1.00)	**(2.03)**
StyleFollowing$_2$	**0.037*****	**0.034*****
	(4.42)	**(4.37)**
StyleFollowing$_3$	**0.030*****	**0.034*****
	(4.29)	**(5.09)**
StyleFollowing$_4$	**0.053*****	**0.040*****
	(6.19)	**(4.96)**
StyleFollowing$_5$	**0.090*****	**0.081*****
	(8.51)	**(8.52)**
Size		0.071**
		(2.48)
BM		-0.158***
		(-6.42)
ROA		2.000***
		(3.72)
StdROA		-3.202***
		(-5.59)
Ret		4.098***
		(14.34)
StdRet		-24.958***
		(-10.40)
Observations	43,645	43,645
Fixed Effects	Firm & yr-qtr	Firm & yr-qtr
Clustering	Firm	Firm
Adjusted R^2	0.581	0.610

Table 3 presents coefficients (t-statistics) for tests of EQ1: Can we identify unique forecast styles? The dependent variable is the consensus analyst forecast accuracy (*Accuracy*). Column 1 (2) presents results without (with) control variables. *StyleFollowing$_x$* is the natural logarithm of one plus the number of analysts following the firm that issue a forecast with x forecast style. *** (**, *) denotes two-tailed significance at the p<0.01 (p<0.05, p<0.10) level for regression coefficients. All variables are defined in Appendix A.

TABLE 4

Does Style Coverage Improve the Information Environment of the Firm? Evidence from Consensus Forecast Dispersion and Accuracy

Panel A: *Consensus Forecast Dispersion*

	Logged coverage variables		*Count-based coverage variables*	
	[1]	[2]	[3]	[4]
Analyst Coverage	-0.122***	-0.165***	-0.007***	-0.010***
	(-10.84)	(-15.25)	(-8.96)	(-12.68)
Style Coverage		**0.129*** **		**0.033*** **
		(13.15)		**(13.36)**
Size	0.016*	0.016*	0.009	0.008
	(1.72)	(1.74)	(1.00)	(0.87)
BM	0.059***	0.059***	0.061***	0.062***
	(6.58)	(6.56)	(6.79)	(6.83)
ROA	-1.080***	-1.104***	-1.147***	-1.174***
	(-7.97)	(-8.16)	(-8.30)	(-8.50)
StdROA	0.997***	1.010***	1.012***	1.024***
	(5.14)	(5.22)	(5.14)	(5.20)
Ret	-1.938***	-1.929***	-1.908***	-1.882***
	(-18.85)	(-18.74)	(-18.27)	(-17.97)
StdRet	7.558***	7.483***	7.709***	7.646***
	(14.02)	(13.88)	(13.89)	(13.73)
Observations	43,645	43,645	43,645	43,645
Fixed Effects	Firm & yr-qtr	Firm & yr-qtr	Firm & yr-qtr	Firm & yr-qtr
Clustering	Firm	Firm	Firm	Firm
Adjusted R^2	0.598	0.602	0.594	0.597

Panel B: *Consensus Forecast Accuracy*

	Logged coverage variables		*Count-based coverage variables*	
	[1]	[2]	[3]	[4]
Analyst Coverage	0.293***	0.275***	0.020***	0.019***
	(9.72)	(9.92)	(8.75)	(8.55)
Style Coverage		**0.053****		**0.016*****
		(2.41)		**(2.67)**
Size	0.044	0.044	0.063**	0.062**
	(1.50)	(1.51)	(2.19)	(2.16)
BM	-0.147***	-0.147***	-0.153***	-0.153***
	(-6.10)	(-6.11)	(-6.25)	(-6.25)
ROA	1.868***	1.859***	1.994***	1.982***
	(3.50)	(3.49)	(3.72)	(3.70)
StdROA	-3.196***	-3.189***	-3.219***	-3.213***
	(-5.64)	(-5.63)	(-5.62)	(-5.61)
Ret	4.236***	4.233***	4.191***	4.198***
	(15.07)	(15.07)	(14.70)	(14.73)
StdRet	-24.589***	-24.606***	-25.031***	-25.041***
	(-10.57)	(-10.58)	(-10.43)	(-10.44)
Observations	43,645	43,645	43,645	43,645
Fixed Effects	Firm & yr-qtr	Firm & yr-qtr	Firm & yr-qtr	Firm & yr-qtr
Clustering	Firm	Firm	Firm	Firm
Adjusted R^2	0.613	0.613	0.610	0.610

Table 4 presents coefficients (t-statistics) for tests of EQ2: Does style coverage improve the information environment of the firm? In Panel A the dependent variable is the consensus analyst forecast dispersion (*Dispersion*). In Panel B the dependent variable is the consensus analyst forecast accuracy (*Accuracy*). For both panels, A*nalyst Coverage* is the number of analysts covering the firm. *Style Coverage* is the number of unique styles of analysts covering the firm. Columns 1 and 3 present results with just *Analyst Coverage*, while Columns 2 and 4 present results after including *Style Coverage*. In Columns 1 and 2, we use the natural logarithm of one plus *Analyst Coverage* and *Style Coverage*. In Columns 3 and 4, *Analyst Coverage* and *Style Coverage* are the unlogged count variables. *** (**, *) denotes two-tailed significance at the p<0.01 (p<0.05, p<0.10) level for regression coefficients. All variables are defined in Appendix A.

TABLE 5

Does Style Coverage Improve the Information Environment of the Firm? Evidence from Subsequent Earnings Announcement Returns

Panel A: *Abnormal Subsequent EA Returns*

Dependent Variable: |*EAReturns*|

	Logged coverage variables		*Not Logged coverage variables*	
	[1]	[2]	[3]	[4]
Analyst Coverage	0.498***	0.595***	0.058***	0.062***
	(9.26)	(9.40)	(10.64)	(10.30)
Style Coverage		**-0.339*** **		**-0.052** **
		(-3.32)		**(-1.97)**
Size	-0.655***	-0.656***	-0.665***	-0.663***
	(-36.18)	(-36.28)	(-36.85)	(-36.76)
BM	0.031	0.034	0.040	0.039
	(0.99)	(1.08)	(1.25)	(1.24)
ROA	2.166**	2.216**	2.213**	2.256**
	(2.06)	(2.10)	(2.11)	(2.15)
StdROA	0.027	0.072	-0.179	-0.132
	(0.02)	(0.05)	(-0.12)	(-0.09)
Ret	-19.458***	-19.487***	-19.200***	-19.246***
	(-15.98)	(-16.02)	(-15.74)	(-15.79)
StdRet	85.573***	85.739***	84.457***	84.631***
	(18.30)	(18.30)	(18.09)	(18.06)
Observations	43,645	43,645	43,645	43,645
Fixed Effects	None	None	None	None
Clustering	EA date	EA date	EA date	EA date
Adjusted R^2	0.110	0.110	0.111	0.111

Panel B: *ERC at Subsequent EA*

Dependent Variable: *EAReturns*

	Logged coverage variables		*Not Logged coverage variables*	
	[1]	[2]	[3]	[4]
EarnSurp	0.408***	0.351***	0.373***	0.316***
	(3.73)	(3.06)	(3.39)	(2.77)
Analyst Coverage	-0.772***	-0.668***	-0.063***	-0.055***
	(-5.27)	(-4.36)	(-4.31)	(-3.62)
Style Coverage		-0.305*		-0.089*
		(-1.73)		(-1.89)
*EarnSurp*Analyst Coverage*	-0.026	-0.095**	-0.005	-0.011***
	(-1.09)	(-2.46)	(-1.49)	(-2.67)
EarnSurp*Style Coverage		**0.127****		**0.033****
		(2.09)		**(2.02)**
Size	-1.189***	-1.191***	-1.214***	-1.210***
	(-7.91)	(-7.91)	(-8.07)	(-8.04)
BM	1.006***	1.006***	1.014***	1.012***
	(9.62)	(9.61)	(9.68)	(9.66)
ROA	33.802***	33.803***	33.525***	33.513***
	(14.58)	(14.56)	(14.48)	(14.44)
StdROA	5.650*	5.691*	5.744*	5.759*
	(1.88)	(1.89)	(1.91)	(1.91)
Ret	-19.379***	-19.468***	-19.297***	-19.448***
	(-10.36)	(-10.40)	(-10.32)	(-10.40)
StdRet	10.633	11.222	11.170	11.742*
	(1.54)	(1.63)	(1.61)	(1.70)
Observations	43,645	43,645	43,645	43,645
Fixed Effects	None	None	None	None
Clustering	EA date	EA date	EA date	EA date
Fully Interacted Model	Yes	Yes	Yes	Yes
Adjusted R^2	0.035	0.035	0.035	0.035

Table 5 presents coefficients (t-statistics) for tests of EQ2: Does style coverage improve the information environment of the firm? In Panel A the dependent variable is the absolute value of buy and hold abnormal returns over day 0 and +1 relative to the earnings announcement for which the analyst forecasted (|*EAReturns*|). In Panel B the dependent variable is signed abnormal returns over the same period (*EAReturns*). In both panels, *Analyst Coverage* is the number of analysts covering the firm. *Style Coverage* is the number of unique styles of analysts covering the firm. Columns 1 and 3 present results with just *Analyst Coverage*, while Columns 2 and 4 present results after including *Style Coverage*. In Columns 1 and 2, we use the natural logarithm of one plus *Analyst Coverage* and *Style Coverage*. In Columns 3 and 4, *Analyst Coverage* and *Style Coverage* are the unlogged count variables. *** (**, *) denotes two-tailed significance at the p<0.01 (p<0.05, p<0.10) level for regression coefficients. All variables are defined in Appendix A.

TABLE 6

Does Style Coverage Improve the Information Environment of the Firm? Evidence from Subsequent Earnings Announcement Bid-ask Spread

Dependent Variable: *AbSpread*

	Logged coverage variables		*Not Logged coverage variables*	
	[1]	[2]	[3]	[4]
Analyst Coverage	0.271*	0.928***	0.104***	0.156***
	(1.94)	(6.10)	(7.09)	(9.73)
Style Coverage		**-2.321*** **		**-0.640*** **
		(-11.00)		**(-11.52)**
Size	-0.738***	-0.747***	-0.854***	-0.835***
	(-16.80)	(-17.05)	(-18.56)	(-18.28)
BM	0.260***	0.278***	0.360***	0.356***
	(3.57)	(3.84)	(4.93)	(4.88)
ROA	-5.601***	-5.255**	-6.576***	-6.042***
	(-2.64)	(-2.49)	(-3.09)	(-2.84)
StdROA	-0.526	-0.215	-1.295	-0.712
	(-0.17)	(-0.07)	(-0.42)	(-0.23)
Ret	-47.958***	-48.150***	-47.137***	-47.701***
	(-16.52)	(-16.67)	(-16.21)	(-16.46)
StdRet	49.399***	50.531***	46.324***	48.472***
	(5.96)	(6.08)	(5.61)	(5.85)
Observations	43,645	43,645	43,645	43,645
Fixed Effects	None	None	None	None
Clustering	EA date	EA date	EA date	EA date
Adjusted R^2	0.046	0.049	0.048	0.052

Table 6 presents coefficients (t-statistics) for tests of EQ2: Does style coverage improve the information environment of the firm? The dependent variable is average abnormal bid-ask spread on the day of and day following the earnings announcement for which the analyst provided a forecast. (*AbSpread*). *Analyst Coverage* is the number of analysts covering the firm. *Style Coverage* is the number of unique styles of analysts covering the firm. Columns 1 and 3 present results with just *Analyst Coverage*, while Columns 2 and 4 present results after including *Style Coverage*. In Columns 1 and 2, we use the natural logarithm of one plus *Analyst Coverage* and *Style Coverage*. In Columns 3 and 4, *Analyst Coverage* and *Style Coverage* are the unlogged count variables. *** (**, *) denotes two-tailed significance at the p<0.01 (p<0.05, p<0.10) level for regression coefficients. All variables are defined in Appendix A.

TABLE 7

The Number of Unique Forecast Styles, Conditional on the Number of Firms Followed by an Analyst

# of firms followed	Number of Styles 1	2	3	4	5	% of Sample
1	13.7% (100.0%)					14%
2	2.2% (22.9%)	7.2% (77.1%)				9%
3	0.5% (6.6%)	4.3% (53.1%)	3.3% (40.3%)			8%
4	0.1% (1.9%)	2.4% (29.7%)	4.3% (54.2%)	1.1% (14.2%)		8%
5	0.0% (0.6%)	1.3% (16.3%)	4.0% (51.7%)	2.3% (29.1%)	0.2% (2.4%)	8%
6	0.0% (0.4%)	0.7% (9.5%)	3.3% (42.0%)	3.2% (40.8%)	0.6% (7.3%)	8%
7	0.0% (0.1%)	0.4% (5.7%)	2.5% (33.3%)	3.6% (48.1%)	0.9% (12.7%)	7%
8	0.0% (0.0%)	0.2% (3.3%)	2.0% (27.8%)	3.5% (48.7%)	1.5% (20.2%)	7%
9		0.1% (2.0%)	1.3% (19.9%)	3.3% (51.4%)	1.7% (26.6%)	6%
10+		0.1% (0.4%)	2.3% (9.5%)	10.5% (43.6%)	11.2% (46.5%)	24%

Table 7 cross-tabulates the number of firms followed by an analyst (rows) with the number of unique styles employed (columns). We collapse these data at the analyst-calendar quarter level (there are 56,890 unique analyst-firm-quarters in our sample). In each cell, we present two values. The first value represents the proportion of the sample in that cell. The second value in parentheses is the percentage of that row falling in that column. For example, 4.3% of the analysts in our sample follow 3 firms and use 2 different styles of forecasting. For analysts following 3 firms, 53.1% employ two different styles.

TABLE 8

Change in Forecast Style Over Time

		Style 1	*Style 2*	*Style 3*	*Style 4*	*Style 5*
		[1]	[2]	[3]	[4]	[5]
a	*Percent that change forecast style from prior quarter*	83.41%	80.18%	78.65%	69.61%	64.77%
b	*Percentage of observations in each style*	10.23%	14.82%	16.95%	26.76%	31.24%
c	*Unconditional Expectation (1 - b)*	89.77%	85.18%	83.05%	73.24%	68.76%
d	*Percent that change as a proportion of unconditional expectation (a / c)*	92.92%	94.12%	94.70%	95.05%	94.19%

Table 8 tabulates the proportion of quarter-over-quarter changes by style. We tabulate these by style since the number of observations in each style is not balanced. We also include the proportion of the sample in each style and the unconditional likelihood of changing (i.e., 1 minus the sample proportion). For all five styles, change rates are less than unconditional rates (i.e., row c < row d).

TABLE 9

Determinants of Analyst Forecast Style

	Style 1	*Style 2*	*Style 4*	*Style 5*
	[1]	[2]	[3]	[4]
VIX_prior30days	0.001	-0.005***	-0.010***	-0.010***
	(0.50)	(-3.56)	(-8.45)	(-9.24)
MktSentiment	0.215***	-0.011	-0.093***	-0.060*
	(4.81)	(-0.27)	(-2.65)	(-1.76)
StdRet	22.221***	5.723***	-0.148	-5.532***
	(25.33)	(6.89)	(-0.20)	(-7.63)
StdROA	1.945***	1.165***	-2.946***	-3.541***
	(5.78)	(3.63)	(-9.82)	(-12.02)
Ret	-3.781***	0.085	3.551***	2.843***
	(-15.70)	(0.37)	(17.43)	(14.36)
Ln(BrokerageSize)	-0.004	0.010	-0.002	-0.023***
	(-0.45)	(1.27)	(-0.33)	(-3.44)
AllStar	-0.017	0.004	-0.042***	0.006
	(-0.86)	(0.23)	(-2.82)	(0.44)
Busyness	-0.026***	-0.022***	-0.036***	-0.047***
	(-4.25)	(-4.05)	(-7.51)	(-10.21)
PriorGoodNewsEA	-0.116***	0.058***	0.130***	-0.051***
	(-8.11)	(4.60)	(11.70)	(-4.71)
FourthQuarterEA	0.072***	-0.011	-0.024**	-0.023*
	(4.57)	(-0.78)	(-1.98)	(-1.92)
Size	-0.025***	0.015***	0.079***	0.062***
	(-5.28)	(3.68)	(21.59)	(17.49)
BM	-0.098***	-0.060***	-0.133***	-0.043***
	(-11.56)	(-7.78)	(-19.77)	(-6.36)
ROA	-0.950***	1.368***	-1.089***	-1.885***
	(-4.18)	(6.13)	(-5.45)	(-9.67)
Observations	368,044	368,044	368,044	368,044
Adjusted R^2	0.613	0.613	0.610	0.610

Table 9 presents coefficients (t-statistics) for determinants of analyst forecast style. Table 9 uses a multinomial logit where the comparison group (i.e., base) is Style 3. Thus, Style 3 is not presented in Table 9. *** (**, *) denotes two-tailed significance at the p<0.01 (p<0.05, p<0.10) level for regression coefficients. All variables are defined in Appendix A.

TABLE 10

Predictability of Analyst Forecast Style

	[1]	[2]	[3]	[4]	[5]
Style 1	**55%**	8%	8%	13%	16%
Style 2	29%	**30%**	8%	15%	18%
Style 3	27%	7%	**31%**	15%	19%
Style 4	23%	8%	9%	**38%**	23%
Style 5	25%	7%	10%	19%	**40%**

Table 10 compares style assignments to predicted assignments based on a Random Forest Classifier. Specifically, each row corresponds to a unique forecasting sample, and columns denote the percentage of observations predicted to be in that style by the Random Forest classifier. The diagonal (bolded) denotes the percentage of observations in each row that were correctly classified.