Lured by the Consensus

Roni Michaely, Amir Rubin, Dan Segal, and Alexander Vedrashko*

Abstract:
We find that investors are fixated on analysts' consensus outputs (earnings forecasts, recommendations, and forecast dispersion), which can be inferior signals compared to the corresponding outputs provided by high-quality analysts, especially when a large number of high-quality analysts follow the firm. This result, which holds at the firm and market level, implies inefficient use of the information contained in analysts' outputs. Further, the post-earnings announcement drift (PEAD) phenomenon occurs only when high-quality analysts are more uncertain about the firm's performance than all analysts following the firm. We conclude that the market's fixation on consensus measures has significant negative economic implications.

Keywords: consensus, analyst quality, forecasts, post-earnings announcement drift, stock recommendations

JEL: G10; G11; G14; G17; G24; M41

---

*Michaely is from The Geneva Finance Research Institute, University of Geneva, (roni.michaely@unige.ch). Rubin is from Simon Fraser University and IDC (arubin@sfu.ca). Segal is from IDC and Singapore Management University (dsegal@idc.ac.il). Vedrashko is from Simon Fraser University (awv@sfu.ca). We thank the discussants and seminar participants at the AsianFA meeting in Tokio, EFMA meeting in Milan, FIRS meeting in Barcelona, MFS meeting in Budapest, Hebrew University, Tel-Aviv University, University of Connecticut, and University of Vienna for helpful comments and suggestions.

Analysts are recognized as important information providers in financial markets. Investors and academics alike use analysts' consensus earnings forecasts as a measure of the market expectations of firms' future earnings. The perceived importance of the consensus forecast has increased in recent years to the extent that even companies' investor relations departments follow the consensus on a continuous basis (Consensus earnings estimates report, 2013). Consensus analysts' recommendations have also become readily and, often, freely available for investors, e.g., on the Nasdaq website. However, by construction, the consensus ignores the possibility that analysts have different abilities and information sets, resulting in varying levels of forecast accuracy and recommendation quality that are persistent over time.[1]

Thus, fixation on the consensus can have negative economic implications: investors may be missing information and incur opportunity costs if they focus on the consensus and disregard the heterogeneity in analyst quality. On the other hand, using the single number of the consensus instead of differentiating among analysts can be optimal for investors. A simple average construct can achieve a better level of predictability than more complex models that incorporate quality differences (e.g., a literature review in Clemen, 1989); for example, the highest quality analyst's forecast can be less accurate than the consensus forecast, which diversifies individual errors.[2] Even if some analysts are consistently more accurate than the consensus, the costs of gathering information and spending cognitive effort on identifying analysts with superior ability may exceed

[1] Individual analysts' forecast accuracy systematically differs for reasons including analysts' varying experience (Mikhail, Walther, and Willis, 1997; Clement, 1999; Hirst, Hopkins, and Wahlen, 2004), aptitude (Jacob, Lys and Neale, 1999), education (Maines, McDaniel, and Harris, 1997; De Franco and Zhou, 2009), brokerage house association and underwriting relationships (Lin and McNichols, 1998; Clement, 1999), proximity to the firm (Malloy, 2005), lead analyst and star status categorization (Stickel, 1992; Cooper, Day, and Lewis, 2001), or work habits (Rubin, Segal, and Segal, 2017).

[2] For example, consider a firm followed by four analysts, whose rankings in terms of forecast accuracy are perfectly persistent over time. Let the top analyst's forecast error (the analyst's forecast minus actual earnings per share) be 2¢, and the remaining three analysts have forecast errors of -4¢, -5¢, and 8¢. The consensus forecast error is 0.25¢, less than the forecast error of the most accurate analyst and that of the average forecast of analysts in the top half of the ranking, who can be called high-quality analysts. Since the persistence of analysts' forecast accuracy is less than perfect in actual markets, relying on the consensus forecast can make even more sense for investors.

the economic benefits of going beyond the consensus, such as when analyst quality heterogeneity is small. In fact, investors' fixation on the mean of analysts' forecast distribution can be an instance of central fixation bias, which describes people's natural tendency to fixate their vision at the center of a group of objects and which can be optimal for initial information processing (Tatler, 2007).

Hence, our objectives are two-fold. We first assess whether investors mainly focus on the consensus output. Second, we examine whether analysts' consensus outputs (e.g., consensus earnings forecasts and consensus recommendations) are the right focal points for investors. Building on the concept of diversification borrowed from asset pricing, we take into account not only heterogeneity and persistence in analyst performance, which motivate the literature contrasting high-quality (HQ; e.g., more accurate, "star", "All-American") and low-quality (LQ) analysts, but also the diversification benefits of relying on the consensus, which can be a first-order effect.[3] Consequently, we advance the literature by considering the market's focus on consensus-type information and addressing the question of market pricing and economic efficiency of consensus outputs. Further, while each of the prior studies examines only one of the outputs of analysts and only at the firm level, we address our research question based on several outputs—earnings forecasts, recommendations, and forecast dispersion—at both the individual firm and market levels, which allows us to reach a general conclusion regarding the phenomenon of investors' consensus fixation.

Overall, we find that investors generally focus on the consensus, which leads to price inefficiency. Differentiating among analysts and following HQ analysts instead of the consensus can provide economic benefits. Our main results can be summarized as follows: first, the market

[3] The literature has found inefficiencies in how the market treats analyst heterogeneity in reacting to earnings forecast revisions (Stickel, 1992; Clement and Tse, 2003; Gleason and Lee, 2003; Chen, Francis, and Jiang, 2005) and recommendations (Michaely and Womack, 1999; Mikhail, Walther, and Willis, 2004; Loh and Mian, 2006; Sorescu and Subrahmanyam, 2006; Ertimur, Sunder, and Sunder, 2007; Fang and Yasuda, 2014; Kucheev, Ruiz, and Sorensson, 2017).

reacts to the consensus forecast rather than to the more accurate average forecast generated by HQ analysts, a categorization based on their past performance. This inefficient use of information in analysts' forecasts results in stock mispricing around earnings announcements. Second, we find that it is not optimal for investors to follow the consensus recommendation revision compared to the average recommendation revision of HQ analysts. Third, the dispersion of HQ analysts' forecasts is more informative about the firm's underlying uncertainty than the dispersion based on all analysts following the firm ("the consensus dispersion"). Fourth, we find that the post-earnings announcement drift (PEAD) arises only when the dispersion of HQ analysts is high relative to the consensus dispersion; thus, investors who do not differentiate among analysts forego significant returns. Finally, we find that HQ analysts' recommendation revisions and forecast dispersion aggregated over all firms in the sample predict the next month's stock market return and volatility, respectively, while the corresponding relations are not observed for the consensus recommendations and consensus forecast dispersion. Taken together, our findings on earnings announcements, recommendations, forecast dispersion, and the PEAD share an underlying economic mechanism, which we refer to as the consensus fixation phenomenon, that causes investors to systematically underweight quality differences among analysts, leading to pricing inefficiency, and to forego investment returns as a result.

Following the literature (e.g., Chen, Francis, and Jiang, 2005; Loh and Mian, 2006), we use the firm-year ranking approach and define HQ and LQ analysts according to their forecast accuracy rankings for the previous year's annual earnings announcement by the firm. We find persistence in forecasting ability across time and firms an analyst covers, indicating that forecasting performance captures analysts' quality. Next, consistent with the diversification effect—the more HQ analysts, the smaller the diversification advantage of the consensus relative to the HQ analysts' average forecast—we find that the average of HQ analysts' forecasts is more accurate than the consensus forecast only when a sufficient number of HQ analysts follow the firm.

Despite the superiority of the average of HQ analysts' forecasts for these announcements, the market reacts more vigorously to earnings surprises that are measured based on the consensus forecast than the average of HQ analysts' forecasts. This finding implies that the attention the market pays to the consensus forecast may be excessive, and investors fail to fully incorporate the information available in HQ analysts' forecasts. The market's fixation on the consensus forecast is inefficient and can be exploited. A strategy based on the difference between the mean forecast of HQ analysts and the consensus prior to the earnings announcement day yields economically and statistically significant abnormal returns over the announcement day and the following trading day.

This phenomenon manifests itself in other aspects of analysts' informational output, such as recommendation revisions and forecast dispersion. We find that the market does not efficiently impound HQ analysts' recommendation revisions into prices. HQ analysts' recommendation revisions are strong predictors of the stock returns in the following month, in contrast to the consensus recommendation revisions. This results in an opportunity cost for investors who do not take into account analyst quality heterogeneity and who use investment strategies based on the consensus recommendation, with a straightforward long-short strategy yielding a 1.1% monthly alpha. Concerning analysts' forecast dispersion, if the average of HQ analysts' forecasts is more informative than the consensus forecast, then the variability of their forecasts should also be a superior measure of uncertainty regarding future firm performance relative to the consensus dispersion. Indeed, we find that the dispersion of HQ analysts' forecasts is a strong predictor of the firm's stock return volatility in the month following the annual earnings announcement month, unlike the dispersion based on all analysts.

Building on studies showing that the PEAD arises in periods of uncertainty (Abarbanell, Lanen, and Verrecchia, 1995; Mendenhall, 2004; Zhang, 2006a; Francis et al., 2007; Hung, Li, and Wang, 2014) and using our finding that HQ analysts' forecast dispersion can proxy for the firm's uncertainty better than the consensus dispersion, we find a significantly greater PEAD

(9.4% after 11 months) when HQ analysts are relatively uncertain than the PEAD in the full sample of earnings announcements, which ignores analyst quality heterogeneity. Overall, these findings indicate that the consensus fixation bias is relevant to both short (e.g., earnings announcements) and long (the PEAD) investment horizon decisions. Investors recognizing the information contained in HQ analysts' forecasts are able to make significantly better investment decisions than those investors making their investment decisions based on the consensus forecast.

Finally, we examine whether the inferiority of consensus outputs is relevant at the market level. Since HQ analysts' recommendation changes are better at the firm level, HQ analysts' average recommendation change across all firms in the sample can also be a superior predictor of the market return. We find that HQ analysts' average recommendation changes predict the market returns for the following month, in contrast to the consensus recommendation changes. For example, a long-short strategy based on the direction of HQ analysts' recommendation revision relative to the consensus revision produces a 6.8% annualized calendar-time alpha in the month following the earnings announcement. The argument for the average dispersion of HQ analysts' forecasts rather than the consensus dispersion predicting market volatility is analogous, and we find that HQ analysts' normalized forecast dispersion predicts the next-month CBOE volatility index (VIX) and changes in the VIX, thereby allowing investors unimpeded by consensus fixation to earn higher returns.[4]

Our findings contribute to the literature across several dimensions. First, we analyze the market efficiency implications of analyst quality differences not previously considered in the literature. These implications include the short- and long-term reaction to earnings announcements,

[4] We also note that aggregating analysts' informational outputs about firms they follow is akin to surveys aggregating the opinions of representative households to measure consumer confidence (Ludvigson, 2004), both of which can be used to predict macroeconomic and market-level variables.

generating a more informative forecast dispersion measure, and finding the predictability of market returns and volatility based on the aggregation of analysts' firm-level outputs.

Second, the study suggests that investors' focus on the consensus earnings forecast may be a result of limited attention (e.g., Hirshleifer, Lim, and Teoh, 2009). Investors may prefer the expediency of the single number of the consensus to the exertion of cognitive effort to assess analyst quality. From this perspective, our finding that investors ignore the valuable information contained in HQ analysts' outputs contributes to the literature that uses subjective expectations of bond risk premia (Buraschi, Piatti, and Whelan, 2017), corrections for biases in individual analysts' forecasts (Chiang et al., 2018), and firm fundamentals (Wahlen and Wieland, 2011; So, 2013) to show that investors can do better than relying on consensus expectations.

The final aspect of our contribution applies to several strands of literature. First, our insight that different analysts bring different amounts of information to the market advances the debate on whether recommendations are informative (Barber et al., 2001; Jegadeesh et al., 2004; Altinkilic and Hansen, 2009; Loh and Stulz, 2011). Second, our findings limit the extent of the PEAD anomaly's challenge to market efficiency (Fama, 1998), in that the PEAD is restricted to periods of high uncertainty regarding the firm's prospects. Third, our findings on the PEAD and aggregate volatility suggest that high- and low-uncertainty firms should be identified based on the forecast dispersion adjusted for analyst quality differences rather than using the consensus-based dispersion alone (e.g., Diether, Malloy, and Scherbina, 2002; Johnson, 2004; Barron, Stanford, and Yu, 2009).

2. Data and variables

We use the sample of annual earnings per share (EPS) estimates and earnings announcements in I/B/E/S during January 1992-December 2015 for companies with daily return data provided by the Center for Research in Security Prices (CRSP) database.[5] The starting year of 1992 is chosen because some analyses require analysts' recommendation data, which begin in 1993. The earnings estimates and actual earnings are adjusted for splits using the daily cumulative adjustment factor from CRSP (Glushkov and Robinson, 2006).

Each year, we rank analysts following the firm based on the closest absolute forecast error, which is the absolute difference between an analyst's earnings forecast closest to the earnings announcement (made at least a day prior to the announcement day) and the actual earnings, divided by the share price at the beginning of the calendar year. We define HQ (LQ) analysts based on whether their absolute closest forecast error for the firm-year is below (above) the median absolute forecast error for the firm-year. The firm-year ranking measure is directly suited for our study compared to various "star" analyst, cross-firm classifications (e.g., the Wall Street Journal's "Best on the Street" or the Thomson Reuters StarMine's "Top Earnings Estimators") in that it has the cross-firm persistence property similar to the star analyst classifications, as we find below, but it is also simpler (uses only analysts' prior accuracy in a given firm) and preserves the sample size because many firms are not followed by such star analysts. In the robustness section, we conduct sensitivity tests, whose results indicate that our findings are not affected by different cutoff values or forecast weights in the definitions of HQ and LQ analysts. It is important to note that advancing the literatures on measuring analysts' quality, determining the subset of superior analysts, or analyzing their relative performance is not among our study's objectives. Instead, we group analysts by their quality with as simple, realistic, and robust alternative to the consensus, and our contribution is to examine the economic implications of investors' ignoring variation in analyst quality and using the consensus data instead.

[5] We focus on annual rather than quarterly earnings for two main reasons. First, fewer analysts provide quarterly forecasts than annual forecasts. Second, annual earnings announcements are typically more informative, including that they are more often supplemented with a conference call and followed by recommendation changes.

From the initial sample, we generate 861,349 firm-year-analyst rankings based on the closest forecast error; the number drops to 804,003 observations once we require firms to have Compustat data. Next, to avoid small sample bias in our ranking when the number of analysts following the firm is small, we exclude firm-years with fewer than four analysts following, thereby reducing the sample to 750,295 observations. In addition, for all analyses except those using stock recommendations (Tables 4 and 8), analysts must appear in the data in two consecutive years for a given annual announcement to be ranked based on the first of the two years, reducing the sample to 485,815 observations.

In the firm-level regressions, we control for the following firm characteristics: size, the annual stock return, the book-to-market ratio, the number of analysts following, and leverage. *Firm size* is the market value of the firm's equity at the end of the month prior to the earnings announcement month. The *annual stock return* is measured based on monthly equity returns in the 12 months prior to the earnings announcement month. The *book-to-market ratio* is computed as stockholder equity minus preferred stock plus deferred taxes at the end of the fiscal year for which the earnings are announced and divided by firm size. The *number of analysts* is the number of analysts who made at least one earnings forecast for the given announcement. *Leverage* is the book value of total liabilities divided by total assets at the end of the fiscal year for which the earnings are announced. Some of the regression models also control for analyst characteristics: (i) *overall tenure* is the number of years since the analyst first appeared in the I/B/E/S file; (ii) *firm-specific tenure* is the number of years since the analyst began covering the company in the I/B/E/S file; (iii) *brokerage house size* is the number of analysts employed by the brokerage firm; and (iv) *firm coverage* is the number of firms covered by the analyst.

In the models predicting market returns and volatility, most of the controls we use follow Li, Ng, and Swaminathan (2013) and are measured in the month prior to the dependent variable's month. The *earnings-to-price ratio* and *dividend-to-price ratio* are calculated from the S&P 500 dividend, earnings, and price data on Robert Shiller's website.[6] The one-month *T-bill rate* and *30-year Treasury yield* are obtained from CRSP. The *term spread* is the difference between the AAA-rated corporate bond yields obtained from the Federal Reserve Bank of St. Louis (FRED) database and the one-month T-bill yield. The *default spread* is the difference between the BAA and AAA corporate bond yields for the last day of the month when both BAA and AAA daily yields exist, obtained from the FRED. *Inflation* is the change in the consumer price index (CPI; all urban consumers, monthly, non-seasonally adjusted) obtained from the FRED. Following Da, Engelberg, and Gao (2015), our regressions of the VIX also control for the perceived economic policy uncertainty (EPU), which is a news-based measure provided by Baker, Bloom, and Davis (2016). *EPU change* is the percentage change in the monthly average of the daily EPU for the month prior to the dependent variable's month. The VIX is from Wharton Research Data Services (WRDS).

3. Individual analysts

3.1. Sample description and tests of persistence in forecasting performance

Our preliminary analysis determines whether HQ and LQ analysts' forecast accuracy is persistent, which is the prerequisite for using groups of analysts ranked according to our quality measure as alternatives to the consensus. The unit of observation in all analyses in this section is the individual analyst. Figure 1 shows the mean absolute forecast errors of HQ and all analysts for each day during the 300 days prior to the earnings announcement. We observe acceleration in the reduction of the mean forecast error around quarterly earnings announcements at the -90, -180, and -270 day marks. The graph shows that the mean absolute forecast error of all analysts is higher than the mean absolute forecast error of HQ analysts in all days prior to the earnings announcement. This result indicates that the analyst ranking is persistent in the time dimension

[6] Available at http://www.econ.yale.edu/~shiller/data.htm.

because the ranking is based on the previous year's accuracy. The mean absolute forecast errors of All and HQ analysts decrease over time to approximately 0.012 and 0.0115, respectively, one day before the earnings announcement. This difference of 4.17% ($\frac{0.0115}{0.0120} - 1$) is economically meaningful and statistically significant (p-value<0.01). Notably, HQ analysts' accuracy 30 days before the announcement is already higher than all analysts' accuracy at the announcement.

Table 1 reports the summary statistics for analysts ranked into the LQ and HQ groups and examines the persistence of their forecast accuracy. In Panel A, we find that relative to LQ analysts, HQ analysts tend to be more experienced overall and within individual firms, be employed by larger brokerage firms, and cover a greater number of firms. To analyze the univariate persistence of analysts' forecast accuracy, we compare HQ and LQ analysts' forecast errors at the earnings announcement, i.e., a year after they were ranked. The absolute forecast errors of HQ analysts remain smaller than those of LQ analysts: the difference is 9% (0.0081/0.0089) and statistically significant. In the last line of Panel A, we find that both HQ and LQ analysts' forecasts exhibit approximately equal magnitudes of optimism bias; the average forecast errors are significantly different from zero, with untabulated p-values<0.01, indicating that HQ analysts' greater forecast accuracy does not appear to be associated with a more positive forecast bias (Lim, 2001).

The regression analysis in Panel B of Table 1 examines the persistence in the quality classification of analysts over time (columns (1)-(4)) and across firms (columns (5) and (6)). In the probit models in columns (1) and (2), the dependent variable is the HQ status in year t, which equals one if the analyst is categorized as HQ and zero otherwise. In columns (3) and (4), the dependent variable is the absolute forecast error, a continuous variable that allows us to control for firm fixed effects in the regression. In columns (1) and (3), we control for firm characteristics, and in columns (2) and (4), we control for both firm and analyst characteristics. The main coefficient of interest is the HQ classification in year t-1. The results show that the coefficient on *HQ status* *(t-1)* is highly significant (p-value<0.01) in all specifications, indicating that analysts' rankings and forecast accuracy are persistent in consecutive years. For example, the unconditional probability of belonging to the HQ group is approximately 50%, and according to columns (1) and (2), accounting for the HQ status in the previous year increases this likelihood by approximately 4.1%. Columns (3) and (4) show that HQ analysts continue to have lower absolute forecast errors in the following year, with their average absolute forecast error being 8.2% lower (0.00072/0.0085) than the average absolute forecast error for all analysts.

We next conduct cross-firm tests to examine whether forecasting performance is persistent not only over time but also across firms that the analysts follow. Not only is this analysis important in its own right, but affirmative findings will also reinforce the argument that some analysts provide superior information than do others; that is, the HQ designation is not a firm-specific attribute but, rather, a characteristic of the analyst. We define an analyst's performance in other firms as HQ if the analyst is classified in the HQ category for the majority of the other firms that he or she follows during the year (excluding this firm).[7] Columns (5) and (6) of Panel B test whether ranking as an HQ analyst in the other firms in year t-1 can predict an analyst's forecasting performance in year t over and above the HQ classification in year t-1 in the same firm. We estimate two probit models where the dependent variable is the HQ status indicator in a given firm in year t. The independent variables of interest are the HQ status indicator of this analyst in the same firm in year t-1 and the *HQ status in other firms* indicator, which is equal to one if this analyst is also HQ in the majority of other firms that he or she followed in year t-1. We find that analysts who were of HQ in the majority of other firms they followed in year t-1 are 5.1% (p-value<0.01) more likely to be HQ in a given firm in year t. The coefficient on the firm-specific HQ designation

[7] If the number of HQ and LQ rankings of the analyst in the other firms is the same, this analyst-year-firm observation cannot be categorized as either HQ or LQ in the other firms and is thus excluded from this analysis (approximately 9% of the observations).

in year t-1 remains positive and significant (p-value<0.01). Hence, the cross-firm findings suggest that analysts' forecasting performance transcends across the stocks that they follow and, further, that HQ analysts are indeed better than their peers in a persistent manner.

3.2. What does persistence in individual forecasting performance mean for the consensus?

Our finding that individual HQ analysts tend to persistently provide more accurate earnings forecasts than LQ analysts seems to imply that the consensus may be an inferior predictor of future earnings than the average of HQ analysts' forecasts. However, the actual extent of how much the average of HQ analysts' forecasts is more accurate than the consensus forecast in a specific firm may depend on the number of HQ analysts following the firm.[8] Table 2 empirically investigates this issue in Panel A and provides statistical tests comparing the absolute standardized unexpected earnings of consensus (SUE of consensus, equal to the difference between the actual earnings and the average forecast provided by all analysts following the firm, normalized by the stock price) with the absolute SUE of HQ analysts, which is based on the difference between the actual earnings and the HQ analysts' average forecast.[9] We find that as the number of HQ analysts following a firm increases, HQ analysts as a group eventually become more accurate than the consensus, confirming our conjecture. Further, when the number of HQ analysts is four or more, the absolute forecast error of HQ analysts is smaller than the consensus. Therefore, it is in these firms that

---

[8] The intuition is simple—in averaging a greater number of forecasts, the noise portion of their individual forecasts cancels out more, leading to a more precise forecast of the true earnings signal. The greater the number of HQ analysts following the firm is, *ceteris paribus*, the smaller the diversification advantage of the consensus is and more likely investors are to obtain a more accurate forecast by following the average forecast of the HQ analysts than they would via the consensus.

[9] Because some HQ and LQ analysts may stop covering the firm after year t-1 and new, unranked analysts may commence coverage, the numbers of HQ and LQ analysts in year t can become too small or too different from each other (e.g., five HQ and one LQ or vice versa), leading to small sample bias and a lack of robustness when the average accuracies of HQ, LQ, and all analysts as groups are compared in the firm-level analysis. To mitigate this issue, we restrict the sample in all firm-level analyses (Tables 2-7) to firm-years in which the numbers of HQ and LQ analysts are not too different in year t. Specifically, we require that neither of these groups exceeds 75% of all analysts providing forecasts for a given announcement.

investors seeking more accurate earnings forecasts should forego the consensus forecast in favor of the average of HQ analysts' forecasts. For the same reason, our sample consists of firms with four or more HQ analysts when we examine whether the market can exploit differential analyst quality in the analysis of recommendation changes, forecast dispersion, and the PEAD.

In Panel B of Table 2, we report firm characteristics for the samples with fewer than four and four or more HQ analysts in year t, corresponding to firm-years for which the consensus forecast is superior or inferior to HQ analysts' average forecast, respectively. Not unexpectedly, firms with four or more HQ analysts are significantly larger. While these firms represent less than 50% of all firms, they account for 87% of the market value of the full sample. Firms with four or more HQ analysts are more levered and perceived by the market to have a greater growth potential (according to the book-to-market ratio of 0.66, compared to 1.36 for firms with few HQ analysts). This association between firm characteristics and the number of analysts following (or the number of HQ analysts) is not surprising and consistent with prior literature. It is important, however, in the context of our findings, suggesting that the sample for testing the inefficiency of consensus fixation is not dominated by small firms or firms followed by few or inexperienced analysts.

4. Earnings announcements

The previous section demonstrates that relying on HQ analysts' earnings forecasts can generate an earnings forecast that is superior to the consensus forecast. To test whether the market is aware of this empirical regularity, we examine whether the immediate reaction to the earnings surprise based on the mean forecast of HQ analysts is greater than the reaction to the earnings surprise based on the consensus forecast. Table 3 reports the regression results in which the dependent variable is the buy-and-hold cumulative abnormal return (BHAR) for the earnings announcement day and the following trading day based on the four-factor model (Fama and French, 1993; Carhart, 1997). The main variables of interest are the coefficients on the SUE based on the consensus and HQ analysts. Table 3 shows that the reaction to the SUE based on the consensus forecast is greater than the reaction to HQ analysts' SUE, with a highly statistically significant difference between the coefficients of 0.103 based on the chi-squared test in the full sample and a slightly smaller but still highly significant difference of 0.06 in the sample of firms with four or more HQ analysts. The coefficient on HQ analysts' SUE is greater and significantly different from the coefficient on LQ analysts' SUE, which suggests that the market is partially aware of the accuracy differences among analysts.[10] This finding is in line with Kirk, Reppenhagen, and Tucker (2014), who show that the market reacts more strongly to the key analyst than to the least influential analyst following the firm. Importantly, the results indicate that the market does not sufficiently recognize analyst quality differences because its reaction to the consensus forecast is significantly stronger even in columns (4)-(6), where HQ analysts are, on average, more accurate than the consensus.

The finding that the market does not give sufficient weight to HQ analysts' forecasts relative to the consensus may have meaningful economic implications. To gauge their magnitude, we first construct a simple measure of earnings surprise based on the difference between HQ analysts' mean forecast and the consensus forecast, labeled *predicted surprise.* The intuition is to replace the actual earnings in the SUE formula with HQ analysts' mean forecast,

$$Predicted\ surprise = \frac{Avg.Forecast^{HQ} - Avg.Forecast^{consensus}}{Price_{t-1}}, \qquad (1)$$

so that *predicted surprise* can be used to predict the SUE of consensus. Investors aware of the quality differences among analysts would be able to use this measure to predict the immediate market reaction to earnings announcements. Given that HQ analysts are more accurate than the consensus and that the market overweights the consensus forecast when it reacts to earnings

[10] As a technical note, because the set of analysts covering a firm tends to change after the ranking in the previous year, the proportion of HQ and LQ analysts varies across firms and years. Therefore, SUE of consensus is not an average of SUE of HQ and SUE of LQ analysts.

surprise, one can expect positive or negative abnormal returns to the earnings announcement when the mean forecast of HQ analysts is greater or smaller than the consensus, respectively. A simple hedging strategy assessing the predictability of the immediate reaction to earnings announcements is to buy the stock at the market close on the day before the announcement when the predicted surprise is positive and to short it when it is negative. Because the market is focused on the consensus, it will be surprised the most when HQ analysts' average forecast is different the most from the consensus. Next, since the consensus includes new analysts, whose quality is currently unknown because they started to follow the firm after the last year's ranking, the difference between HQ and LQ analysts' forecasts measures analyst heterogeneity not priced because of consensus fixation. Therefore, we also consider a strategy using an alternative definition for *predicted surprise* based on the normalized difference between HQ and LQ analysts' mean forecasts:

$$Predicted\ surprise = \frac{Avg.Forecast^{HQ} - Avg.Forecast^{LQ}}{Price_{t-1}} \qquad (2)$$

The potential returns that investors can earn using these alternative predicted surprise measures represent the returns foregone for investors if they use the surprise based on the difference between the consensus and actual earnings, thereby not incorporating analysts' quality differences into prices.

We report the empirical results in Table 4. The analysis is based on two variations of the signal based on *predicted surprise*: the *positive predicted surprise* and *big positive predicted surprise* indicators. *Positive predicted surprise* is equal to one if *predicted surprise* is positive and zero otherwise. A stronger signal, the *big positive predicted surprise* indicator is one (zero) if *predicted surprise* is above (below) the median of its positive (negative) values in the previous year, and is set to missing otherwise. Using the values of *predicted surprise* measured in the previous year ensures that our analysis is out-of-sample. We regress the two-day cumulative

BHAR on each of these indicators and control variables. The coefficients on the predicted surprise indicators are positive and significant in all specifications, reaching 0.0019 in column (3), and the statistical significance of the predicted surprise indicators is greater for the definition based on the difference between HQ and LQ analysts' forecasts. The last line of the table reports the two-day abnormal returns of a hedging strategy that is long if the predicted surprise indicator of that column is equal to 1 and short if it is equal to 0. All returns are statistically significant and reach 0.24% for *big positive predicted surprise* based on the difference between HQ and LQ analysts' forecasts. These returns can be high enough relative to the transaction costs (Novy-Marx and Velikov, 2016) because *predicted surprise* achieves its highest values when HQ analysts are most accurate, i.e., in firms followed by many analysts (according to Table 2), implying relatively small transaction costs for these larger firms.

The overall conclusion from Tables 3 and 4 is that the market seems to overreact to the actual earnings' deviations from the consensus, compared to deviations from HQ analysts' average estimate. This fixation on the consensus forecast presents an opportunity cost for investors and prevents stock prices from efficiently reflecting the available information.

5. Stock recommendations, forecast dispersion, and implications for the PEAD

The persistence in analysts' forecasting performance over time and across stocks suggests that HQ analysts have superior ability; thus, it is possible that they issue superior stock recommendations compared to the consensus recommendation based on all recommendations for the firm.[11] Further, given that HQ analysts are better at forecasting future earnings, the dispersion in their forecasts is likely to contain more relevant information than the dispersion of the forecasts

[11] A relation between earnings forecasts and recommendations has been examined in the prior literature (e.g., Francis and Soffer, 1997; Bradshaw, 2004; Loh and Mian, 2006; Ertimur, Sunder, and Sunder, 2007).

of the entire set of analysts following the firm. In this section, we empirically examine these predictions.

5.1. Stock recommendations

To gauge whether investors can benefit from being aware of the differences in analysts' forecasting ability, we examine how future returns are associated with the recommendation revisions of different subsets of analysts. If investors internalize and act on the quality differences among analysts, which would imply that they are not fixated on the consensus recommendation, then no relation between future returns and recommendation revisions should obtain regardless of the quality of the analysts making the recommendations.

A recommendation is an integer between 1 and 5, where 1 is "strong buy", 5 is "strong sell", and 3 is "hold". For ease of interpretation, we measure recommendation revisions as the negative of the current recommendation of the analyst minus the previous recommendation of the analyst; thus, a positive recommendation revision is an upgrade. The recommendation revision for the firm is the average of individual analysts' revisions. The sample consists of the recommendation revisions made during the month of the annual earnings announcement. The earnings announcement month has several advantages making it the best time frame to examine whether the market efficiently incorporates its knowledge of analyst quality into its reaction to recommendations. First, the month with the annual announcement has the most information for analysts to process during the year because the information in earnings announcements has a major influence on recommendation revisions (Yezegel, 2015). Second, analysts of different quality types face the same information set when making recommendations that month, in contrast to recommendations at random dates during the year. Finally, and perhaps most importantly, it is the earnings announcement month that reveals that the market is fixated on the consensus forecast and does not recognize superior HQ analysts; thus, we expect this pattern to be prominent for HQ analysts' stock recommendations during this month, as well.

We first examine the immediate market reaction to recommendation revisions, which is inherently an individual analyst-level rather than analyst group-level event, i.e., there is no consensus recommendation to consider in the analysis of the immediate reaction. The untabulated regressions of the immediate market reaction to HQ and LQ analysts' individual recommendation revisions yield results that are consistent with the finding for earnings announcements: investors recognize, at least to an extent, the more accurate forecasters by reacting more strongly to the recommendation revisions of HQ analysts relative to those of LQ analysts. However, the important question remains whether the market fully incorporates quality differences into prices or is fixated on the consensus recommendation.

Table 5 reports the results concerning a delayed response to recommendation revisions by different analyst types. For each calendar month with an annual earnings announcement by a given firm, we average all, HQ, and LQ analysts' recommendation revisions and regress the next months' stock return on these recommendation revision averages. This analysis predicting equity returns one calendar month ahead enables us to use all revisions during the current month because, by the end of that month, investors have learned and updated the analyst quality classification based on this new earnings announcement.[12] The investment delay from the revision date to the end of the revision month provides investors with sufficient time to react to the revision and, because such a

[12] The updated analyst ranking also allows all the revisions to be used because all analysts following the firm are ranked as of the end of the announcement month, i.e., there are no revisions by unranked, newcomer analysts who have just started following the firm between the earnings announcement and the end of the month. The sample has exactly 50% HQ and 50% LQ analysts; thus, the requirement for earnings announcements that at most 75% of forecasts are made by one analyst type is not binding for recommendation revision analyses.

delay reduces the market response the next month, leads to conservative return estimates for the next month (Barber et al., 2001).[13]

The regression results in columns (1)-(3) of Table 5 reveal that the coefficients on the consensus and HQ recommendation revisions are positive and significant. The greater coefficient for HQ analysts and not significant coefficient for LQ analysts imply that price momentum following consensus recommendation revisions (e.g., Jegadeesh et al., 2004) is driven by HQ analysts' recommendation revisions.[14] This supports the argument in the literature that analyst recommendations bring new information to the market (Jegadeesh et al., 2004; Loh and Stulz, 2011).

To examine whether investors should follow the HQ recommendation revision rather than the consensus recommendation revision, we regress the next month's return on the difference between HQ analysts' and consensus recommendation revisions. The intuition for this test is that the gain for investors is the greatest when HQ analysts are relatively more informed, i.e., when their average recommendation is more different from the consensus recommendation. The regression result in column (4) shows that following HQ analysts' upgrades or downgrades net of the consensus revision leads to higher out-of-sample returns for investors.

To further assess the additional value inherent in incorporating analyst quality differences into investment decisions, we examine the returns of a long-minus-short strategy in the month following the revisions, where the long (short) position is in the firms for which the consensus, HQ, LQ, or HQ minus the consensus mean recommendation revision variable is positive (negative) during the earnings announcement month. The last two lines of Table 5 report both the event-time

[13] The results for predicted monthly returns in Table 5 are unaffected by using the subsamples of recommendation revisions made either before or after the earnings announcement during that calendar month or coinciding with the announcement's two-day window. We also reach the same outcome by conducting event-time analysis using returns over the periods (2,32) and (2,62) days following the revision.

[14] Loh and Mian (2006) and Ertimur, Sunder, and Sunder (2007) contrast market responses to the recommendations of more and less accurate analysts based on a partially contemporaneous rather than predictive relation between the recommendations of the two quality groups and stock returns.

average long-short return and the four-factor alpha based on the calendar-time approach for this strategy. In particular, with HQ analysts' recommendation revisions, the resulting one-month strategy return is 0.85% and highly statistically significant, in contrast to the consensus recommendation revisions, for which the hedging strategy yields a statistically nonsignificant 0.36%. For the HQ net of the consensus trading signal, the return is even higher, at 1.24% per month. The calendar-time approach finds that investors achieve a significant alpha of 1.06% (12.72% annualized) when they rely on HQ analysts' recommendation revision net of the consensus recommendation revision.

Overall, the predictive relation between analyst recommendation revisions and equity returns in the subsequent month is driven by HQ analysts' mean recommendation rather than the consensus recommendation. Hence, our findings suggest that analyst quality based on earnings forecasts generalizes to recommendation revisions and that investors can earn significant additional returns by taking into account analyst quality differences instead of relying on the consensus recommendation. These conclusions are entirely consistent with the notion that treating all analysts as equal and, thus, using the consensus measure of analyst output can lead to inefficient pricing.

5.2. Analysts' forecast dispersion and volatility

Analysts' forecast dispersion has been widely used as a proxy for uncertainty about firms' future prospects. We conjecture that just as HQ analysts' superior earnings forecasts and recommendations indicate that they have superior information concerning firm value, these analysts' forecast dispersion similarly contains more accurate information about future uncertainty. We examine whether the disagreement about the firm's prospects among HQ analysts is a superior predictor of uncertainty surrounding the firm's future performance, measured by future return volatility, relative to the disagreement among all analysts, which we call the consensus dispersion.

Table 6 reports the results of regressing stock return volatility during the month following the earnings announcement month on the standard deviation of analysts' forecasts before the earnings announcement. To avoid stale forecasts and to make forecasts more comparable in terms of their proximity to the announcement, we use only the forecasts in the 60 days prior to the announcement.[15] We consider the dispersion of forecasts by all, HQ, and LQ analysts separately, whose coefficients are the variables of interest. HQ analysts' forecast dispersion is statistically significant, while the consensus dispersion is not significant. These findings suggest that only HQ analysts' forecasts capture variation in uncertainty, which is associated with future equity volatility in a given firm. This result extends the earlier findings of a contemporaneous relation between analysts' forecast dispersion, stock volatility, and the option value of the firm (Ajinkya and Gift, 1985). The next two subsections illustrate how investors can benefit from utilizing HQ analysts' dispersion instead of the consensus dispersion.

5.3. Post-earnings announcement drift

Our results concerning the differences in the properties of forecast dispersion of all and HQ analysts' forecasts have important implications for the PEAD anomaly. First, the model proposed by Abarbanell, Lanen, and Verrecchia (1995) predicts that when forecast dispersion is high, investors place less weight on the forecasts relative to their private information, so that investors reduce their immediate response to earnings surprise. Consequently, following the

[15] The length of the measurement period for forecast dispersion significantly varies in the literature. For instance, it can be one month (Krishnaswami and Subramaniam, 1999), four months (Zhang, 2006b), six months (Babenko, Tserlukevich, and Vedrashko, 2012), and up to one year since the previous earnings announcement (Diether, Malloy, and Scherbina, 2002). Our choice of 60 days ensures that we use only the annual earnings forecasts made after the previous quarterly earnings announcement. Our results are unaffected by using a different period length.

announcement, as investors receive more information over time, prices adjust in a manner potentially resulting in a PEAD. Similarly, Zhang (2006a) argues that investor underreaction to public information is more significant when uncertainty is high and finds that analysts' forecast dispersion predicts a price drift following analysts' forecasts.[16] Second, our finding of a relation between HQ analysts' dispersion and return volatility, combined with a relation between return volatility and the PEAD documented in Mendenhall (2004), posit a link between HQ analysts' dispersion and the PEAD. For the full set of analysts, this relation is expected to be weaker because of a weaker relation between consensus dispersion and return volatility.

Therefore, we examine whether investors can better exploit the PEAD when the dispersion of HQ analysts' forecasts is high and relative to the consensus dispersion than when they disregard analyst quality differences.[17] We calculate the PEAD using the calendar-time approach. To be able to relate our results to the PEAD literature, we use the consensus earnings surprise measure to assign announcing stocks to a long or short portfolio at the end of each month, depending on whether the earnings surprise (SUE of consensus) is positive or negative, respectively. The stocks are then held in the two portfolios for horizons from 1 to 11 months to avoid overlapping with the next annual earnings announcement. The monthly PEAD is the alpha from regressing the monthly value-weighted portfolio returns on the four Fama-French and Carhart factors.[18] The cumulative PEAD is the monthly alpha for a given horizon (the number of months that the stock is held in the

[16] In related studies, Francis et al. (2007) find a positive relation between uncertainty, which they measure with the unexplained portion of working capital accruals, and the PEAD. Hung, Li, and Wang (2014) find that exogenously reduced information uncertainty leads to a lower PEAD.

[17] For an association between the PEAD and HQ analysts' forecast dispersion to exist, investors do not need to be aware or take into account HQ analysts' forecast dispersion because they can learn about the firm's uncertainty level, including its systematic component, from a variety of sources. HQ analysts' forecast dispersion effectively reflects that information.

[18] We obtain similar results using equal-weighted portfolios. Separately, we also note that because of our sample's requirement that four or more HQ analysts follow the firm, the sample consists of relatively large firms, thereby reducing the likelihood that any findings may be attributed to illiquidity (Sadka, 2006).

long or short calendar-time portfolio) multiplied by the horizon length in months. The resulting relation between forecast dispersion and the PEAD is presented in Figure 2 and Table 7.

Figure 2 reports the cumulative long PEAD portfolio return minus the short PEAD portfolio return for the sample of announcements with high uncertainty, defined as announcements for which HQ analysts' forecast dispersion is greater than the consensus dispersion, the low-uncertainty sample, in which HQ analysts' forecast dispersion is lower than the consensus dispersion, and the full-sample, i.e., regular, PEAD. The high-uncertainty PEAD is clearly greater than the full-sample PEAD, and the low-uncertainty PEAD is below the full-sample PEAD. The largest cumulative long-short abnormal return of the PEAD strategy is 11.2%, reached eight months following high-uncertainty announcements. Table 7 reports the monthly alphas on long, short, and long-minus-short strategies for the subsamples with high- and low-uncertainty announcements. We find that the low-uncertainty PEAD alpha (approximately 60% of the announcements) is not statistically significant except for the 11-month horizon. In contrast, when the forecast dispersion of HQ analysts is greater than the consensus dispersion, the abnormal return of the long-minus-short portfolio is highly statistically significant at almost all horizons. To determine the additional return that investors can earn from this strategy over and above the return from the simple PEAD strategy, we regress the long-short post-announcement raw return for the high-uncertainty portfolio on the four Fama-French and Carhart factors and the long-short post-announcement raw return for the portfolio including all announcements. The resulting monthly alphas for different horizons are reported in column (7). Investors can earn up to 9.4% cumulative alpha (at the eight-month horizon) over and above the regular PEAD alpha. Overall, the PEAD is observed primarily during periods of high information uncertainty, as determined by the forecast dispersions of HQ analysts, which allows investors who are not fixated on the consensus forecast and recognize analyst quality differences to achieve better investment performance.

6. The content of HQ analysts' information output at the aggregate

Our findings that HQ analysts' have superior ability to forecast earnings, issue more informative recommendations, and have better forecasts of volatility have immediate implications for aggregate returns and volatility: we should expect that the superior abilities at the firm level manifest themselves at the aggregate market levels, thereby benefiting investors who overcome their consensus fixation. Accordingly, we average the changes in recommendations of all, HQ, and LQ analysts across all firms that made an annual earnings announcement this month and examine whether the consensus or HQ analysts' average recommendation revisions predict future market returns. The aggregation argument works similarly for forecast dispersion, in that aggregating HQ analysts' dispersion across all firms in the sample should result in a dispersion measure reflecting the degree of uncertainty in the market.[19]

In Table 8, Panel A, we report the estimation results for the relation between revisions in stock recommendations and future market returns. To that end, each month, we calculate a simple average of HQ, LQ, and all analysts' recommendation changes across all firms that announce annual earnings, resulting in one monthly time series for the average recommendation change in the market. The dependent variable is the value-weighted market return in the month following the recommendation change month. The *consensus mean recommendation revision* is the mean of all recommendation changes during the month in which the firm's earnings are announced. The *HQ (LQ) mean recommendation revision* is an analogous variable that is based only on recommendation changes of HQ (LQ) analysts. The *HQ-consensus mean recommendation revision* is equal to the difference between the HQ analysts' and consensus mean recommendation revisions when the HQ analysts' mean revision is either positive and greater or negative and

[19] Notably, the aggregation principle does not assume and the aggregate findings do not indicate that individual HQ analysts have superior macroeconomic knowledge or ability to predict market-level developments (e.g., Hutton, Lee, and Shu, 2012).

smaller than the consensus revision; otherwise, this variable is zero. The control variables follow Li, Ng, and Swaminathan (2013) and also include the previous month's market return to account for return momentum.

The regression results reveal that the coefficient on the consensus recommendation revision is positive and weakly statistically significant, indicating that the information in recommendation revisions is not fully internalized by the market. However, this coefficient's significance is entirely due to the mean recommendation revision by all HQ analysts because LQ analysts' coefficient is not statistically significant. The coefficients on the HQ revision net of the consensus revision in columns (4) and (8) are greater than the other coefficients, suggesting that when HQ analysts are comparatively more informed and, thus, their recommendations deviate more from the consensus, investors can benefit more from following the mean of HQ analysts' recommendation revisions. To estimate the economic magnitude of this market return predictability, we conduct an out-of-sample trading strategy analysis, reported in the last line of Panel A. The long and short trading signals are based on the historical variation in monthly mean recommendation changes as follows: if the average recommendation revision for a given month is greater (smaller) than the median of the monthly average recommendation revisions over the previous 24 months, i.e., the current recommendation revision is more optimistic (pessimistic) than they were in the recent past, investors should buy (short) the market value-weighted index and hold it for one month.[20] For the variable based on the difference between HQ and consensus mean recommendation revisions in the last column of Panel A, one should buy (short) the market when this variable is positive (negative), i.e., when HQ analysts, on average, upgrade and are more optimistic (downgrade and are more pessimistic) than the consensus mean revision. We benchmark the performance of these

[20] The results are unaffected by selecting a longer historical window of up to five years. The 24-month window that we use minimizes the number of months lost to initialize this out-of-sample analysis, while providing enough observations to robustly calculate the median of monthly mean recommendation changes.

strategies against the market return and report the market model alphas. Only HQ analysts' recommendations produce statistically significant alphas, with a particularly large alpha based on HQ analysts' mean recommendation when it is different from the consensus recommendation, at 0.57% per month (6.84% annualized). These findings suggest that investors can earn economically meaningful returns if they follow the recommendations of a subset of HQ analysts and not the consensus recommendation.

In Panel B of Table 8, we examine another market-level aggregate relation—whether HQ analysts' forecast dispersion aggregated across all firms is associated with systematic uncertainty in the economy captured by current or future market volatility. We use the VIX to measure expected market volatility. To capture the degree to which HQ analysts' uncertainty is different from the consensus (all analysts') uncertainty, we define the dispersion ratio for HQ analysts as the HQ analysts' forecast dispersion divided by the forecast dispersion of all analysts following the firm. The dispersion ratios are analogous to the measures of uncertainty in the PEAD analysis. A bonus feature of HQ and LQ analysts' forecast dispersions being normalized in the dispersion ratios is that they are made comparable across firms and years and can thus be suitably aggregated to measure market uncertainty. A market-level measure of analyst uncertainty is created by firm value-weighting the dispersion ratios across firms each month, and we also include the results for aggregate dispersions for completeness.

Panel B reports the regressions of monthly VIX returns (this month's VIX divided by the last month's VIX, minus one) on the dispersions of all, HQ, and LQ analysts' forecasts and the HQ and LQ analysts' dispersion ratios measured before the earnings announcement in the previous month.[21] The results indicate that a higher dispersion ratio for HQ analysts predicts a greater VIX return in the following month. The LQ dispersion ratio does not capture the level of uncertainty in

[21] We also obtain very similar results using the VIX values in lieu of the returns.

the market, nor do the non-normalized forecast dispersion variables. To estimate the opportunity cost for investors not recognizing the limitation of relying on the consensus forecast dispersion, the last line of Panel B reports the performance of a long-short strategy for the VIX identical in design to the long-short strategy for the recommendations market return index. If the forecast dispersion measure in the corresponding column in a given month is above the median of its monthly values during the previous 24 months, then the strategy is to long the VIX return in the next month. If the forecast dispersion measure is below the median, then the VIX return is shorted. To benchmark the performance of these monthly strategy return series, we regress it on the market return and report the alphas. The only statistically significant (p-value<0.01) abnormal return is for the trading signal based on HQ analysts' dispersion ratio; the strategy yields an economically significant return of 3.32% per month.[22] We conclude that when HQ analysts are, on average, more uncertain than other analysts about the prospects of the firms they follow, investors should expect an increase in market-wide uncertainty.

Together, consistent with the firm-level findings, the aggregate results indicate that the market's fixation on the consensus recommendation and the consensus forecast dispersion results in a delayed incorporation of information into prices, both at the individual firm level and at the market level. Investors who recognize quality differences among analysts can benefit from exploiting this market inefficiency.

7. The definition of high-quality analysts

[22] While this alpha is much higher than the typical alphas for stock-based long-short strategies, it is credible because the market may not be able to fully arbitrage it away. The transaction costs of implementing a VIX strategy with complete end-of-month rebalancing using VIX futures would be much higher than transaction costs for stocks because of a significant contango in the VIX futures market, resulting in costs of approximately 14% per month (Nadig, 2016). Further, the VIX strategy's average monthly raw return and volatility are 2.3% and 19.6%, respectively.

Based on the combined findings of sections 4-6, we reject the null of the economic efficiency of consensus fixation using a simple measure of analyst quality differences. In this section, our robustness analysis examines the sensitivity of our finding of the inefficiency of consensus fixation to the definition of HQ analysts. The firm-year definition used throughout this paper, which splits analysts into two groups at the median based on the accuracy of their closest estimate to the annual EPS announcement, is only one of many ways of ranking analysts and generating an alternative to the consensus forecast. A ranking that, like ours, is based on the last year's forecast accuracy is used, for instance, in Loh and Mian (2006) and supported by Sinha, Brown, and Das (1997), who find it to be superior to rankings based on more years, and Carpenter and Lynch (1999), who find it to be relatively less exposed to survivorship bias.[23] Other ranking methods can include such variations as using different forecast accuracy cutoffs between the two groups and giving different weights to forecasts.[24] We note that our analysis of the inefficiency of consensus fixation does not require finding the subset of analysts who beat the consensus by the

[23] More generally, rankings based on past performance are common for analyst forecast persistence studies (Stickel, 1992; Sinha, Brown, and Das, 1997) and in a number of other areas, such as the mutual fund and pension fund performance forecasting (Hendricks, Patel, and Zeckhauser, 1993; Carhart, 1997; Tonks, 2005) and economic forecasting (Aiolfi and Timmerman, 2006) literatures. Brown (2004) finds that models built on past forecasting performance predict analysts' forecasting accuracy as well as a model based on analysts' individual characteristics, such as that in Clement (1999).

[24] A somewhat more different approach is to rank analysts in a given year by averaging their forecast errors across the firms they follow. This method underlies various "star" analyst rankings, which are cross-firm by design. This alternative ranking procedure would avoid losing the observation of the first year when an analyst begins covering a firm, as we could use the analyst's ranking in other firms. However, the cross-firm ranking approach has several pitfalls. First, an aggregated ranking across firms can be misleading if an analyst's superiority is mainly firm- or industry-specific (e.g., Kadan et al., 2012). Second, with the cross-firm ranking, we end up with some firms followed almost exclusively by either HQ or LQ analysts and even populated by just one analyst quality type, coinciding with the consensus, which would undermine our study's objective because it relies on comparing the average outputs of different analyst types to the consensus in each firm. While the cross-firm ranking approach is not suitable for this study, its ranking characteristics are effectively provided by our firm-year ranking measure: Table 1 shows that an analyst's forecast accuracy in a given firm and the accuracy measure based on all firms covered by the analyst are highly correlated.

greatest margin or whose other outputs result in the greatest additional returns relative to the returns based on the consensus measure.[25]

The first two alternative ranking procedures that we consider in this section define HQ analysts as those in the top 70% (with the HQ/LQ proportion at 70%/30%) and the top 30% (with the HQ/LQ proportion at 30%/70%) of forecast accuracy distribution. Appendix A examines a broad range of possible cutoff values for HQ analysts and finds that analyst forecasting performance is persistent over time for all cutoffs. The third alternative definition for HQ and LQ analysts is computed using all forecasts, rather than the most recent forecast, and gives different weights to forecasts, depending on how long the forecast has been outstanding. The details of this time-weighted analyst quality measure are also provided in Appendix A.

Next, we repeat the key firm-level analyses of the paper with the three alternative quality definitions. If HQ is defined as the top 70% of analysts, then the smallest number of HQ analysts following the firm is three for their average forecast accuracy to be superior to the consensus forecast. In that subsample, the market reaction to the consensus forecast is greater than that for HQ analysts. Importantly, the market does not sufficiently recognize analyst quality differences and is fixated on the consensus forecast, resulting in predictable mispricing at the earnings announcement day. These results for forecasts are also robust to defining the top 30% of analysts as HQ. The results on return and volatility predictability for recommendations and forecast dispersion are little affected by the first two alternative definitions.[26] With the time-weighted analyst quality definition, HQ analysts' average forecast is more accurate than the consensus when four or more analysts follow the firm. The design of the time-weighted measure results in a

[25] For these reasons, this study is not related to a literature in economics and statistics aiming to achieve better forecasting performance than the consensus forecast by proposing different methods of combining forecasts (e.g., Conroy and Harris, 1987; Clemen, 1989; Brown, Gay, and Turac, 2008).

[26] For HQ defined as the most accurate 70% of analysts, firms with three or more HQ analysts constitute approximately 2/3 of all firms, indicating that our findings about the inefficiency of consensus fixation are not limited to large firms.

relatively small sample size, which could explain that while the market reacts more to the consensus than to the average of HQ analysts' forecasts based on this measure, this difference is not statistically significant. However, this result does not mean that the market sufficiently recognizes analyst quality differences because it would then react more strongly to the average of HQ forecasts than it would to the consensus. The replication results of the other tables hold for the time-weighted quality measure as consistently as for the first two alternative cutoff measures. Overall, we conclude that the definition of HQ analysts and the smallest required number of HQ analysts associated with this definition are immaterial for our conclusion about the inefficiency of consensus fixation. More generally, any alternative to the consensus forecast can work for this analysis as long as it is more accurate than the consensus forecast.

8. Conclusion

Investors can benefit from disregarding the consensus and using a signal with greater information content. However, the market does not demonstrate an appropriate awareness of these significant deficiencies in the consensus: the market reacts more strongly to deviations from the consensus earnings forecast than it does to deviations from HQ analysts' forecasts. In addition, HQ analysts' stock recommendations and forecast dispersion predict the first two moments of firm and stock market returns better than the corresponding consensus recommendation and forecast dispersion of all analysts. In short, the persistence of analysts' differential ability along multitude dimensions is not sufficiently recognized by the market, resulting in inefficient pricing after earnings announcements and stock recommendations changes. Overall, our findings suggest that the market's fixation on the consensus forecast is not justified. The forecasts and other information outputs of HQ analysts demonstrably provide superior, more accurate information than the consensus figures.

We show that sell-side research analysts' superior forecasting ability is reflected not only in individual stocks but also at the aggregate market level. This result is particularly surprising, as the limits to arbitrage are much narrower at that level because the aggregate stock market is much more liquid than individual stocks. This finding implies that the fixation on the consensus is not limited to individual investors but affects institutional investors, as well. This conclusion is corroborated by our finding that firms with a greater number of HQ analysts are larger and have greater analyst following, which are associated with greater institutional holdings. Consequently, the question arises whether investors' fixation on the consensus can be remedied. Tracking analyst performance by the investing public is a necessary step towards this objective. News and information provider outlets, both electronic and non-electronic, can also play a major role in this regard by changing how they report and present data on analysts' expectations to help investors circumvent their cognitive constraints and fully take into account the variation in analyst quality.

Appendix A

The appendix provides details on how several alternative classifications of analysts into the HQ and LQ groups affect the persistence in analyst forecasting performance. The ranking procedure sorts analysts in a given firm-year based on their absolute forecast error. In general, HQ analysts are those who are ranked in the top p percent of analysts, while LQ analysts are those in the bottom $(1-p)$ percent. If analysts' forecasting performance were uncorrelated across years, then the fractions of analysts who preserve their ranking in two consecutive years as HQ and LQ would be p^2 and $(1-p)^2$, respectively, or $p^2+(1-p)^2$ of all analysts.

Figure A.1 plots the fraction of analysts who retain their rankings in consecutive years and the expected fraction assuming no performance correlation across years. In this figure, to rank analysts up to the decile on the horizontal axis, the sample is constrained to firms that are followed by ten or more analysts. We find that with all cutoff values of p on the horizontal axis, the actual fraction of persistent forecasting performance is above the expected fraction and that all the differences are statistically significant (p-value<0.01). For example, when we classify the top 10% of analysts following a firm in a given year as HQ (p=10%) and the bottom 90% as LQ, the expected fraction given random assignment is $0.9^2+0.1^2=0.82$. The figure shows that the actual fraction is greater than that at 0.843. The overall finding is that for all of the cutoff values, there is a sizeable persistent component; thus, for accuracy persistence, which exact cutoff value we choose to partition HQ and LQ analysts should make little difference.

The two alternative cutoffs for HQ analysts—the top 70% (so that the LQ analysts are the bottom 30%) of analysts and the top 30% (the LQ analysts are the bottom 70%) of analysts, which are considered in Section 7—are symmetric around the cutoff at the median used throughout this paper. Because analysts are ranked in year t-1, the proportion of HQ and LQ analysts following the firm tends to become different at the year t announcement, potentially resulting in too few HQ or LQ analysts for the firm. To avoid small sample bias, these alternative analyst quality definitions require the same restrictions on the sample as those used in the rest of this paper, with an additional condition. For the robustness analysis, we require that the proportion of HQ and LQ analysts does not change by more than a 20% margin from t-1 to t. The motivation is obvious: for the definition of HQ analysts as the top 30% in year t-1, the additional condition ensures that their fraction always stays above 10% of analysts covering the firm in year t; and if HQ analysts are defined as the top 70% in year t-1, then their fraction can be between 50% and 90% of analysts in year t, while the fraction of LQ analysts never falls below 10% of analysts covering the firm in year t.

The third ranking is based on the time-weighted absolute forecast error, which is computed using all forecasts by the analyst during the 300 days prior to the annual earnings announcement. This measure is similar to the time-weighted bid-ask spread that is commonly used in the market microstructure literature (e.g., McInish and Wood, 1992; Lee, Mucklow, and Ready, 1993; Bessembinder, 1999; Rubin, 2007). For a given announcement and analyst, it is computed as follows:

$$TWFE_t = \frac{FE_{300} \times d_1 + \sum_{j=2}^{n}(FE_j \times d_j)}{300} \tag{3}$$

where $TWFE_t$ is the time-weighted absolute forecast error of the analyst in year *t*; FE_{300} is the absolute forecast error based on the forecast outstanding on the 300th day prior to the earnings announcement; d_1 is the number of days that this forecast is outstanding (from the 300th day prior to the earnings announcement to the earliest of the earnings announcement day or the following earnings forecast revision day); *n-1* is the number of estimates issued by the analyst between the 299th day prior to the earnings announcement and the earnings announcement day; FE_j is the absolute forecast error of forecast *j;* and d_j is the number of days that the forecast has been outstanding. The advantage of the time-weighted measure is that it captures the analyst's ability over a four-quarter period, instead of at a single point just before the annual earnings announcement. However, the measure's disadvantage is that it excludes analysts who have not provided an annual forecast 300 or more days prior to the annual earnings announcement day. The sample therefore shrinks by approximately 65% relative to the sample used throughout this paper, thus both reducing the power of our empirical analysis with this quality measure and possibly biasing the results because it remains unknown whether the lack of early forecasts is due to the analyst's poor ability or a neutral reason, such as common practice in the given industry or firm being analyzed. These disadvantages are among the reasons the analyst classification used in our paper equal-weights the forecasts made by the best performing subset of analysts, which is certainly more common in the literature and analogous to the approach examined, for instance, in Aiolfi and Timmerman (2006). Further, studies focused on the forecast weighting topic find that the simple averaging of expert forecasts is commonly equivalent or more optimal than more sophisticated weighting methods for various economic series (Genre et al., 2013).

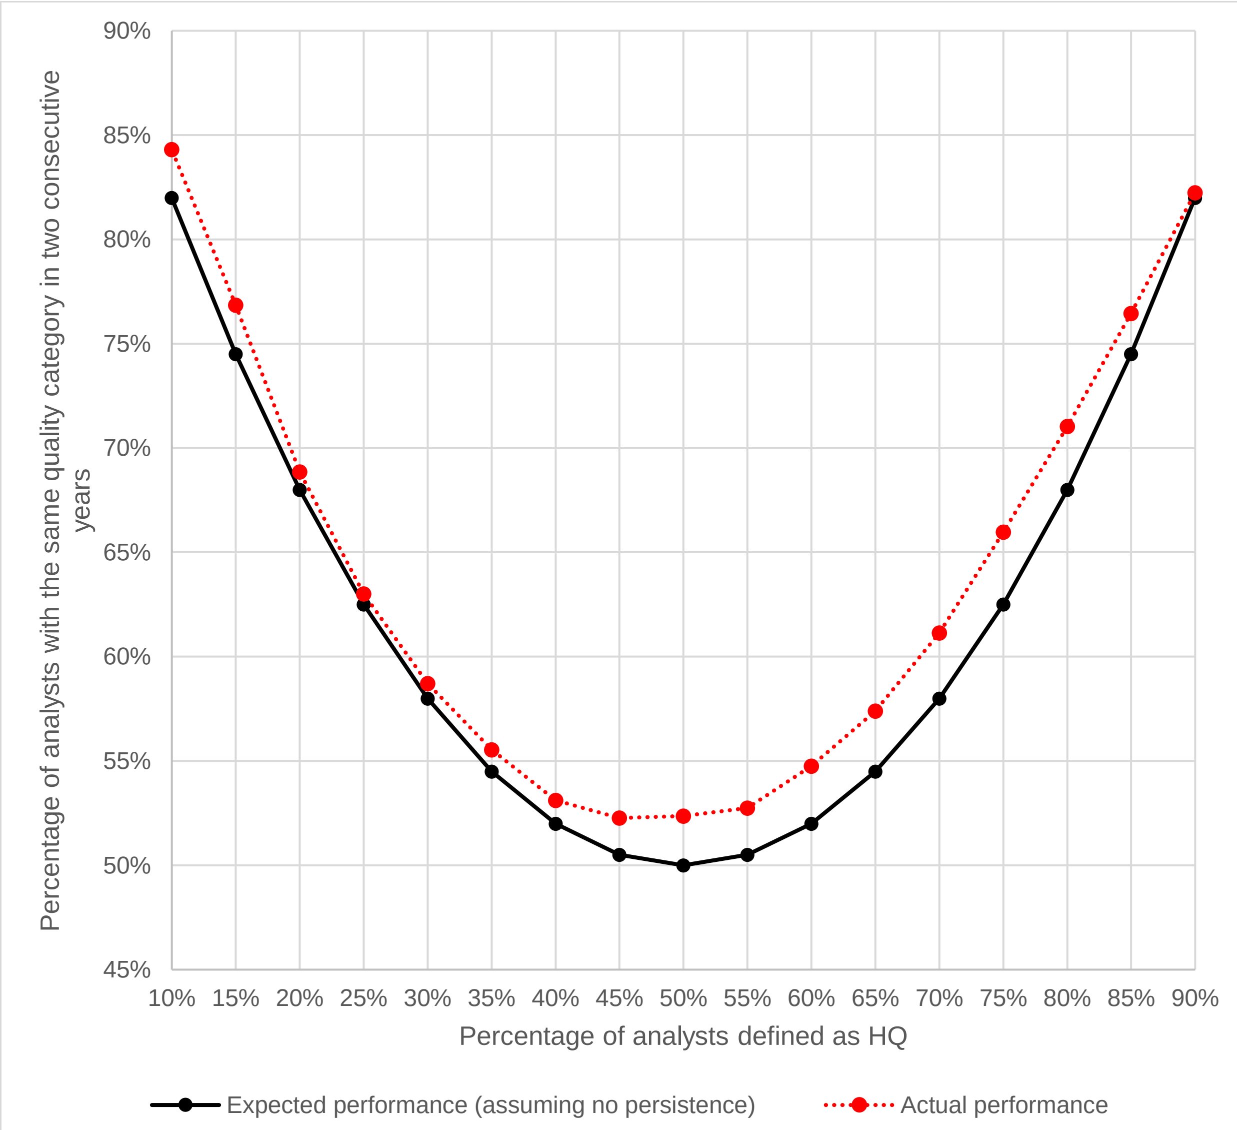

Figure A.1: Persistence in analysts' forecasting performance. This figure depicts how the fraction of analysts retaining their ranking of either high- or low-quality (HQ and LQ) in terms of forecast accuracy in two consecutive years depends on the cutoff percentile in the definition of HQ analysts. HQ analysts are those whose closest absolute forecast errors are less than the absolute forecast error at the cutoff percentile (the horizontal axis) of the distribution of forecast errors for the firm's annual earnings announcement in year t-1. The closest absolute forecast error is the absolute difference between an analyst's forecast estimate closest to the earnings announcement prior to the announcement day and the actual annual earnings, divided by the share price at the beginning of the calendar year. *Expected performance assuming no persistence* is the fraction of analysts who would have the same forecast performance category in two consecutive years if their performance were uncorrelated between years.

References

Abarbanell, J.S., Lanen, W.N. Verrecchia, R.E., 1995. Analysts' forecasts as proxies for investor beliefs in empirical research. Journal of Accounting and Economics, 20(1), 31-60.

Aiolfi, M., Timmermann, A., 2006. Persistence in forecasting performance and conditional combination strategies. Journal of Econometrics, 135(1), 31-53.

Ajinkya, B.B., Gift, M.J., 1985. Dispersion of financial analysts' earnings forecasts and the (option model) implied standard deviations of stock returns. The Journal of Finance, 40(5), 1353-1365.

Altinkiliç, O. and Hansen, R.S., 2009. On the information role of stock recommendation revisions. Journal of Accounting and Economics, 48(1), 17-36.

Babenko, I., Tserlukevich, Y., Vedrashko, A., 2012. The credibility of open market share repurchase signaling. Journal of Financial and Quantitative Analysis, 47(5), 1059-1088.

Baker, S.R., Bloom, N. and Davis, S.J., 2016. Measuring economic policy uncertainty. The Quarterly Journal of Economics, 131(4), 1593-1636.

Barber, B., Lehavy, R., McNichols, M., Trueman, B., 2001. Can investors profit from the prophets? Security analyst recommendations and stock returns. The Journal of Finance, 56(2), 531-563.

Barron, O.E., Stanford, M.H. and Yu, Y., 2009. Further evidence on the relation between analysts' forecast dispersion and stock returns. Contemporary Accounting Research, 26(2), 329-357.

Bessembinder, H., 1999. Trade execution costs on Nasdaq and the NYSE: A post-reform comparison. Journal of Financial and Quantitative Analysis, 34(3), 387-407.

Bradshaw, M.T., 2004. How do analysts use their earnings forecasts in generating stock recommendations? The Accounting Review, 79(1), 25-50.

Brown, L.D., 2004. How important is past analyst forecast accuracy? Financial Analysts Journal 57, 44-49.

Brown, L.D., Gay, G.D. and Turac, M., 2008. Creating a "smart" conditional consensus forecast. Financial Analysts Journal, 64(6), 74-86.

Buraschi, A., Piatti, I., Whelan P., 2017. Expected Term Structures, University of Chicago working paper.

Carhart, M.M., 1997. On persistence in mutual fund performance. The Journal of Finance, 52(1), 57-82.

Carpenter, J.N., Lynch, A.W., 1999. Survivorship bias and attrition effects in measures of performance persistence. Journal of Financial Economics, 54(3), 337-374.

Chiang, Chin-Han; Fan, Jianqing; Hong, Harrison; Tu, Jun, 2018. Robust measures of earnings surprises. Journal of Finance, forthcoming.

Clemen, R.T., 1989. Combining forecasts: A review and annotated bibliography. International Journal of Forecasting, 5(4), 559-583.

Clement, M.B., 1999. Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? Journal of Accounting and Economics, 27(3), 285-303.

Consensus earnings estimates, Special 12-page report, April 2013, IR Magazine, www.InsideInvestorRelations.com

Conroy, R. and Harris, R., 1987. Consensus forecasts of corporate earnings: Analysts' forecasts and time series methods. Management Science, 33(6), 725-738.

Cooper, R.A., Day, T.E. and Lewis, C.M., 2001. Following the leader: a study of individual analysts' earnings forecasts. Journal of Financial Economics, 61(3), 383-416.

Da, Z., Engelberg, J. and Gao, P., 2015. The sum of all fears investor sentiment and asset prices. Review of Financial Studies, 28(1), 1-32.

De Franco, G., Zhou, Y., 2009. The performance of analysts with a CFA designation: the role of human-capital and signaling theories. The Accounting Review, 84(2), 383-404.

Diether, K.B., Malloy, C.J., Scherbina, A., 2002. Differences of opinion and the cross section of stock returns. The Journal of Finance, 57(5), 2113-2141.

Ertimur, Y., Sunder, J., Sunder, S.V., 2007. Measure for measure: The relation between forecast accuracy and recommendation profitability of analysts. Journal of Accounting Research, 45(3), 567-606.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3-56.

Fang, L.H. and Yasuda, A., 2014. Are stars' opinions worth more? The relation between analyst reputation and recommendation values. Journal of Financial Services Research, 46(3), 235-269.

Francis, J., Lafond, R., Olsson, P., Schipper, K., 2007. Information uncertainty and post-earnings-announcement-drift. Journal of Business Finance and Accounting, 34(3-4), 403-433.

Francis, J. and Soffer, L., 1997. The relative informativeness of analysts' stock recommendations and earnings forecast revisions. Journal of Accounting Research, 35(2), 193-211.

Glushkov, D., Robinson, D., 2006. A Note on IBES Unadjusted Data, WRDS Documentation on IBES, https://wrds-web.wharton.upenn.edu/wrds/support/Data/_001Manuals%20and%20Overviews/_003I-B-E-S/index.cfm.

Genre, V., Kenny, G., Meyler, A., Timmermann, A., 2013. Combining expert forecasts: can anything beat the simple average? International Journal of Forecasting, 29(1), 108-121.

Hendricks, D., Patel, J., Zeckhauser, R., 1993. Hot hands in mutual funds: Short-run persistence of relative performance, 1974–1988. The Journal of Finance, 48(1), 93-130.

Hirshleifer, D., Lim, S., Teoh, S. H., 2009. Driven to distraction: Extraneous events and underreaction to earnings news. Journal of Finance 64, 2289-2325.

Hirst, D.E., Hopkins, P.E., Wahlen, J.M., 2004. Fair values, income measurement, and bank analysts' risk and valuation judgments. The Accounting Review, 79(2), 453-472.

Hung, M., Li, X., Wang, S., 2014. Post-earnings-announcement drift in global markets: Evidence from an information shock. Review of Financial Studies, 28 (4), 1242-1283.

Hutton, A.P., Lee, L.F., Shu, S.Z., 2012. Do managers always know better? The relative accuracy of management and analyst forecasts. Journal of Accounting Research, 50(5), 1217-1244.

Jacob, J., Lys, T.Z., Neale, M.A., 1999. Expertise in forecasting performance of security analysts. Journal of Accounting and Economics, 28(1), 51-82.

Jegadeesh, N., Kim, J., Krische, S.D., Lee, C., 2004. Analyzing the analysts: When do recommendations add value? The Journal of Finance, 59(3), 1083-1124.

Johnson, T. 2004. Forecast dispersion and the cross-section of expected returns. The Journal of Finance 59 (5): 1957–78.

Kadan, O., Madureira, L., Wang, R. and Zach, T., 2012. Analysts' industry expertise. Journal of Accounting and Economics, 54(2-3), 95-120.

Kirk, M. P., Reppenhagen, D.A., Tucker, J.W., 2014. Meeting individual analyst expectations, Accounting Review 89, 2203–2231.

Krishnaswami, S., Subramaniam, V., 1999. Information asymmetry, valuation, and the corporate spin-off decision. Journal of Financial Economics, 53(1), 73-112.

Kucheev, Y.O., Ruiz, F. and Sorensson, T., 2017. Do Stars Shine? Comparing the Performance Persistence of Star Sell-Side Analysts Listed by Institutional Investor, the Wall Street Journal, and StarMine. Journal of Financial Services Research, 52(3), 277-305.

Lee, C.M., Mucklow, B. and Ready, M.J., 1993. Spreads, depths, and the impact of earnings information: An intraday analysis. Review of Financial Studies, 6(2), 345-374.

Li, Y., Ng, D.T., Swaminathan, B., 2013. Predicting market returns using aggregate implied cost of capital. Journal of Financial Economics, 110(2), 419-436.

Lim, T., 2001. Rationality and analysts' forecast bias. The Journal of Finance, 56(1), 369-385.

Lin, H.W., McNichols, M.F., 1998. Underwriting relationships, analysts' earnings forecasts and investment recommendations. Journal of Accounting and Economics, 25(1), 101-127.

Loh, R.K., Mian, G.M., 2006. Do accurate earnings forecasts facilitate superior investment recommendations? Journal of Financial Economics, 80(2), 455-483.

Loh, R.K., Stulz, R.M., 2011. When are analyst recommendation changes influential? Review of Financial Studies, 24(2), 593-627.

Ludvigson, S.C., 2004. Consumer confidence and consumer spending. Journal of Economic Perspectives, 18(2), 29-50.

Maines, L.A., McDaniel, L.S., Harris, M.S., 1997. Implications of proposed segment reporting standards for financial analysts' investment judgements. Journal of Accounting Research, 35, 1-24.

Malloy, C.J., 2005. The geography of equity analysis. Journal of Finance, 60(2), 719-755.

McInish, T.H. and Wood, R.A., 1992. An analysis of intraday patterns in bid/ask spreads for NYSE stocks. Journal of Finance, 47(2), pp.753-764.

Mendenhall, R.R., 2004. Arbitrage risk and post-earnings-announcement drift. Journal of Business 77, 875-894.

Michaely, R., Womack, K.L., 1999. Conflict of interest and the credibility of underwriter analyst recommendations. Review of Financial Studies, 12(4), 653-686.

Mikhail, M.B., Walther, B.R., Willis, R.H., 1997. Do security analysts improve their performance with experience? Journal of Accounting Research, 35, 131-157.

Nadig, D., 2016. Understanding VIX ETFs: Careful What You Wish For. http://www.etf.com/sections/blog/understanding-vix-etfs, September 2.

Novy-Marx, R., Velikov, M. 2016. A taxonomy of anomalies and their trading costs. Review of Financial Studies 29, 104-147.

Rubin, A., 2007. Ownership level, ownership concentration and liquidity. Journal of Financial Markets, 10(3), 219-248.

Rubin, A., Segal, B., Segal D., 2017. The interpretation of unanticipated news arrival and analysts' skill, Journal of Financial and Quantitative Analysis 52(4), 1491-1518.

Sadka, R., 2006. Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. Journal of Financial Economics 80, 309-349.

Sinha, P., L. Brown, Das S., 1997. A re-examination of financial analysts' differential earnings forecast ability. Contemporary Accounting Research 14, 1-42.

So, E., 2013. A new approach to predicting analyst forecast errors: Do investors overweight analyst forecast? Journal of Financial Economics 108, 615-640.

Sorescu, S., Subrahmanyam, A., 2006. The cross section of analyst recommendations. Journal of Financial and Quantitative Analysis, 41(1), 139-168.

Stickel, S., 1992. Reputation and performance among security analysts. Journal of Finance 47, 1811-36.

Tatler, B.W., 2007. The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. Journal of Vision, 7(14):4, 1-17.

Tonks, I., 2005. Performance persistence of pension-fund managers. The Journal of Business, 78(5), 1917-1942.

Wahlen, J.M. and Wieland, M.M., 2011. Can financial statement analysis beat consensus analysts' recommendations? Review of Accounting Studies, 16(1), 89-115.

Yezegel, A., 2015. Why do analysts revise their stock recommendations after earnings announcements? Journal of Accounting and Economics, 59(2), 163-181.

Zhang, X., 2006a. Information uncertainty and stock returns. Journal of Finance, 61(1), 105-137.

Zhang, X., 2006b. Information uncertainty and analyst forecast behavior. Contemporary Accounting Research, 23(2), 565-590.

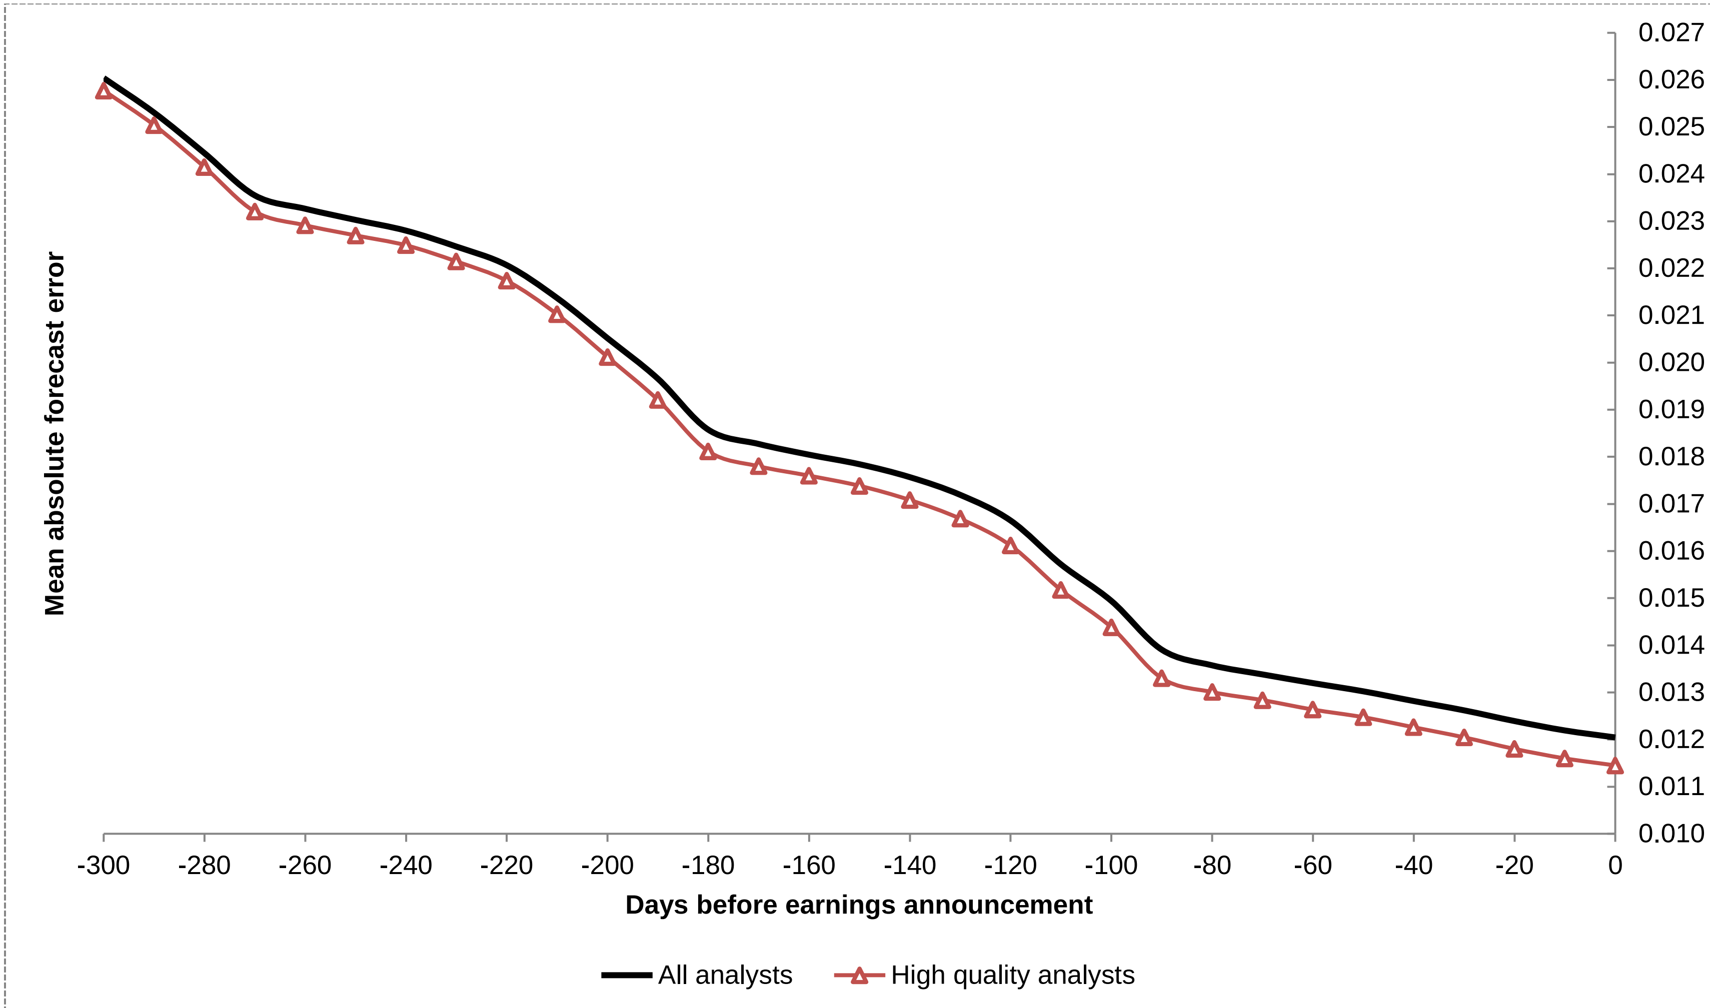

Figure 1: Absolute forecast errors up to 300 days before the annual earnings announcement day. The absolute forecast error is the absolute difference between an analyst's forecast estimate and the actual annual earnings, divided by the share price at the beginning of the calendar year. The absolute forecast error at day -t prior to the earnings announcement date is calculated as the mean absolute forecast error based on all annual earnings forecasts outstanding for the firm as of that day, averaged across the firm's announcements and then averaged across firms. HQ analysts are those whose closest absolute forecast errors are below the median closest absolute forecast error for the earnings announcement in the previous year, where the closest absolute forecast error is that closest to the announcement day prior to that day.

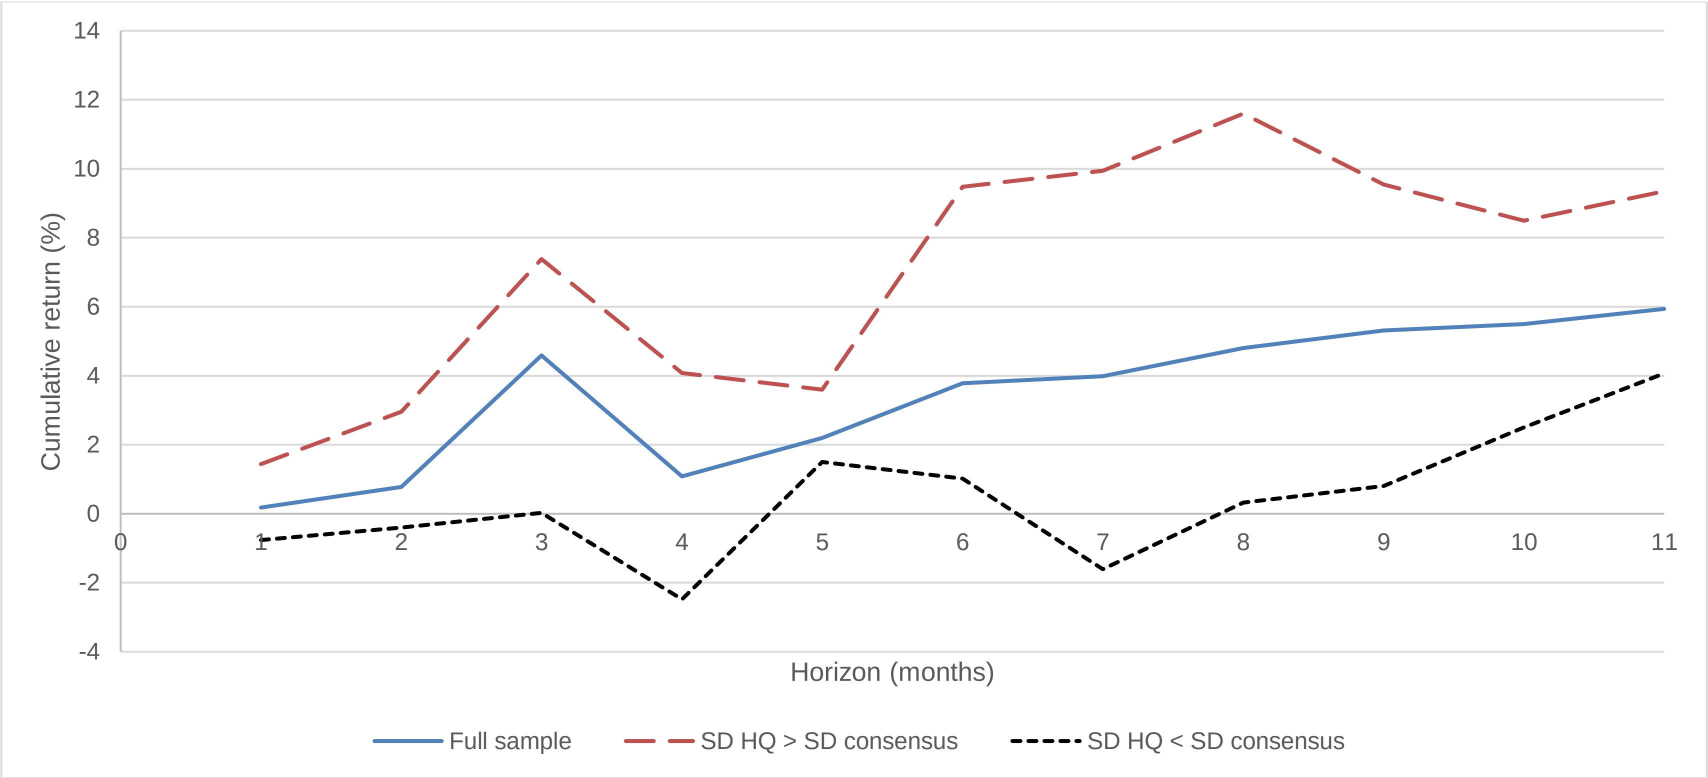

Figure 2: Cumulative post-announcement drifts and analysts' uncertainty

This figure shows the cumulative returns for 1- to 11-month horizons following earnings announcements. Each month, stocks enter a calendar-time long (short) portfolio, depending on whether their earnings surprise is positive (negative), where earnings surprise is defined based on the consensus estimate in Table 2. The horizontal axis is the drift's horizon, which is the number of months a stock is held in the calendar-time portfolio. The monthly long-minus-short value-weighted portfolio return is regressed on the four Fama-French-Carhart factors. The cumulative drift for a given drift horizon on the vertical axis is calculated as the regression intercept (monthly alpha) multiplied by the portfolio's horizon length in months. The graphs are shown for the full sample and two subsamples of firms in which the standard deviation of HQ analysts' forecasts (*SD HQ*) is either greater or smaller than that of all analysts' forecasts (*SD consensus*) for the firm-year. Analysts are ranked as HQ based on their previous year's absolute forecast errors, defined in Table 1.

Table 1: Individual analyst analysis: characteristics and persistence of forecast accuracy

Panel A conducts univariate analysis for high- and low-quality (HQ and LQ) analysts. HQ (LQ) quality analysts are those whose closest absolute forecast errors are below (at or above) the median closest absolute forecast error for the firm's earnings announcement. The closest *absolute forecast error* for an analyst is the absolute difference between the analyst's forecast closest to the earnings announcement prior to the announcement day and the actual annual earnings, divided by the share price at the beginning of the calendar year. *Overall tenure* is the number of years since the analyst first appeared in the I/B/E/S file. *Firm-specific tenure* is the number of years since the analyst began covering the specific firm in the I/B/E/S file. *Brokerage house size* is the number of analysts in the analyst's brokerage house. *Firms covered* is the number of firms covered by the analyst. Panel B shows probit models predicting the *HQ status* indicator, which equals one if the analyst is ranked as HQ in a given firm in the current year and zero otherwise (columns (1), (2), (5), and (6)) and the regressions for the analyst's closest absolute forecast error in a given firm in the current year in columns (3) and (4). *HQ status in other firms* equals one (zero) if the analyst is of HQ in the majority of the other firms that the analyst follows during the year; analysts who have an equal number of HQ and LQ rankings in other firms are excluded. *Firm size* is the log of the firm's market value of equity, equal to the stock price times the number of shares outstanding at the end of the month prior to the annual earnings announcement. *Annual return* is the annual return of the firm's equity over the 12 months prior to the earnings announcement month. *Leverage* is the book value of total liabilities divided by the book value of total assets, and *Book-to-market* is the book value of common equity divided by the market value of equity at the end of the fiscal year. *Number of analysts* is the number of analysts following the firm. All independent variables are measured prior to the announcement date. The probit coefficients are reported as marginal probability effects. All models include the intercept. Robust standard errors are clustered by firm. *z*-statistics for probit models and *t*-statistics elsewhere are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Panel A. HQ and LQ analyst characteristics

Analyst or announcement characteristic	HQ analysts	LQ analysts	Difference (t-statistic)
Overall tenure	7.07	7.00	0.07*** (4.73)
Firm-specific tenure	3.04	2.97	0.07*** (8.61)
Brokerage house size	65.76	63.04	2.72*** (19.14)
Firms covered	17.60	17.55	0.05* (1.79)
Absolute forecast error	0.0081	0.0089	-0.008*** (-12.83)
Forecast error	0.00185	0.00181	0.00004 (0.66)

Panel B. Persistence of analysts' forecasting performance

	HQ status		Absolute forecast error		HQ status	
	(1)	(2)	(3)	(4)	(5)	(6)
HQ status (year t-1)	0.0414***	0.0407***	-0.0007***	-0.0007***	0.0531***	0.0382***
	(25.54)	(25.13)	(-16.14)	(-15.91)	(32.95)	(22.84)
HQ status in other firms (year t-1)					0.0387***	0.0514***
					(23.14)	(31.83)
Firm size	0.0034***	0.0011***	-0.0061***	-0.0061***		
	(10.52)	(3.11)	(-25.12)	(-25.15)		
Annual return	-0.0003	0.0005	-0.0009***	-0.0009***		
	(-0.50)	(0.88)	(-5.68)	(-5.67)		
Leverage	0.0006	0.0003	0.00573***	0.0057***		
	(0.37)	(0.20)	(6.58)	(6.55)		
Book-to-market	0.0001	0.00003	0.00002	0.00002		
	(1.40)	(0.54)	(1.48)	(1.48)		
Number of analysts	0.0007***	0.0007***	0.00022***	0.0002***		
	(12.40)	(14.52)	(10.78)	(10.69)		
Overall tenure		0.0007***		-0.00004***		0.0002
		(4.33)		(-7.12)		(1.12)
Firm-specific tenure		0.0023***		0.00001		0.003***
		(9.66)		(1.45)		(12.90)
Brokerage house size		0.0001***		-0.000001		0.0001***
		(10.63)		(-0.26)		(7.23)
Firm coverage		-0.0007***		0.00003***		-0.0005***
		(-12.32)		(8.40)		(-8.63)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects			Yes	Yes		
Observations (analyst-years)	485,815	485,815	485,815	485,815	443,262	443,262
Adj. R-squared			0.344	0.344		

Table 2: The number of HQ analysts and the accuracy of their forecasts relative to the consensus

Panel A compares the accuracy of the HQ analysts' average forecast and the consensus forecast sorted by the number of HQ analysts following the firm in a given earnings announcement year. Analysts are ranked as HQ based on their previous year's absolute forecast errors, defined in Table 1. The SUE of consensus (SUE of HQ analysts) is standardized unexpected earnings, equal to the difference between the actual earnings and the average forecast provided by all analysts (HQ analysts) following the firm normalized by the stock price at the beginning of the year. Absolute SUE is the absolute value of SUE. *Accuracy improvement* is the percentage reduction from the absolute SUE of consensus to the absolute SUE of HQ analysts. *t*-statistics in the last column are for the difference in the means test between the absolute SUEs of the consensus and HQ analysts. Panel B reports sample characteristics, defined in Table 1, for two samples sorted based on whether the firm is followed by fewer than four HQ analysts or four or more HQ analysts in a given announcement year. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Panel A. The smallest number of HQ analysts needed to improve over the consensus forecast

Number of HQ analysts	Absolute SUE of Consensus	Absolute SUE of HQ analysts	Accuracy improvement	t-statistics Abs. SUE difference
1 or more	0.00656	0.00678	-3.31%	-8.63***
2 or more	0.00589	0.00595	-1.08%	-3.19***
3 or more	0.00514	0.00513	0.19%	0.54
4 or more	0.00461	0.00455	1.17%	2.99***
5 or more	0.00422	0.00415	1.52%	3.51***
6 or more	0.00404	0.00396	1.96%	3.95***
7 or more	0.00386	0.00377	2.35%	4.47***
8 or more	0.00377	0.00367	2.69%	4.60***
9 or more	0.00355	0.00346	2.61%	3.91***
10 or more	0.00346	0.00337	2.59%	3.44***

Panel B: Comparing samples in which HQ analysts' average forecast is either less or more accurate than the consensus

	Number of HQ analysts		Difference
	Fewer than 4	4 or more	(t-statistics)
Firm characteristics:			
Firm size ($M)	1,372	10,522	9,149*** (43.78)
Number of analysts	6.3	16.7	10.4*** (210.1)
Book-to-market	1.34	0.66	-0.67*** (-6.49)
Leverage	0.53	0.57	0.04** (14.93)
Annual return (%)	21.7	20.4	-1.3* (-1.74)
Avg. number of firms per year	1136	891	
Analyst characteristics:			
Overall tenure	6.91	7.42	0.51*** (8.66)
Firm-specific tenure	2.33	3.43	1.10*** (29.48)
Brokerage house size	57.1	69.4	12.3*** (18.73)
Firms covered	17.7	18.2	0.5*** (3.73)

Table 3: The immediate reaction to earnings announcements

This table reports the earnings response coefficients for measures of earnings surprise (SUE) based on the forecasts of all analysts following the firm (consensus) and on the forecasts of high- and low-quality (HQ and LQ) analysts. Analysts are ranked as HQ or LQ based on their previous year's absolute forecast errors, defined in Table 1. The dependent variable is the buy-and-hold abnormal return (based on the four-factor Fama-French and Carhart model) for the earnings announcement day and the following trading day. The SUE of Consensus and SUE of HQ and LQ analysts are defined in Table 2. All other variables are defined in Table 1. Columns (1)-(3) use the entire sample of earnings announcements, and columns (4)-(6) use the sample of earnings announcements by firms followed by at least four HQ analysts. All independent variables other than SUE are measured prior to the announcement date. The intercept and year fixed effects are included in all regressions. Robust standard errors are clustered by firm. *t*-statistics are provided in parentheses. The last two lines report the p-values for the chi-squared tests of the equality of the coefficients on SUE measures for the three analyst groups. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

	Full Sample			4 or more HQ analysts		
	(1)	(2)	(3)	(4)	(5)	(6)
SUE of Consensus	0.7245***			0.7526***		
	(13.62)			(5.20)		
SUE of HQ analysts		0.6211***			0.6927***	
		(12.93)			(5.12)	
SUE of LQ analysts			0.5691***			0.6044***
			(12.78)			(4.85)
Firm size	-0.0003	-0.0002	-0.0001	-0.0007	-0.0007	-0.0007
	(-0.71)	(-0.55)	(-0.37)	(-1.19)	(-1.15)	(-1.14)
Annual return	-0.0006	-0.0006	-0.0006	0.0007	0.0007	0.0007
	(-0.98)	(-0.89)	(-0.87)	(0.68)	(0.72)	(0.71)
Leverage	0.0059***	0.0057***	0.0057***	0.0030	0.0029	0.0028
	(3.73)	(3.55)	(3.56)	(1.24)	(1.19)	(1.12)
Book-to-market	0.00002	0.00001	0.00001	-0.0002***	-0.00018***	-0.00023***
	(0.38)	(0.29)	(0.28)	(-9.20)	(-7.67)	(-11.14)
Number of analysts	0.00002	0.00002	0.00001	-0.00004	-0.00003	-0.00004
	(0.25)	(0.28)	(0.08)	(-0.42)	(-0.40)	(-0.45)
Observations	44,709	44,709	44,709	20,221	20,221	20,221
Adj. R-squared	0.015	0.013	0.013	0.011	0.010	0.009
p-value (SUE of HQ analysts vs. SUE of consensus)		0.000			0.009	
p-value (SUE of HQ analysts vs. SUE of LQ analysts)			0.02			0.01

Table 4: Predicting market reaction on the earnings announcement day

The dependent variable is the buy-and-hold abnormal return (based on the four-factor Fama-French and Carhart model) for the earnings announcement day and the following trading day. Analysts are ranked as HQ or LQ based on their previous year's absolute forecast errors, defined in Table 1. *Predicted surprise* is equal to (HQ analysts' average forecast minus the consensus forecast) in columns (1) and (2) and (HQ analysts' average forecast minus LQ analysts' average forecast) in columns (3) and (4), normalized by the stock price at the beginning of the year. The *positive predicted surprise* indicator equals one if *predicted surprise* is positive and zero if it is negative. *Big positive predicted surprise* equals one if *predicted surprise* is greater than the median of positive values of *predicted surprise* and zero if *predicted surprise* is smaller than the median of the negative values of *predicted surprise* in year t-1. All independent variables are measured prior to the announcement date, and the regressions include the intercept and year fixed effects. Robust standard errors are clustered by firm. *t*-statistics are provided in parentheses. The last line of the table provides the two-day holding returns of a trading strategy that is long if the predicted surprise indicator variable in that column is equal to 1 and short if it is equal to 0. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

	Predicted surprise: HQ average – Consensus		Predicted surprise: HQ average – LQ average	
	(1)	(2)	(3)	(4)
Positive predicted surprise	0.0015*		0.0019**	
	(1.92)		(2.48)	
Big positive predicted surprise		0.0007*		0.0008**
		(1.71)		(1.98)
Firm size	0.0003	-0.00002	0.0003	-0.00002
	(0.76)	(-0.03)	(0.75)	(-0.04)
Annual return	-0.0002	-0.0006	-0.0002	-0.0010
	(-0.25)	(-0.67)	(-0.24)	(-1.15)
Leverage	0.0040**	0.0071***	0.0040**	0.0072***
	(2.45)	(2.89)	(2.45)	(2.90)
Book-to-market	0.00001	0.00001	0.00001	0.00001
	(0.03)	(0.15)	(0.03)	(0.11)
Number of analysts	-0.00001	-0.00006	-0.00001	-0.00002
	(-0.02)	(-0.51)	(-0.04)	(-0.15)
Observations	44,709	20,999	44,709	20,605
Adj. R-squared (%)	0.086	0.078	0.171	0.230
Two-day long-short strategy returns (%)	0.14*	0.20*	0.19**	0.24*
	(1.88)	(1.64)	(2.52)	(1.94)

Table 5: Returns following recommendation revisions

The dependent variable is the firm's stock return in the calendar month following the month with a recommendation revision. The sample consists of all recommendation revisions in the month when the annual earnings is announced by the firm. A recommendation is an integer from 1 to 5, where 1 is strong buy, 5 is strong sell, and 3 is hold. A recommendation revision is the negative of the difference between the current and the previous recommendations of an analyst; thus, a positive (negative) recommendation revision is an upgrade (downgrade). The *consensus, HQ, and LQ recommendation revision* variables are the averages of individual analysts' revisions of all, HQ, and LQ analysts following the firm, respectively, during the earnings announcement month. Analysts are ranked as HQ or LQ based on this year's absolute forecast errors, defined in Table 1. *HQ-consensus* is equal to HQ analysts' average recommendation revision minus the consensus recommendation revision if HQ analysts' average recommendation revision is either positive and greater or negative and smaller (more negative) than the consensus recommendation revision; otherwise, this variable is zero. The other independent variables are defined in Table 1 and measured prior to the earnings announcement. All regressions include the intercept, and robust standard errors are clustered by firm. The last two lines provide the event-time raw long-short strategy returns and calendar-time alphas of the value-weighted long-short portfolio based on the four-factor Fama-French and Carhart model in the month following the month with the revision. The long (short) position is in the firms for which the mean recommendation revision is positive (negative), according to the consensus, HQ, and LQ analysts' recommendation revision variables in columns (1)-(3), respectively. In column (4), the position in the stock is long (short) if *HQ-consensus* is positive (negative); otherwise, the stock does not enter either portfolio. *t*-statistics are provided in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
Consensus recommendation revision	0.0015*			
	(1.67)			
HQ analysts' recommendation revision		0.0025**		
		(2.09)		
LQ analysts' recommendation revision			0.0004	
			(0.30)	
HQ-consensus				0.0158***
				(2.62)
Lagged dependent variable	0.0141	0.0143	0.0158	0.0161
	(1.01)	(1.04)	(1.15)	(1.19)
Firm size	-0.0043***	-0.0043***	-0.0043***	-0.0043***
	(-4.26)	(-4.25)	(-4.20)	(-4.20)
Leverage	0.0008	0.0008	0.0007	0.0008
	(0.18)	(0.19)	(0.16)	(0.19)
Book-to-market	-0.00001	-0.00001	-0.00001	-0.00001
	(-0.34)	(-0.30)	(-0.32)	(-0.32)
Number of analysts	0.0003*	0.0003*	0.0003*	0.0003*
	(1.80)	(1.81)	(1.78)	(1.78)
Year fixed effects	Yes	Yes	Yes	Yes
Observations	21,381	21,381	21,381	21,381
Adj. R-squared	0.0430	0.0420	0.0419	0.0421
Long-short raw return next month (%)	0.36	0.85***	0.14	1.24***
	(1.61)	(3.11)	(0.49)	(3.34)
Long-short monthly calendar-time alpha (%)	-0.16	0.32	-0.31	1.06*
	(-0.45)	(0.79)	(-0.71)	(1.80)

Table 6: Forecast dispersion and future return volatility

The dependent variable is the standard deviation of the firm's daily returns in the month following the annual earnings announcement month. Consensus, HQ, and LQ analysts' forecast dispersion are the standard deviation of the consensus, HQ, and LQ analysts' forecasts, respectively, normalized by the stock price and using each analyst's closest forecast issued during the 60 days prior to the earnings announcement. Analysts are ranked as HQ or LQ based on their previous year's absolute forecast errors, defined in Table 1. The other independent variables are defined in Table 1 and measured prior to the announcement date. All regressions include the intercept. Robust standard errors are clustered by firm. *t*-statistics are provided in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)
Consensus dispersion	0.0730		
	(1.63)		
HQ analysts' dispersion		0.1315**	
		(2.55)	
LQ analysts' dispersion			0.0340
			(0.79)
Lagged dependent variable	0.5179***	0.5141***	0.5206***
	(12.69)	(12.82)	(12.63)
Firm size	-0.0007	-0.0006	-0.0008*
	(-1.59)	(-1.40)	(-1.73)
Annual return	0.0002	0.0002	0.0002
	(0.53)	(0.56)	(0.46)
Leverage	-0.0007	-0.0007	-0.0005
	(-0.29)	(-0.31)	(-0.23)
Book-to-market	0.0015	0.0016	0.0015
	(1.25)	(1.41)	(1.16)
Number of analysts	-0.00004	-0.00004	-0.00003
	(-1.22)	(-1.40)	(-1.13)
Year fixed effects	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Observations	4,812	4,812	4,812
Adj. R-squared	0.691	0.693	0.691

Table 7: Post-earnings announcement drift and analysts' uncertainty

The table reports the monthly abnormal returns for 1- to 11-month horizons following annual earnings announcements. Announcements are divided into two subsamples in which the standard deviation of HQ analysts' forecast errors is greater (the *high-uncertainty* sample) or smaller (the *low-uncertainty* sample) than the standard deviation of forecast errors for all analysts following the firm. Analysts are ranked as HQ based on their previous year's absolute forecast errors, defined in Table 1. A stock is assigned to the long or short portfolio, depending on whether its earnings surprise is positive or negative, respectively, where earnings surprise is defined based on the consensus estimate in Table 2. A stock enters a portfolio at the beginning of the month following the month of the announcement and is held for the length of the horizon. In columns (1)-(6), the monthly value-weighted portfolio returns are regressed on the four Fama-French-Carhart factors to obtain the reported drift, which is the intercept of the regression (monthly alpha). For the monthly alphas reported in column (7), the high-uncertainty long-short post-announcement raw return is regressed on the full-sample long-short post-announcement raw return and the four Fama-French-Carhart factors. *, **, and *** represent 10%, 5%, 1%, significance based on the regression *t*-statistics, respectively.

PEAD horizon (months)	High uncertainty			Low uncertainty			High uncertainty relative to full sample
	Long	Short	Long-Short	Long	Short	Long-Short	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1	1.09	-0.35	1.44*	-0.46	0.30	-0.76	1.18**
2	1.36***	-0.12	1.48**	0.13	0.33	-0.20	1.17*
3	1.87***	-0.58	2.46***	0.20	0.19	0.01	1.48**
4	1.06***	0.04	1.02	-0.27	0.35	-0.62	0.81
5	0.90***	0.18	0.72	-0.05	-0.35	0.30	0.43
6	1.19***	-0.40	1.58***	-0.03	-0.20	0.17	1.19***
7	0.97***	-0.35	1.42***	-0.16	0.07	-0.23	1.10***
8	0.97***	-0.48*	1.45***	-0.18	-0.22	0.04	1.18***
9	0.50***	-0.56**	1.06***	-0.14	-0.22	0.09	0.80**
10	0.42***	-0.44**	0.85***	-0.08	-0.33	0.25	0.59**
11	0.41***	-0.44**	0.85***	0.03	-0.4**	0.37*	0.32*

Table 8: Analysts' outputs at the aggregate level

The dependent variables are the monthly value-weighted market returns (Panel A) and the returns on the VIX (Panel B) for the calendar month following the month with the earnings announcement. The main explanatory variables in Panel A are based on the recommendation revisions during the earnings announcement month, defined in Table 5. The *consensus, HQ, and LQ mean recommendation revision* variables are the sums of all, HQ, and LQ analysts' revisions, respectively, divided by the number of revisions for all firms with an earnings announcement in that calendar month. The *HQ-consensus mean recommendation revision* is the *HQ mean recommendation revision* minus the *consensus mean recommendation revision* if the *HQ mean recommendation revision* is either positive and greater or negative and smaller than the consensus recommendation revision; otherwise, this variable is zero. In Panel B, *forecast dispersion* is a value-weighted average of firm-level analysts' forecast error dispersions prior to the earnings announcement, defined in Table 6, with the firms' market capitalizations as the weights. Analysts are ranked as HQ or LQ based on this year's (Panel A) and the previous year's (Panel B) absolute forecast errors, defined in Table 1. The *dispersion ratio of HQ (LQ) analysts* is the value-weighted average of firms' ratios of *HQ (LQ) analysts' dispersion* to *Consensus dispersion*. The control variables are the monthly *earnings-to-price ratio, dividend-to-price ratio, term spread, default spread, one-month T-bill rate, 30-year Treasury yield, rate of inflation*, and change in economic policy uncertainty (*EPU*), defined in Section 2. The models use Newey-West standard errors with three lags in Panel A and Huber-White robust standard errors in Panel B. The last row in Panel A reports the market model alphas obtained as follows. In Panel A, the alphas are from regressing a long-short market-holding strategy return on the market return, where the market-holding strategy is the following. In columns (5)-(7), long (short) the next month's market return if the consensus, HQ, or LQ analysts' mean recommendation revision of that column is greater (smaller) than the median of the respective mean recommendation revision over the previous 24 months; in column (8), long (short) the next month's market return if the *HQ-consensus mean recommendation revision* is positive (negative), and have a zero position in the market index if the *HQ-consensus* is zero. The last line in Panel B provides the alphas from regressing the VIX strategy return on the market return, where the VIX strategy return is a long or short VIX return, depending on whether the column's *dispersion* or *dispersion ratio* last month is greater or smaller than the median of the *dispersion* or *dispersion ratio* over the previous 24 months, respectively. *t*-statistics are provided in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Panel A: Predicting market returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Consensus mean	0.016*				0.015*			
recommendation revision	(1.85)				(1.87)			
HQ analysts' mean		0.017**				0.017**		
recommendation revision		(2.25)				(2.44)		
LQ analysts' mean			0.005				0.004	
recommendation revision			(0.91)				(0.70)	
HQ-consensus mean				0.027*				0.028*
recommendation revision				(1.89)				(1.89)
Lagged dependent	0.073	0.073	0.070	0.068	0.050	0.055	0.047	0.050
variable	(0.93)	(0.96)	(0.90)	(0.91)	(0.73)	(0.80)	(0.68)	(0.75)
Earnings-to-price ratio					-0.148	-0.120	-0.175	-0.140
					(-0.51)	(-0.41)	(-0.60)	(-0.47)
Dividend-to-price ratio					1.913**	1.869**	1.968**	1.936**
					(2.45)	(2.40)	(2.48)	(2.45)
Term spread					0.679**	0.665*	0.647*	0.614*
					(1.99)	(1.94)	(1.93)	(1.81)
Default spread					-1.733	-1.583	-1.857*	-1.678
					(-1.65)	(-1.51)	(-1.75)	(-1.58)
One-month T-bill yield					3.308	3.136	3.323	3.025
					(1.34)	(1.24)	(1.36)	(1.18)
Long-term T-bond yield					0.091	0.100	0.090	0.105
					(1.28)	(1.44)	(1.26)	(1.52)
Inflation					0.470	0.437	0.544	0.499
					(0.64)	(0.61)	(0.74)	(0.70)
Intercept	0.009***	0.010***	0.008***	0.008***	0.028	0.026	0.026	0.022
	(3.44)	(3.59)	(2.78)	(2.93)	(1.08)	(1.02)	(1.03)	(0.90)
Observations	265	265	265	265	265	265	265	265
Adj. R-squared	0.09	0.018	-0.001	0.011	0.028	0.037	0.018	0.031
Monthly alpha (%)					0.42	0.49*	0.31	0.57***
					(1.44)	(1.65)	(1.06)	(2.60)

Panel B: Predicting changes in market volatility

	(1)	(2)	(3)	(4)	(5)
Consensus dispersion	2.104				
	(0.35)				
HQ analysts' dispersion		5.345			
		(0.64)			
LQ analysts' dispersion			2.625		
			(0.39)		
Dispersion ratio of HQ analysts				0.107*	
				(1.91)	
Dispersion ratio of LQ analysts					-0.010
					(-0.28)
Lagged dependent variable	-0.046	-0.045	-0.047	-0.040	-0.047
	(-0.36)	(-0.36)	(-0.36)	(-0.30)	(-0.36)
EPU change	0.060	0.062	0.061	0.068	0.059
	(1.00)	(1.04)	(1.02)	(1.16)	(0.98)
Value-weighted market return	0.576	0.592	0.591	0.518	0.545
	(1.40)	(1.44)	(1.43)	(1.24)	(1.33)
Earnings-to-price ratio	-1.109	-1.075	-1.078	-1.560	-1.195
	(-0.80)	(-0.77)	(-0.77)	(-1.07)	(-0.86)
Dividend-to-price ratio	1.244	0.865	1.092	3.567	1.919
	(0.31)	(0.22)	(0.26)	(0.88)	(0.49)
Term spread	1.130	1.020	1.044	1.828	1.363
	(0.57)	(0.51)	(0.51)	(0.87)	(0.67)
Default spread	-4.540	-4.347	-4.523	-6.245	-5.034
	(-0.85)	(-0.82)	(-0.85)	(-1.14)	(-0.94)
One-month T-bill yield	7.197	6.844	6.563	8.362	7.834
	(0.52)	(0.49)	(0.47)	(0.60)	(0.55)
Long-term T-bond yield	-0.150	-0.176	-0.134	-0.246	-0.155
	(-0.37)	(-0.43)	(-0.33)	(-0.61)	(-0.38)
Inflation	-2.114	-2.338	-2.194	-1.974	-1.832
	(-0.53)	(-0.58)	(-0.54)	(-0.50)	(-0.46)
Intercept	0.127	0.120	0.123	0.075	0.150
	(0.77)	(0.73)	(0.74)	(0.47)	(0.87)
Observations	197	197	197	197	197
Adj. R-squared (%)	0.038	0.041	0.038	0.056	0.037
Obs. in Long portfolio	115	124	74	141	64
Obs. in Short portfolio	97	88	138	71	148
Monthly alpha (%)	1.85	2.04	-0.57	3.32***	-0.56
	(1.34)	(1.60)	(-0.40)	(2.68)	(-0.41)