Regulation Fair Disclosure and the Private Information of Analysts

Eric Zitzewitz[1]

Stanford GSB

April 2002

[1]Stanford Graduate School of Business, 518 Memorial Way, Stanford, CA 94305. Tel: (650) 724-1860. Email: ericz@stanford.edu. The author would like to thank Jay Coughenour and participants at the Stanford GSB Applied Micro brownbag for helpful suggestions and comments. The data for this project were provided by I/B/E/S as part of a program to encourage academic research.

Abstract

This paper reports evidence that Regulation Fair Disclosure has had its desired effect of reducing selective disclosure of information about future earnings to individual analysts without reducing the total amount of information disclosed. In particular, it finds that multi-forecast days, which typically follow public announcements or events, now account for over 70 percent of the new information about earnings, up from 35 percent before Reg FD. This result is obtained by applying a new methodology from Zitzewitz (2001a) for measuring the information content of individual forecasts. These results are strongest for the fourth quarter of 2000, when the SEC Chairman who introduced Reg FD was still in office; since the change in administration, some of the initial effects of Reg FD appear to have been reversed.

Regulation Fair Disclosure and the Private Information of Analysts

On August 10, 2000, the Securities and Exchange Commission (SEC) approved Regulation Fair Disclosure (Reg FD), which became effective on October 23rd of the same year. This regulation prohibits the selective disclosure of material information to financial professionals including analysts; companies with material information to disclose must now do so in a press release or conference call that is simultaneously open to all investors.

Reg FD has been controversial. Proponents have argued that the selective disclosure that Reg FD prohibits has at least three negative effects. First, selective disclosure creates incentives for analysts to optimistically bias their opinions in order to maintain access to information. Second, selective disclosure creates the opportunity for analysts' favored clients to earn trading profits at the expense of uninformed investors and thus might be considered "unfair." Third, by increasing the asymmetry of information, selective disclosure can reduce liquidity and increase firms' cost of capital.

In response, opponents of Reg FD have argued that informal conversations between analysts and management are essential to the transmission of information. Opponents argue that while Reg FD may make the dissemination of information more simultaneous and uniform, less information will be disseminated and/or information will be disseminated less frequently. This could potentially increase stock price volatility, particularly following earnings announcements.

As one might expect, proponents are disproportionately groups of small investors and capital-raising corporations. The former SEC chairman who introduced Reg FD, Arthur Levitt, noted that "two thirds of our letters [in favor of Reg FD] came from Fools," subscribers to the Motley Fool, a popular website for individual investors, and that "without them, Reg FD would not have happened."[1] A survey of CFOs in December 2001 supported Reg FD by a 64-21 margin,[2] while separate surveys in mid-to-late 2001 found that 80 and 90 percent of public companies claimed Reg FD has had a positive-to-neutral effect on their business.[3]

Opponents include analysts, their employers, and their institutional clients. A survey of sell-side analysts presented at a congressional hearing in May 2001 found that 69 percent believed that Reg FD had adversely affected the advice they provide to institutional clients.[4] Another survey of analysts and institutional investors found that more than 80 percent believed Reg FD gave companies an excuse to minimize communications and more than 50 percent claimed that conference calls and private discussions were less useful than before Reg FD.[5]

Although the decision to implement Reg FD has already been taken, this is still very much a live debate. In part, this is because opinions on Reg FD are correlated with partisan affiliation. Reg FD was issued in the last months of a Democratic administration; the chairman of the SEC has since been replaced by a Republican appointee, Harvey Pitt, who was reportedly critical of Reg FD when it was introduced.[6] Mr. Pitt has stated that while the principle behind Reg FD is "unassailable," the details are not. Among the details that could be adjusted is the definition of materiality and the fact that companies are currently allowed to make "inadvertent" selective disclosures so long as they make the information public within 24 hours, allowing the recipient analyst(s) and favored clients up to 24 hours in which to make what are essentially (but not legally considered) insider trades. There is also a question of whether the regulation will be enforced; there has been wide speculation in the financial press that Mr. Pitt will be less aggressive in enforcement in general than Mr. Levitt was.

This paper contributes to the debate by using a new methodology for measuring the information content of individual analyst's quarterly earnings forecasts [outlined in Zitzewitz (2001a)] to examine how the information environment has changed since Reg FD was implemented. Several recent studies have examined how Reg FD has affected the total amount of information that is disclosed to the market prior to an earnings release. This paper goes a step further and asks whether Reg FD has affected *how* the information is disclosed and, in particular, which analysts get early access to it.

Recent studies have examined how stock price volatility, liquidity, and earnings forecast accuracy have changed following Reg FD; these studies have mostly concluded that Reg FD has not had the negative effects that opponents predicted. Heflin, Subramanyam, and Zhang (2001) find that daily stock price volatility was higher in the fourth quarter of 2000 than it was in the fourth quarter of 1999, but that there was no increase in volatility around earnings announcements, as one might expect if Reg FD had constrained the pre-announcement flow of information. Eleswarapu, Thompson, and Venkataraman (2001) find that spreads and price impact were essentially unchanged following Reg FD.[7] Shane, Soderstrom, and Yoon (2001) find that earnings forecasts made within 20 days of an earnings announcement were as accurate following Reg FD as they were before, although the accuracy of longer-term forecasts declined.[8]

Like earlier studies, this paper finds that Reg FD has not reduced the total pre-announcement flow of information. But this does not mean that the regulation has had limited impact. Using the new methodology mentioned above to measure the information content of individual analyst forecasts reveals that the regulation had a dramatic impact in the way intended by regulators, especially in the fourth quarter of 2000, but that some of that impact has been reversed following the change in leadership at the SEC.

In order to determine whether Reg FD has reduced the private disclosure of information to analysts, this study examines the information content of forecasts and forecast revisions made on single and multi-forecast days. When many analysts issue or update their forecasts on the same day, it is usually not by coincidence, but rather to incorporate some new public information (e.g., earnings guidance, some other news event). Likewise, when a single analyst updates her forecast for a *well-covered* firm, it suggests that the analyst is doing so based on information that was not simultaneously available to other analysts. The information could have been distilled from publicly available information through the diligence of the analyst, or it could have been gained in conversations with either the company or its business partners. The correlation between single vs. multiple-forecast days and private vs. public information is not necessarily perfect, of course, but as Section 4.3 discusses, there is reason to believe that is it reasonably high both before and after Reg FD.

One might find this approach indirect, and wonder whether a better approach would be for the researcher read through news items, company press releases, and conference call transcripts to determine which ones contained new information about next quarter's earnings, and then study the frequency of these announcements and the information content of forecasts that followed them. A problem with this approach is, aside from its being extremely laborious, is that potentially relevant information is released very often, and so the researcher would have to make many ad hoc judgements about relevance. She could potentially use stock market returns to help, but even then, since many news items are informative only about future earnings, the researcher would have to make judgements about whether the news is informative about near-term earnings. Relative to this fairly problematic approach, we can view the approach in this paper as letting analysts' decisions about when to update their forecasts tell us when significant and relevant public information has been released.

Prior to Reg FD, about 70 percent of analyst forecasts occurred on days in which no other analyst issued a forecast, these solo forecasts contained about 65 percent of the new information about earnings. The remaining 30 percent of forecasts occurred on days in which many analysts updated their forecasts, and these forecasts contained about 35 percent of new information about earnings. Since Reg FD was implemented, solo forecasts have accounted for only 50 percent of forecasts and 27 percent of new information. Over 70 percent of new information is now being revealed on multi-forecast days. The results are even more extreme for the fourth quarter of 2000, when 16 percent of new information was revealed on single-forecast days. This evidence is consistent with firms complying with Reg FD, especially in the last quarter of 2000, when Arthur Levitt was still SEC Chairman.

An issue with this study, as well as with the three studies discussed above, is that the time period following Reg FD is short and hardly typical. The NASDAQ collapse started in earnest in October 2000, and the economy officially entered recession in March 2001. As one might expect, there was a lot of new information about earnings revealed in this period, much of it negative. I attempt to control for this issue by constructing matched samples based on the ex-post difference between actual earnings and expected earnings six months prior. This matching analysis suggests that the results are robust to this issue.

The importance of controlling for the total amount of new information about earnings is highlighted by contrasting our results with a study by Mohanram and Sunder (2002) that is similar in spirit to this paper, but uses a very different methodology. Mohanram and Sunder use a methodology by Barron, et. al. (1998) to infer analysts' information environment from their forecasts, in particular whether the amount of private information available to analysts has changed since Reg FD. The Barron, et. al. methodology assumes that analysts forecast simultaneously and that they report rationally constructed posterior expectations, and it infers high amounts of private information from high forecast dispersion. An issue with this approach is that high forecast dispersion could also be the result of either 1) exaggeration or anti-herding by analysts[9] or 2) the public consensus expectation changing rapidly over time, as it did as the economy entered recession in late 2000 and early 2001.

The methodology used in this paper explicitly controls for these two issues, and reaches a different conclusion. Mohanram and Sunder find that forecast accuracy has decreased and forecast dispersion has increased after Reg FD, and the Barron, et. al. methodology suggests that this is due to less public information and more private information after Reg FD. This study instead concludes that the share of new information that is private has fallen dramatically, suggesting that the decline in forecast accuracy and increase in forecast dispersion after Reg FD are due to the increased arrival rate of information about earnings as the economy entered recession.

The remainder of this paper is divided into six sections. The next section presents the methodology for measuring the information content of a forecast, which builds on the methodology in Zitzewitz (2001a). The second section examines the robustness of the measure to measurement error and to differences in analyst's forecasting strategies. The third section describes the data and some implementation issues, while the fourth and fifth sections present the results and some robustness checks. A conclusion follows.

1 Measuring forecast information content

This section describes the methodology for measuring the information content of forecasts that will then be used to measure the share of information that becomes available on single-forecast days. The first subsection specifies the information environment in which analysts operate. The next subsection describes how to construct a measure of forecast information content using a history of forecasts and actual earnings.

1.1 Information environment

Assume that analysts issue forecasts of a random variable A (for actual earnings). Discrete packets of information about A, which I will call signals, arrive sequentially. I assume that a signal is either observed by one analyst or by many: think of a signal observed by one analyst as a private conversation or piece of information uncovered through diligent analysis of publicly available data and of a signal observed by many analysts as resulting from a public announcement or event. After observing a signal, the analyst or analysts have the opportunity to issue or update their forecasts of A; for reasons I will discuss below, one can think of analysts as choosing to do so when their signal is especially informative about A. If a signal is not significant enough to be immediately incorporated into a forecast, it becomes public knowledge, and the next analyst(s) to forecast get(s) credit for it.

I will call each time at which a signal arrives a time period, indexed by t. Prior to the signal arrival, all analysts share a common prior that $A \sim N(C_t, \Sigma_t^{-1})$ where C_t is the mean of the common prior and Σ_t is the precision of the prior. I will refer to the mean of the common prior as the "consensus."[10]

After observing the common prior, one or more analysts receive new information about A, which I model as a noisy signal $S_t \sim N(A, p_t^{-1})$ with precision p_t, where I will assume without loss of generality that the noise in the signal is uncorrelated with the error in the consensus prior expectation, $E[(S_t - A)(A - C_t)] = 0$.[11] An analyst who observes this signal can perform Bayesian updating and construct a posterior about A (DeGroot, 1970):

$$
\begin{aligned}
A &\sim N[E_t, (\Sigma_t + p_t)^{-1}] \qquad (1) \\
E_t &= C_t + (S_t - C_t) \cdot \frac{p_t}{p_t + \Sigma_t}.
\end{aligned}
$$

It will be convenient to think about the signal as s_t, the difference between the posterior and prior expectations:

$$s_t = E_t - C_t = (S_t - C_t) \cdot \frac{p_t}{p_t + \Sigma_t} \sim N[0, a_t] \quad (2)$$

$$a_t = \frac{p_t}{\Sigma_t(p_t + \Sigma_t)} = \Sigma_t^{-1} - (\Sigma_t + p_t)^{-1}. \quad (3)$$

In two separate ways, a_t can be thought of as the value of the signal. First, as shown in (3), a_t is the difference between the mean-squared error of the prior and the posterior. It can thus be thought of as the reduction in an analyst's uncertainty about A resulting from observing the signal. Second, as discussed in Zitzewitz (2001a) and Zitzewitz (2001b), a_t is proportional to the value of early access to the signal to a mean-variance trader investing in a security whose returns are linear in A. Thus, under some assumptions, a_t can be thought of as proportional to the monetary value of the signal to an analyst who is selling early access to it.

1.2 Estimating forecast information content

This subsection describes a methodology for estimating a, the average information content of a set of forecasts. This methodology can be used to measure the share of information that reaches the market on single-forecast and multi-forecast days. The methodology is described in much more detail in Zitzewitz (2001a), this section only provides a summary.

Estimating a from a track record of forecasts and actual earnings is a two-step process. First, we need to understand how analysts convert posterior expectations into forecasts, then we can use this understanding to measure when and how information becomes available.

1.2.1 Estimating exaggeration/herding

An analyst who receives a signal can convert her new expectation E_t into a forecast F_t using a forecasting function:

$$F_t = g(E_t, C_t, \cdot). \tag{4}$$

The inverse of this function with respect to E_t can be written

$$E_t = g^{-1}(F_t, C_t, \cdot). \tag{5}$$

As discussed in Zitzewitz (2001a), one can estimate this inverse forecasting function non-parametrically across multiple forecasts using the fact that the error in the rationally constructed posterior expectation $e_t = A - E_t$ must be uncorrelated with everything known at time of forecasting, including F_t, C_t, and all other arguments of g^{-1}. In other words, one can estimate the regression equation:

$$\begin{aligned} A &= g^{-1}(F_t, C_t, \cdot) + e_t. \qquad (6) \\ e_t &= A - E_t \perp F_t, C_t, \cdot \end{aligned}$$

When I do this in Zitzewitz (2001a), I find that the average forecasting function is essentially linear and of the form

$$F_t - C_t = c + b \cdot (E_t - C_t) \tag{7}$$

with $b > 1$ and $c > 0$. The fact that $b > 1$ implies that analysts exaggerate the difference between their rationally constructed posterior expectations and the prior consensus (i.e., they anti-herd); the fact that c is slightly positive implies that in addition analysts are slightly optimistically biased.[12] The regression I actually estimate across multiple observations of A_i is:

$$\begin{aligned} A_i - C_{it} &= \alpha + \beta(F_{it} - C_{it}) + \varepsilon_{it}. \qquad (8) \\ \varepsilon_{it} &= A_i - E_{it} \\ \alpha &= -\frac{c}{b} \\ \beta &= b^{-1} \end{aligned}$$

The fact that the error term, the expectational error of the analyst at time of forecasting, must be uncorrelated with all information known at time of forecasting means that (8) can be estimated consistently using standard techniques.[13]

The right-hand side of (8) can potentially include other control variables known at time of forecasting such as multiple measures of C_t, and forecast, firm, and analyst characteristics. We can also allow $\beta = b^{-1}$ to vary with these variables. In Zitzewitz (2001a), we found that the inclusion of control variables did not affect estimates of β so long as C was constructed properly. We also found some evidence of heterogeneity in β: estimated past exaggeration by an analyst was the best predictor of future exaggeration, analysts exaggerated more on single-forecast than multi-forecast days, and that, controlling for these relationships, other forecast, firm, and analyst characteristics had only small correlations with exaggeration. Controlling for heterogeneity in β will be important to our estimates of information content, as discussed below.

1.2.2 Estimating information content

Once we estimate (8), we can estimate E and from that $a = Var(E - C)$. We can write this estimator of a as:

$$\widehat{a} = \widehat{\beta}^2 \cdot \widehat{Var}(F_{it} - C_{it}) = \frac{[\widehat{Cov}(A_i - C_{it}, F_{it} - C_{it})]^2}{\widehat{Var}(F_{it} - C_{it})} \rightarrow Var(E_{it} - C_{it}) = a. \qquad (9)$$

Intuitively, $\widehat{Var}(F_{it} - C_{it})$ captures the average deviation of forecasts from the prior consensus. Forecasts can be different from the prior consensus either because new information has been observed, or because forecasters are exaggerating (simply adding noise would be considered a form of exaggeration here). Adding the $\widehat{\beta}^2$ term corrects for this exaggeration, so $\widehat{\beta}^2 \cdot \widehat{Var}(F_{it} - C_{it})$ can be viewed as forecast deviation corrected for exaggeration.

1.2.3 Estimating the information content of simultaneous forecasts

When forecasts occur either simultaneously or close enough to each other that we cannot determine the order, estimating the information content of each individual forecast is complicated. In the data used in this paper, this occurs when forecasts are issued on the same day. In these cases, we can estimate the information content of each forecast as if it were the first forecast that day, but we cannot add these estimates together, since some of the new information may be duplicated.

Fortunately, for our purposes, it is sufficient to estimate the information content of the forecasts combined. If we assume that exaggeration factor b used by forecasters forecasting on the same day is the same, we can take the simple average of all forecasts made on the same day and then estimate $\widehat{a}$ for this new average forecast. In the data, forecasts made on multiple are usually very similar to each other; the standard deviation of these forecasts is 6 basis points of market cap, compared with a standard deviation of $\overline{F} - C$ of 22 basis points. This suggests that the estimated information content on multi-forecast days is not likely to be sensitive to a non-equal weighting of these forecasts. When I control for heterogeneity in b in section 5 below, I also form weighted average forecasts based on the b predicted for the individual analysts; doing so does not change the results.

2 Robustness of information measure

Before using the estimator $\widehat{a}$ to analyze how forecast information content has changed since the enactment of Reg FD, I consider the robustness of this estimator to various issues. To summarize, the measure of information content is robust to changes in optimism or exaggeration by a consistent factor. It will be downwardly biased when optimism or exaggeration vary in a way that is observed by forecast users but not by the econometrician, and it will be upwardly biased when the prior consensus measure used by the econometrician contains measurement error. Controlling for heterogeneity in exaggeration or optimism and for any potential consensus mismeasurement are therefore important issues that we will address when presenting the results.

In this section, I will abbreviate: $s = E - C$, $x = F - C$, $y = A - C$. From (6) and (7) above note that $x = b \cdot s + c$ and $y = s + e$ and that $E(s) = 0$, $E(e) = 0$, $E(se) = 0$.

2.1 Robustness to consistent exaggeration or optimism

Suppose analysts all increase their exaggeration factor to $b' = b \cdot \gamma$ with $\gamma > 1$. The variance of x rises by a factor of γ^2, while β falls by a factor of γ. The value of $\widehat{\beta}^2 \cdot \widehat{Var}(x)$ is thus unchanged.

Likewise, a constant increase in c affects the estimated α in (8), but not the estimated β and not the variance of x, so the estimated forecast information content is unaffected.

2.2 Robustness to variation in optimism

I will consider variation in optimism to be variation in c that is uncorrelated with s and e, as variation in c in a way that is correlated with s (i.e., being more optimistic when forecasting above the consensus) is what we mean by exaggeration, and variation in c in a way that is correlated with e is ruled out by the definition of e, since c is known at time of forecasting. Variation in c that is uncorrelated with s or e increases $Var(x)$ without affecting $Cov(x,y)$. Since $a = \beta^2 \cdot Var(x) = Cov(x,y)^2 \cdot Var(x)^{-1}$, variation in c reduces measured forecast information content.

How one interprets this depending on whether we think that users of forecasts have better information about variation in c than the econometrician. If forecast users do not anticipate the variation in c, then the variation adds noise to the forecasts; this additional noise should reduce their value to users for the same reason it reduces our measure of forecast information. If, however, the users anticipate and correct for the variation in c and we do not, then our estimate of c will be downwardly biased. In order to avoid this bias, I will attempt to control for forecast, firm, and analyst characteristics that are known to forecast users and may be correlated with forecast optimism.

2.3 Robustness to variation in exaggeration

Suppose that I forecasts are made by analysts who exaggerate using factor b and the other J forecasts are made by analysts who use factor $b' = \gamma \cdot b$. The first set are indexed by i, the second set by j. For simplicity, assume that c is constant and equal to zero. We also assume that for $k \neq i$, $E(s_i s_k)$:

$$\begin{aligned}
\widehat{a} &= \frac{[\widehat{Cov}(x,y)]^2}{\widehat{Var}(x)} = \frac{[\sum b \cdot s_i \cdot (s_i + e_i) + \sum b' \cdot s_j \cdot (s_j + e_j)]^2}{(I+J) \cdot [\sum b^2 \cdot s_i^2 + \sum b'^2 \cdot s_j^2]} \qquad (10) \\
&= \frac{1}{I+J} \cdot \frac{(\sum s_i^2 + \gamma \cdot \sum s_j^2)^2}{\sum s_i^2 + \gamma^2 \cdot \sum s_j^2} \\
&= \frac{\sum s_i^2 + \sum s_j^2}{I+J} - \frac{(\gamma-1)^2 \cdot \sum s_i^2 \cdot \sum s_j^2}{(I+J) \cdot [\sum s_i^2 + \gamma^2 \cdot \sum s_j^2]} \\
&\rightarrow a \cdot \left(1 - \frac{(\gamma-1)^2 \cdot I \cdot J}{(I+J) \cdot (I + \gamma^2 J)}\right) \leq a.
\end{aligned}$$

In other words, if exaggeration is inconsistent, then $\widehat{a}$ will be downwardly biased by a factor that is increasing as the ratio of exaggeration factors becomes more different from 1. The limit of this bias term as γ approaches infinity and zero is $\frac{J}{I+J}$ and $\frac{I}{I+J}$, respectively. Intuitively, the information from forecasts with lower exaggeration factors is "lost": if everyone is screaming, a whisper will not be heard, and information will be lost.

As with variation in optimism, if variation in exaggeration is not anticipated by forecast users, then the information loss captured by the bias term is real, and we would like to include it in our measure of information content. If variation in exaggeration is anticipated by users but not by the econometrician, then this term captures measurement error. We will attempt to examine the importance of this measurement error by predicting exaggeration based on observables (forecast, firm, and analysts characteristics) and then adjusting forecasts based on predicted exaggeration.

2.4 Robustness to consensus mismeasurement

If the prior consensus C is mismeasured, this can bias the estimate of β and $Var(x)$. The most likely way for the consensus to be mismeasured is for the consensus measure to exclude information that is known to the analysts, such as earlier signals that were too small to justify incorporation into a forecast. In this case, the difference between the true and the measured consensus, $C - C'$, would be known to the analysts and would therefore be uncorrelated with the expectational error in the prior or any posterior formed using that prior: $A - C \perp C' - C$ and $A - E \perp C' - C$. Instead of $y = A - C$ and $x = F - C$, I would construct our measure of information value using $y + n$ and $x + n$, where $n = C - C'$ and $E(n) = E(yn) = E(xn) = 0$. The estimator of a would be

$$\begin{aligned} \widehat{a} \; &\rightarrow \; \frac{[Cov(x+n,y+n)]^2}{Var(x+n)} = \frac{[Cov(x,y)+Var(n)]^2}{Var(x)+Var(n)} \\ &= \underbrace{\frac{Cov(x,y)^2}{Var(x)}}_{a} + Var(n)\cdot\underbrace{[1-\frac{Var(x)}{Var(x)+Var(n)}\cdot(1-\beta)^2]}_{>0,<1} \end{aligned} \tag{11}$$

This estimator of information content would be biased upward since we would be treating the information in n as new information.

As with variation in exaggeration, the interpretation of the bias due to consensus mismeasurement depends on whether forecast users observe the "true" or the "mismeasured" consensus. If clients observe the true consensus, then a reflects the true new information content of the forecasts to these users and $\widehat{a}$ would be upwardly biased. If, however, users observed only the "mismeasured" consensus, then the "bias" would reflect genuine additional information learned from the forecast. As before, the approach will be to attempt to control for consensus mismeasurement by conditioning on variables that are observable to forecast users at time of forecasting.

3 Data and implementation issues

In order to estimate forecast information content and the share of total information disseminated on single-forecast days, I need data on individual analyst forecasts, and I need to construct a measure of the consensus prior expectation before each analyst forecasts. This section discusses these two issues.

3.1 Data

The data for this study are from the I/B/E/S Detail History dataset of analysts' earnings forecasts.[14] The I/B/E/S data is free of survivorship bias and most analysts who make publicly available earnings forecasts provide their forecasts to I/B/E/S.[15]

Since past research has shown that the predictive power of long-term earnings forecasts is very low (e.g., Crichfield, et. al, 1978), we restrict our sample to quarterly earnings forecasts made up to 6 months prior to earnings release.

We want to measure the new information content in analyst's forecast relative to the current consensus; it is therefore important to know the dates and order in which forecasts were made public with some precision. This has only been possible with I/B/E/S data since the early 1990s. I/B/E/S dates forecasts using the date it was entered into the I/B/E/S system. It has been well documented (e.g., by O'Brien, 1988) that the lags between a forecast becoming public and its entry into the I/B/E/S system were substantial in the 1980s (i.e., up to a month). In the 1980s, analysts mailed their forecasts, often in monthly batches, to I/B/E/S where they were hand entered into the system. Since 1991-92, however, almost all analysts have entered their forecasts directly into the I/B/E/S system on the day they wish to make their forecast widely available (Kutsoati and Bernhardt, 1999). Current practice for analysts is now usually to publicly release forecasts within 24 hours of providing them to clients. I/B/E/S analysts have real-time access to each other's forecasts through this system, so an analyst entering a forecast into the system on Wednesday knows about forecasts entered on Tuesday and could potentially revise her forecast to incorporate their information. An additional advantage of the post-92 data is the shift from retrospective data entry by a specialist to real-time data entry by either the analyst or her employee should have considerably reduced data-entry-related measurement error.

In addition to limiting the sample to firm-quarters ending in 1/93-6/01, we also eliminate firm-quarter combinations with share prices of less than $5 or market capitalizations of less than $100 million (in 1999 CPI-deflated dollars) at the beginning of the 180-day forecasting window. This eliminates about 7 percent of the potential sample; I eliminate these equities primarily because extreme outliers were concentrated among these stocks and using these sample restrictions removes enough outliers that the results are no longer sensitive to the treatment of the remaining outliers.[16]

In addition, I exclude firm-quarter combinations with fewer than five analysts forecasting and exclude the first forecast(s) following a prior quarter's earnings announcement. This eliminates another 17 percent of the potential sample. Since I identify private and public information based on whether individual or multiple analysts updated their forecasts on a given day, excluding firms covered by very few analysis is important; further restrictions of the sample along this dimension do not affect the results.[17] Likewise, since I construct a measure of the prior consensus expectation from the time series of forecasts, I exclude the first forecast after an earnings announcement since for these forecasts there is likely to be substantial amount of public information about future earnings that is not incorporated in past forecasts. Table 1 reports summary statistics for the sample for the major variables used in the analysis.

3.2 Measuring the consensus

As discussed in section 2.4, measuring the consensus well is important, since measurement error can upwardly bias our estimates of forecast information content. Zitzewitz (2001a) experiments with three different methodologies for estimating the prior consensus: 1) taking the mean of all outstanding forecasts, 2) taking the mean of the three most recent forecasts, and 3) estimating the relationship between a series of forecasts and future earnings econometrically, and then applying the estimated model out of sample.

While measuring the consensus as the equal-weighted mean of all forecasts is a common approach, an issue with this measure is that if more information is available to later forecasters or if analysts incorporate prior forecasts into their estimates, then a properly constructed expectation should put much more weight on later forecasts. Measuring the consensus as the mean of the last three forecasts addresses this issue in a fairly *ad hoc* fashion; the econometric expectation described below address this issue in a less *ad hoc* manner.

The idea behind using the econometric expectation of earnings as a consensus measure is to let the data tell us how much extra weight to put on more recent forecasts. In the process, we also let the data tell us how to adjust for biases in the forecasts (e.g., optimism and exaggeration/herding). We divide the sample into subsamples for 1993-97 and 1998-01 and estimate the model:

$$A_i - M_{it} = a + b(\overline{F}^1_{i,t} - M_{it}) + c(\overline{F}^2_{i,t} - M_{it}) + dM_{it} + e \cdot days \tag{12}$$

where A_i is actual earnings, M_{it} is the mean of all outstanding forecasts prior to forecast t, $\overline{F}^1_{i,t}$ and $\overline{F}^2_{i,t}$ are the mean of all forecasts on the two most recent days before the date of forecast t on which forecasts were made, and *days* is the number of days between the forecast and earnings announcement. Including M_{it} adjusts for the fact that forecasted earnings-price ratios are negative correlated with earnings surprise, while including *days* controls for the fact that earnings forecasts get less optimistically bias as the earnings announcement approaches.

The coefficients estimated for one-half of the data are used to generate predicted values of A for the other half, and these predicted values are our econometric estimate of the expectation of A given all public forecasts made before forecast t. We split the sample in this way in order to avoid a data-snooping bias in our consensus measure. We allow the coefficients to vary depending on the number of forecasts made on each of the two most recent days – the idea is that an average of multiple forecasts should be given more weight, but how much more weight depends on the correlation of the private information embodied in these forecasts, so we want the data to tell us how much more weight (Table 2).[18]

We use the econometric expectation of earnings as our primary consensus measure since of the three measures, it should most closely approximate the optimal expectation of earnings given the history of public forecasts. In practice, however, the results are very similar if we use the mean of the last three forecasts.

4 Results

The section first presents evidence that the amount of information about earnings that gets out before the earnings announcement has not decreased since the implementation of Reg FD. It then presents evidence that the way in which the information gets out has changed dramatically: the share of information embodied in forecasts on single-forecast days has declined from 65 percent before Reg FD to 27 percent after.

4.1 Effect of Reg FD on total information released

As discussed in Section 1.1, our measure of forecast information content, a, can also be thought of as a measure of the reduction in the mean squared error of the consensus expectation (Equation 3). These MSE reductions are additive, so we can discuss the reduction in MSE from the beginning of the 180-day forecasting window to the end as the total amount of information embodied in analyst forecasts, and the MSE of the last consensus posterior expectation as the amount of information about earnings that was never embodied in an analyst forecast. We can also discuss the share of potential MSE reduction that occurs during the 180-day window as the share of information about earnings that gets released to analysts.

Figure 1 plots the MSE of the posterior expectation by the number days remaining before the earnings announcement. MSE reduction is roughly linear, suggesting that information is released more or less uniformly during the 180-day period. The MSE at the end of the 180-day window is about 35 percent of the MSE at the beginning, which implies that 65 percent of the information about earnings that could be potentially revealed to analysts gets revealed.

Table 3 presents estimates of the beginning and ending MSE and MSE reduction by quarter. Average mean squared error reduction remained roughly constant at about 60 percent both before and after the introduction of Reg FD. It did decline to 43 percent in the fourth quarter of 2000, but was above 80 percent the following two quarters. This suggests that if Reg FD had a "chilling" effect on the pre-announcement flow of information, it did so only in the first quarter in which it was introduced.

4.2 Effect of Reg FD on the way information is released

In addition to measuring the share of total potential information about earnings that is released before earnings announcements, we can also measure the share of the information that is released that becomes available on single and multiple-forecast days. A measure of the share of information on single-forecast days is given by

$$\lambda_S = \frac{n_S \cdot \overline{a}_S}{n_S \cdot \overline{a}_S + n_M \cdot \overline{a}_M}, \tag{13}$$

where n_S and n_M are the number of single and multi-forecast days, respectively, and $\overline{a}_S$ and $\overline{a}_M$ is the average information content of the forecasts made on these days.

Table 4 presents estimates of $\overline{a}_S$, $\overline{a}_M$, and λ_S in different time periods. Estimates are reported for five different time periods. Results from the three post-Reg FD quarters for which we have data (2000Q4, 2001Q1, and 2001Q2) are compared to the same three quarters from the year earlier and to the entire 1993-99 period. In addition, we compare results from the fourth quarter of 2000, the one quarter in which Reg FD was implemented and a Democratically-controlled SEC was encharged with enforcing it, with the first two quarters of 2001.

For each period and for single and multi-forecast days, we report estimates of the exaggeration coefficient β, the standard deviation of the difference between forecasts and the prior consensus $F - C$, and our measure of forecast information content, $\widehat{a} = \widehat{\beta}^2 \cdot \widehat{Var}(F - C)$. We then use these estimates to calculate λ_S. Figure 2 presents a graph of the time series of λ_S.[19]

The results suggest that λ_S dropped from approximately 65 percent before Reg FD to 27 percent afterwards. The drop in λ_S was most extreme in the fourth quarter of 2000, when it fell to 16 percent, but it recovered to 35 percent in the first two quarters of 2001.

Deriving standard errors for a and λ_S is not straightforward, so instead we bootstrap them. The bootstrap sampling is done on firm-quarter combinations in order to control for clustering of errors. We can also bootstrap p-values for differences in λ_S between comparisons and p-values for the significance of these differences. When we do this, we find that the post-Reg FD reduction in λ_S is clearly significant at conventional levels, but the increase in λ_S after the change in administration is only borderline significant. The bootstrapped standard errors also give an indication of the sensitivity of the results to outliers, since with a large number of firm quarters, since approximately $e^{-1} = 37$ percent of the bootstrap samples will exclude a particular anomalous firm-quarter, 37 percent will include it once, and the remaining 26 percent will include it twice or more.

4.3 Interpreting the results

The reason we are interested in how the share of information release on single-forecast days changed following Reg FD is that it might provide evidence about whether Reg FD is having its intended effect of encouraging firms to substitute public for selective disclosure. This would be the case if the private-information share of the information contained in single forecasts was significantly higher than for multiple forecasts, and if the private-information shares and single and multiple forecast days did not change dramatically after Reg FD. This subsection reviews characteristics of the data that suggest that single vs. multiple and private vs. public are highly correlated, and considers how the results would be affected if the correlation were not perfect.

Why might single vs. multiple and private vs. public not be perfectly correlated? An analyst forecasting based on public information might still be a single forecaster if: 1) there are very few analysts covering the stock, 2) the analyst updates her forecast more slowly than other analysts, or 3) the analyst updates her forecast much more quickly than other analysts. As mentioned above, to control for the first possibility, we limit the sample to firms covered by at least 5 analysts – further restricting the sample to firms covered by between 10 and 25 analysts does not qualitatively affect the results. If a single forecaster forecasts the day after many forecasters using the same information, the prior consensus used to estimate the information content of the single forecast will be updated to include the information from the multi-forecast day, and thus the estimated information content of the single forecast should be close to zero. The effect on the estimated λ_S would be essentially zero, since adding a solo-forecast day with no information content would not change $n_S \cdot \overline{a}_S$.

On the other hand, if one analyst beat all the other analysts covering the firm by a day in releasing her forecast, then a substantial part of the public information will be misclassified as private information. For a firm covered by many analysts, however, this should not be very likely. Only about 6 percent of solo forecasts were followed within one or two business days by a multi-forecast day, and this share was slightly lower before Reg FD than it was after (5.6 percent vs. 6.4 percent). This suggests that a decrease in public information appearing in solo forecasts was not likely to be responsible for the decline in λ_S from 65 to 27 percent after Reg FD.

An analogous problem would be analysts forecasting on the same day using private information. This could be the case either if: 1) analysts forecast on the same day coincidentally or because of overlapping calendars (e.g., all analysts issue forecasts on Fridays), or 2) analysts rush to copy the forecasts of other analysts on the same day that they are released. As mentioned above, forecasts on multiple days tend to be close to one another, with a standard deviation (SD) of 6 basis points of market cap, and far from the prior consensus (SD = 22 basis points). They are not exaggerated the way forecasts on single-forecast days are: the estimated β's for single forecasts on multiple-forecast days are 0.8 before Reg FD and 1.0 after, whereas they are about 0.5 for single forecast days.[20] The information content of forecasts on multi-forecast days is also much higher, both before and especially after Reg FD. This is all consistent with multi-forecast days being symptomatic of the public release of a significant amount of information, and not with analysts forecasting using private information coincidentally forecasting on the same day. In addition, both single and multi-forecast days appear evenly distributed through both the week and month.[21]

If analysts were able to copy the forecasts of other analysts on the same day that they were released, then this could cause private information to be released on a multi-forecast day. This might explain why λ_S declined gradually from 1993-99; as information technology allowed analysts more rapid access to each other's forecasts, copying a forecast on the same day may have become easier. It seems hard to imagine that this explains the dramatic drop in λ_S in 2000Q4, however. As Table 4 reports, the share of multi-forecast days did not change very much after Reg FD; what mostly changed was the amount of information embodied in forecasts made on those days. This is inconsistent with the change in λ_S being due to an increase in forecast copying.

5 Robustness checks

Although I performed many robustness checks that I mention where relevant, this section discusses two that are especially important: 1) controlling for other differences between pre and post-FD periods and 2) controlling for heterogeneity in optimism and exaggeration.

5.1 Controlling for other post-FD differences

As mentioned in the introduction, a potential criticism of other studies that examine the pre and post-Reg FD period is it is difficult to control for other differences in these periods. In particular, the economy was entering recession in late 2000 and early 2001 and thus there was more new information that usual about earnings, most of which was negative. If firms tended to announce negative news to all analysts at once but convey positive news privately, then this might help explain the decline in λ_S after October 2000.

In addition, since 1993 there has been a gradual increase in the number of analysts covering firms. While the number of analysts covering firms did not sharply increase with the implementation of Reg FD, it would be nice to control for the fact that any increase might mechanically increase the share of information that appears to be released on multi-forecast days. In order to control for these issues, I repeat the analysis in Section 4 comparing the Post-FD sample with a matched sample of firm-quarter combinations that are selected to have the same number of analysts covering and the same ex-post 180-day earnings surprise.

Table 5 repeats the analysis in Table 4 for the post-FD sample, for the unmatched pre-FD sample, and for three different matched pre-FD samples. Matched samples are constructed using both grid and propensity scoring methods.[22] Matching on analyst coverage and 180-day surprise affects the results only slightly, reducing the pre-FD estimate of λ_S from 64 to 62 percent, compared with 27 percent after the introduction of Reg FD. This suggests that post-FD reduction in λ_S is due to a reduction of λ_S for firm-quarters with a given analyst coverage and a given ex-post earnings surprise, not due to a shift the mix.

5.2 Controlling for heterogeneity

A second issue to control for is the potential that the results are due to a change in the bias due to heterogeneity in optimism (a heterogeneity) and exaggeration (b heterogeneity). Heterogeneity in either optimism or exaggeration that is observed by analysts' audiences but not by the econometrician can cause underestimates of information content. Heterogeneity in b is much more important than a heterogeneity as a potential source of bias, and heterogeneity is much more pronounced on single-forecast days than on multi-forecast days, when most analysts use $b \approx 1$ and $a \approx 0$. As a result, one's main concern should be that an increase in b heterogeneity after the implementation of Reg FD is causing an underestimate of $\overline{a}_S$ and thus λ_S in this period.

One approach for analyzing the potential effects of a and b heterogeneity is to predict a likely a and b for each forecast based on observable characteristics of the forecast, and then perform our analysis on forecasts adjusted for predictable optimism and exaggeration. Since we are most concerned about heterogeneity that is observed by an analyst's audience, this approach is attractive in that by removing heterogeneity that can be predicted by the econometrician, we are hopefully removing most of the heterogeneity that can be predicted by the audience.

In particular, I estimate the following interaction equation:

$$\begin{aligned} A_i - C_{it} &= (\gamma Z_{it}) + (\delta Z_{it}) \cdot (F_{it} - C_{it}) \qquad (14) \\ \alpha_{it} &= -\frac{a}{b} = \gamma Z_{it} \\ \beta_{it} &= b^{-1} = \delta Z_{it}, \end{aligned}$$

where Z_{it} is a vector of forecast, firm, and analyst characteristics. Results from estimating (14) on the entire sample are presented in Table 6. The best predictor of future exaggeration is past exaggeration by a given analyst, and the best predictor of an optimistic bias is past optimism. Analysts at larger brokerage (which are usually the more prestigious) exaggerate less and are less optimistic than analysts at smaller brokerages. Analysts are more optimistic when forecasts are disperse (Ackert and Athanannos, 1997), and they exaggerate more when forecasting immediately after another analyst. Analysts also exaggerate positive news more than negative (Zitzewitz, 2001a).

Exaggeration and optimism are then predicted for each forecast in a given year using (14), where γ and δ are estimated using a sample including all other years. Using these estimated $\widehat{\alpha}_{it}$ and $\widehat{\beta}_{it}$, I then construct an exaggeration and optimism-adjusted forecast:

$$\widetilde{F}_{it} = C_{it} + \widehat{\alpha}_{it} + \widehat{\beta}_{it} \cdot (F_{it} - C_{it}). \qquad (15)$$

Table 7 repeats the analysis in Table 4 using the adjusted forecasts $\widetilde{F}_{it}$. The main impact of controlling for heterogeneity in forecasting function is to increase the information content of solo forecasts after Reg FD. Analysts who had exaggerated less historically had relatively more informative forecasts after Reg FD than before. Controlling for heterogeneity in exaggergation results in these more informative forecasts receiving more weight, so the average solo forecast appears more informative, and the share of information released on single-forecast days increases slightly, from 27 to 33 percent, although the difference between the pre and post-FD periods remains significant.

6 Conclusion

This paper has presented evidence that while the total flow of earnings-related information to the clients of analysts has not been reduced since the implementation of Regulation Fair Disclosure, the way in which that information reaches them has. In particular, the share of new information embodied in forecasts made on single-forecast days has declined from 65 percent before Reg FD to 27 percent after Reg FD. Multi-forecast days typically indicate a public announcement or event about earnings, since most analysts tend to update their forecasts after such announcements. Accordingly, the dramatically increase in the share of information released on multi-forecast days suggests that much more information is being disclosed through public announcements, which was the intended effect of Reg FD.

It is perhaps not surprising that the effects of Reg FD appear to be the strongest during the fourth quarter of 2000, before Arthur Levitt, who championed Reg FD, was replaced as SEC Chairman by Laura Unger and Harvey Pitt, who were fairly vocal opponents. This could be evidence that firms began to take Reg FD less seriously once the agency response for enforcing it was run by opponents of the regulation. It could also be the case that firms initially overreacted to Reg FD, or that analysts needed a quarter to learn how to obtain private information in the Post-FD environment. Given the fact that we only have time-series identification, it is difficult to rule these alternative stories out.

The selective disclosure of private information to analysts, who in turn share the information preferentially with institutional clients, creates a source of asymmetric information in the market that should generate trading profits for informed investors at the expense of creating a wedge between the returns to uninformed investors and the cost of capital to firms. Since Reg FD appears to be having its intended effect of reducing selective disclosure, it should have benefitted individual investors and capital-raising firms at the expense of analysts and their employers and clients.

Whether or not Reg FD enhances long-run market efficiency depends in part on one's view of the analyst profession. One view of analysts is that they are toll takers who "blackmail" firms by demanding access to inside information in exchange for optimistically biased coverage. Firms would rather not provide this information, since for the reasons discussed above it increases their costs of raising capital, but given that most firms are cooperating and receiving optimistic coverage, they do not dare attempt to raise capital with merely unbiased or even negatively biased coverage. Rational investors pay attention to analysts' opinions, since they contain the information obtained through this blackmail. In this environment, information received through selective disclosure could very well be a substitute for information received through diligent analysis of public information, so eliminating selective disclosure could have two beneficial effects: 1) eliminating a source of asymmetric information and 2) creating more incentive for diligent research by the analysts.

In an alternative view, young analysts engage in diligent analysis and develop reputations.[23] Once analysts develop reputations for having informed opinions, they can leverage their reputation and force companies to disclosure inside information for the reasons described above. In this environment, Reg FD might have the short-term benefit of reducing asymmetric information, but at the longer-term cost of discouraging entry into the analyst profession and investments in reputation-building through diligent analysis. In this sense, Reg FD might be like a decision to void a class of patents, which would have the static benefits of reducing monopoly power but at the dynamic cost of discouraging innovation.

A question that the methods used in this paper could conceivably answer is which analysts (and which types of analysts) have been hurt most by Reg FD and thus, which were benefitting the most from selective disclosure before the regulation. Unfortunately, the three quarters of post-FD data available when this paper was written were not enough for meaningful results at anything below the aggregate levels. Assuming Reg FD remains in place, this will obviously change with time.

References

Abarbanell, J. and Lehavy, R.: 2000, Differences in commercial database reported earnings: Implications for inferences concerning analyst forecast rationality, the association between prices and earnings, and firm reporting discretion. University of North Carolina Mimeo.

Ackert, L. F. and Athanassakos, G.: 1997, Prior uncertainty, analyst bias, and subsequent abnormal returns, *Journal of Financial Research* **20**, 263–273.

Agrawal, A. and Chadha, S.: 2002, Who is afraid of reg fd? the behavior and performance of sell-side analysts following the sec's fair disclosure rules. University of Alabama Mimeo.

Barron, O. E., Kim, O., Lim, S. C. and Stevens, D. E.: 1998, Using analysts' forecasts to measure properties of analysts' information environment, *The Accounting Review* **73**, 421–433.

Crichfield, T., Dyckman, T. R. and Lakonishok, J.: 1978, An evaluation of security analysts' forecasts, *Accounting Review* **53**, 651–668.

Eleswarapu, V. R., Thompson, R. and Venkataraman, K.: 2001, The impact of regulation fair disclosure: Trading costs and information asymmetry. Southern Methodist University Mimeo.

Heckman, J. J., Ichimura, H. and Todd, P.: 1998, Matching as an econometric evaluation estimator, *Review of Economic Studies* **65**, 261–294.

Heflin, F., Subramanyam, K. and Zhang, Y.: 2001, Stock return volatility before and after regulation fd. Purdue and USC Mimeo.

Kutsoati, E. and Bernhardt, D.: 1999, Can relative performance evaluation explain analysts' forecasts of earnings. Tufts University Mimeo.

Mohanram, P. S. and Sunder, S. V.: 2002, Has regulation fair disclosure affected financial analysts' ability to forecast earnings? NYU Mimeo.

O'Brien, P.: 1988, Analysts' forecasts as earnings recommendations, *Journal of Accounting and Economics* **10**, 53–83.

Philbrick, D. R. and Ricks, W. E.: 1991, Using value line and ibes analyst forecasts in accounting research, *Journal of Accounting Research* **29**, 397–417.

Rosenbaum, P. and Rubin, D. B.: 1983, The central role of the propensity score in observational studies for causal effect, *Biometrika* **70**, 41–55.

Shane, P., Soderstrom, N. and Yoon, S. W.: 2001, Earnings and price discovery in the post-reg fd information environment: A preliminary analysis. University of Colorado Mimeo.

Zitzewitz, E.: 2001a, Measuring herding and exaggeration by equity analysts. MIT Mimeo.

Zitzewitz, E.: 2001b, Opinion-producing agents: career concerns and exaggeration. MIT Mimeo.

Notes

[1] *Wall Street Journal*, "'Fool' Followers Suffer, Not Gladly," 7/2/2001, C1.

[2]Survey of CFOs cited in *Institutional Investor*, "Making Peace with Reg FD," December 2001, p. 32.

[3]Two separate surveys by PriceWaterhouseCoopers, cited in *Financial Executive*, "Reg FD Beneficial, Say Tech Leaders," July/August 2001, p. 62 and *Wall Street Journal*, "Disclosure Rule is Receiving Mixed Reviews," 10/30/2001, p. B11, respectively.

[4]A Securities Industry Association survey cited in *Wall Street Journal*, "Deals & Deal Makers: Rule of Fair Disclosure Hurts Analysts, House Subcommittee is Told at Hearing," 5/18/2001, p. C15.

[5]AIMR survey cited in *Wall Street Journal*, "Disclosure Rule is Receiving Mixed Reviews," 10/30/2001, p. B11.

[6]*Strategic Finance*, "Fallout From Fair Disclosure," November 2001, p. 63. Commissioner and former acting Chairman Laura Unger, a Republican appointed by President Clinton, has been an even more outspoken opponent of Reg FD, arguing that "anecdotally, there seems to be a correlation between the quality of information under this new [regulation] and the drop in tech stocks" (Fortune, "Is Reg FD wrecking your portfolio," 4/16/2001, p. 392).

[7]Eleswarapu, et. al. (2001) actually find that spreads and price impact decreased after Reg FD, but only for stocks that were decimalized. For stocks that were not decimalized, there was no significant change.

[8]Agrawal and Chadha (2002) and also report that average overall forecast accuracy has declined; they do not distinguish between short and long-term forecasts.

[9]Zitzewitz (2001a) documents that analysts forecast more than twice as far away from the prior consensus on average as an rationally constructed posterior expectation.

[10]Although the financial press commonly calculates the consensus as the mean of all outstanding forecasts, the concept they are after is the best estimate given all public information. C_t refers to the latter; as we discuss below, one can usually construct a better estimate of C_t than the mean of all outstanding forecasts.

[11]This assumption is without loss of generality since for any signal S_t' we can always define a new signal $S_t = A - (C_t - A) \cdot Cov(S_t' - A, C_t - A) \div Var(C_t - A)$ that will have the desired orthogonality property.

[12]In a estimation of (8) for solo forecasts from 1993-99, we find an α of 2 basis points (i.e., 2 cents of EPS for a stock with a $100 share price) and a β of 0.5. In some sense, the exaggeration implied by the $\beta < 1$ is a more important phenomenon than the slight average optimistic bias. Considering that the 46 percent of forecasts that are more than 4 basis points below the prior consensus are pessimistically, rather than optimistically biased, exaggeration seems to be more important than the average optimism bias in determining the sign of the expected error in a particular forecast.

[13]Notice that this is still true even if analysts' decisions to issue forecasts depend on the strength of signal received or other factors, since all of these must be known at time of forecasting and thus be uncorrelated with the error term. Also notice that e_{it} will be correlated for multiple forecasts of the same A_i. As discussed in Zitzewitz (2001a), one needs to correct for this when calculating standard errors.

[14]In contrast to some prior studies but consistent with Zitzewitz (2001a) and Zitzewitz (2001b), this study uses I/B/E/S actual earnings rather than COMPUSTAT earnings. Although the basic results are not sensitive to this choice, we use I/B/E/S actuals because they are recorded on same basis that analysts make their forecasts for I/B/E/S. If potential clients use the I/B/E/S data to evaluate analysts, they are most likely to compare I/B/E/S forecasts with I/B/E/S actuals, and therefore the I/B/E/S actual is the number the analyst should focus on forecasting with her

I/B/E/S forecast. Past researchers, e.g. Philbrick and Ricks (1991), had noted problems with I/B/E/S actuals and recommended using Compustat actuals. Abarbanell and Lehavy (2000) find that the quality of I/B/E/S actuals has improved significantly since 1992 and that after 1992 earnings response coefficients are significantly higher when I/B/E/S forecasts are matched with I/B/E/S rather than Compustat actuals.

[15]For example, 96 percent of the analysts ranked by Institutional Investor from 1996-99 appear in the IBES dataset.

[16]The companies covered by I/B/E/S analysts with market capitalizations under $100 million tend to be formerly larger cap companies that have experienced stock price declines. These companies tend to have very large variances in earnings/price ratios, which increases the heteroskedasticity in sample and makes it difficult to distinguish data entry errors from true outliers due to a low market cap in the denominator. Eliminating penny stocks from the sample also has the advantage of reducing discreteness problems that result from analysts forecasting and companies reporting earnings as a whole number of pennies per share.

[17]Specifically, I replicated the results in this paper for a sample of firm-quarters with coverage by between 10 and 25 analysts. Results were qualitatively similar, although they were obviously noisier since I was using only half the data.

[18]In these regressions, as in all the regressions in the paper, actual and forecasted earnings-per-share are normalized by the share price and observations are then weighted by their market capitalization to reduce heteroskedasticity. Results from equal-weighted regressions are qualitatively similar but have larger heteroskedasticity-robust standard errors. An analysis of the residuals reveals that the variance of the residual increases roughly proportionally with market cap, so weighting observations by market-cap is very close to the FGLS estimator.

[19]I calculate λ_S by the quarter of earnings being forecast, rather by the time period

in which the forecast is made, for a number of reasons. It is important to control for the correlation of errors within each firm-quarter combination mentioned in footnote 13 and for other differences between the pre and post-FD sample (e.g., ex-post 180-day surprise). This is most easily done if all observations from the same firm-quarter are used to calculate a given value of λ_S, since it allows to select on firm-quarter when bootstrapping standard errors or performing a matching analysis. Even though the beginning of the 180-day forecasting window for the fourth quarter of 2000, mid-August 2000, was two months before Reg FD became effective, newspaper reports suggested that many firms began attempting to comply with Reg FD before it became effective. An attempt to calculate λ_S monthly suggests that it began to decline from pre-FD levels in September 2000 and began to increase again in February 2001, but these monthly estimates are too imprecise for these differences to be statistically significant.

[20]The betas reported in Table 4 for multiple forecast days are for the average forecast that day, i.e. from a regression of $A - C$ on $\overline{F} - C$, not $A - C$ on $F - C$, and thus they are higher than the single-forecast betas referred to above.

[21]About 21 percent of all forecasts are each made on Tuesdays, Wednesdays, and Thursdays, with slightly smaller shares on Mondays and Fridays. Likewise, forecasts are also evenly distributed throughout the month.

[22]In nearest-neighbor matching, each observation in the post-FD sample is matched with the observation in the pre-FD sample with the closest values of the matching variable(s). Propensity score matching matches on the probability that an observation is from the post-FD sample, conditional on its observable characteristics (Rosenbaum and Rubin, 1983). As Heckman, Ichimura, and Petra (1998) discuss, when the relationship between the observables and the expected post-FD single-forecast information share is adequately summarized in the propensity score, propensity score

matching will be more efficient.

[23]The development of this alternative view benefitted from discussions with Jeremy Bulow.

Table 1. Summary statistics

	Mean	SD	Percentiles 5	Median	95
Forecast data (quarterly EPS/share price; in basis points)					
Actual earnings (A)	111	215	-127	122	331
Forecast (F)	122	174	-74	129	320
Consensus before forecasting (C)	126	168	-61	131	320
A - F	-11	126	-116	0	64
F - C	-4	47	-50	0	32
A - C	-15	129	-130	1	64
Forecast characteristics					
Number of analysts covering firm	13.7	7.1	5	12	28
Number of forecasts made on day	1.9	2.1	1	1	6
Single-forecast day?	0.66	0.47			
Days before earnings announcement	97	58	14	91	170
Firm characteristics					
Market capitalization (in $billions)	9.8	26.6	0.2	2.3	42.8
SD(prior forecasts)/price (basis points)	88	1675	0.1	5.4	241.7
Analyst characteristics					
Number of analysts in brokerage firm	58	43	8	52	142
Observations	683,158				
Firms	4,141				
Firm*quarter combinations	49,091				
Analysts	6,083				

This table reports summary statistics for the sample of earnings forecasts used in the paper: the I/B/E/S sample of quarterly earnings forecasts for quarters ending between 1/93 and 6/01 made within 6 months of earnings reporting date. The sample excludes the first set of forecasts made after a prior quarter's earnings announcement and firms that are covered by fewer than 5 analysts, have a 6-month-prior market cap below $100 million (in 99$), or have a share price below $5. All earnings variables are per share, divided by the 6-month prior share price.

Table 2. Regressions used for constructing the econometric expectation of earnings
Dependent variable: ACT - MEAN

No. estimates on second most recent day		Number of estimates on most recent day One	Two	Three+
One	LAST - MEAN	0.46	0.77	0.90
		(0.05)	(0.12)	(0.16)
	LAST2 - MEAN	0.30	0.30	0.14
		(0.06)	(0.07)	(0.14)
	Obs.	438,758	86,728	70,505
Two	LAST - MEAN	0.41	0.73	0.61
		(0.06)	(0.13)	(0.13)
	LAST2 - MEAN	0.56	0.75	0.33
		(0.14)	(0.19)	(0.17)
	Obs.	48,281	16,042	17,808
Three+	LAST - MEAN	0.55	0.57	0.23
		(0.16)	(0.22)	(0.24)
	LAST2 - MEAN	0.75	0.50	0.86
		(0.18)	(0.27)	(0.22)
	Obs.	19,766	10,732	22,817

Variable definitions (all earnings variables are per share, divided by the share price)

ACT	Actual I/B/E/S earnings
MEAN	Mean of all outstanding forecasts made prior to current day
LAST	Mean of all forecasts on the most recent day on which estimates were made
LAST2	Mean of all forecasts on the second most recent day on which estimates were made

The expectation of actual earnings (ACT) given the prior forecasting history is estimated by regressing ACT - MEAN on the mean of all forecasts (MEAN), the differences between this mean and the mean of all forecasts made on the two most recent days on which forecasts were made (LAST - MEAN and LAST2 - MEAN, respectively), and the number of days before an earnings announcement (Equation 12, described in Section 4.2). The coefficients are allowed to vary with the number of estimates on each day (each cell is a separate regression); in most cases they increase with the number of estimates on that day as one might expect. Regressions are market-cap-weighted and include a constant term. The coefficient on MEAN (not shown) is roughly -0.01, while the coefficient on days is -0.06 basis points. The coefficients shown are for the entire 1993-01 period, however to avoid a data snooping bias, the consensus measure is constructing by splitting the sample into 1993-96 and 1997-01 and using the coefficients from one half of the sample to construct the consensus measure for the other half.

Table 3. Share of potential information about earnings released during 180-day forecasting window, by quarter

		Average forecast error (in basis points)						
		First prior consensus			Last posterior expectation			MSE reduction
Period	Firm quarters	Mean	SD	MSE	Mean	SD	MSE	
1993-99	39,839	-5.5	62	3,847	3.5	37	1,384	64%
1999Q4, 2000Q1-Q2 (pre-Reg FD)	5,421	4.1	48	2,294	6.2	31	1,020	56%
2000Q4, 2001Q1-Q2 (post-Reg FD)	3,920	-12.7	73	5,562	2.0	46	2,097	62%
1993Q1	1,470	-11.3	68	4,805	2.0	49	2,417	50%
1993Q2	1,308	-6.4	80	6,483	-1.8	65	4,286	34%
1993Q3	1,631	-9.3	71	5,092	2.3	41	1,685	67%
1993Q4	1,860	-8.8	82	6,770	0.7	62	3,811	44%
1994Q1	1,834	-8.5	78	6,129	2.9	46	2,132	65%
1994Q2	2,005	-2.9	60	3,576	1.3	36	1,324	63%
1994Q3	2,049	-2.5	78	6,069	3.5	63	4,035	34%
1994Q4	2,069	-15.0	139	19,459	-1.0	85	7,236	63%
1995Q1	1,992	-4.1	68	4,680	2.7	39	1,501	68%
1995Q2	2,082	-0.2	82	6,805	3.0	37	1,393	80%
1995Q3	2,179	-5.1	78	6,182	3.8	33	1,097	82%
1995Q4	2,275	-13.5	65	4,351	-0.1	36	1,329	69%
1996Q1	2,250	-11.5	67	4,645	4.5	29	887	81%
1996Q2	2,360	-2.4	60	3,561	4.6	25	671	81%
1996Q3	2,498	-4.4	62	3,870	4.3	44	1,998	48%
1996Q4	2,559	-3.5	60	3,592	4.1	40	1,578	56%
1997Q1	2,527	2.8	43	1,827	6.6	22	524	71%
1997Q2	2,660	-1.8	41	1,683	3.9	20	424	75%
1997Q3	2,628	-3.3	44	1,919	4.7	21	455	76%
1997Q4	2,788	-6.2	53	2,845	3.0	31	943	67%
1998Q1	2,798	-9.1	37	1,472	3.2	16	265	82%
1998Q2	2,894	-9.0	50	2,614	2.7	27	719	73%
1998Q3	2,948	-15.5	59	3,673	1.7	30	883	76%
1998Q4	2,891	-11.6	66	4,491	2.7	49	2,420	46%
1999Q1	2,488	-3.3	54	2,912	6.2	27	783	73%
1999Q2	2,692	-0.4	43	1,876	4.5	19	374	80%
1999Q3	2,630	-0.1	50	2,496	4.1	22	492	80%
1999Q4	2,714	-0.1	56	3,160	5.3	42	1,775	44%
2000Q1	2,575	6.7	47	2,287	7.6	32	1,063	54%
2000Q2	2,708	5.9	39	1,524	5.8	16	305	80%
2000Q3	2,647	1.6	59	3,485	5.1	40	1,608	54%
2000Q4	2,573	-8.8	88	7,865	0.6	67	4,492	43%
2001Q1	2,311	-12.4	59	3,588	3.3	24	566	84%
2001Q2	861	-20.8	65	4,717	2.5	16	246	95%

This table compares the mean-squared error (MSE) of the consensus expectation of earnings at the beginning of the 180-day forecasting window and with the MSE at the end (i.e., immediately before the earnings announcement). As discussed in Section 2.1 and illustrated in equation (3), the reduction in MSE during the forecasting window can be thought of as a measure of the information content of all forecasts made during this period and the MSE at the end of the window can be thought of as a measure of the information that could have been embodied in a forecast but was not. The percentage reduction in MSE can thus be thought of as the share of information that could potentially become public that does. This share does not appear to change following the introduction of Reg FD in the fourth quarter of 2000.

Table 4. Forecast information content for single and multi-forecast days by time period

				From a regression of A - C on F - C					Single-forecast day share of		
Time period	Single or multi-forecast day?	Obs. (days)	Forecasts/day	Constant (basis points)	Beta	SD(F - C) (basis points)	Beta*SD (basis points)	Beta^2*Var (basis points)	Days	Forecasts	Info
1993-99	Multi	64,262	2.67	-2.1	0.992	18.6	18.5	342	85%	69%	63%
				(0.6)	(0.039)	(1.0)	(1.0)	(36)			(3%)
	Single	376,373	1.00	-1.9	0.513	19.5	10.0	100			
				(0.5)	(0.028)	(0.4)	(0.6)	(12)			
1999Q4, 2000Q1-Q2	Multi	10,752	3.12	5.2	0.940	12.5	11.7	137	80%	56%	65%
(pre-Reg FD)				(1.1)	(0.082)	(0.5)	(1.0)	(23)			(4%)
	Single	42,853	1.00	5.4	0.542	14.7	8.0	64			
				(0.9)	(0.043)	(0.7)	(0.6)	(9)			
2000Q4, 2001Q1-Q2	Multi	9,068	3.55	-8.6	1.404	19.2	26.9	725	77%	49%	27%
(post-Reg FD)				(1.7)	(0.156)	(1.7)	(3.7)	(208)			(11%)
	Single	30,937	1.00	-8.0	0.476	18.5	8.8	77			
				(1.6)	(0.111)	(1.5)	(1.7)	(29)			
2000Q4	Multi	3,640	3.41	-6.2	1.624	19.4	31.5	995	79%	52%	16%
(Reg FD, Levitt SEC)				(3.4)	(0.305)	(1.7)	(7.9)	(559)			(12%)
	Single	13,596	1.00	-6.9	0.358	20.3	7.3	53			
				(2.9)	(0.184)	(2.3)	(3.3)	(45)			
2001Q1-Q2	Multi	5,428	3.65	-10.0	1.281	19.0	24.4	594	76%	47%	35%
(Reg FD, Unger-Pitt SEC)				(1.8)	(0.118)	(2.2)	(3.2)	(160)			(6%)
	Single	17,341	1.00	-8.7	0.577	17.3	10.0	99			
				(2.1)	(0.081)	(1.6)	(1.2)	(25)			
Tests of changes in shares			Share before	Share after	Difference	Std. Error	p-value				
Pre-Reg FD vs. Post-Reg FD			65%	27%	38%	10%	<.0001				
Reg FD, Levitt SEC vs. Reg FD, Unger-Pitt SEC			16%	35%	-18%	14%	0.096				

This table reports estimates of average forecast information content for single and multi-forecast days in different time periods. Forecast information content is estimated as Beta*Var(F - C), where Beta is the coefficient from a regression of actual earnings less the prior consensus (A - C) on the difference between the mean of all forecasts on that day less the prior consensus (F - C), and Var(F - C) is the variance of the right-hand-side variable. Bootstrapped standard errors adjusted for clustering within firm-quarter combinations are reported in parathenthis. P-values of tests are also bootstrapped.

Table 5. Using matching techniques to control for differences in Reg FD sample

Panel A. Summary statistics of post-FD and matched pre-FD samples

	Matching method/variable	Number of analysts covering	180-day earnings surprise (Basis points)
Post-FD		14.2	-40.0
Pre-FD	Unmatched	12.5	-25.7
Pre-FD	Nearest neighbor: no. of analysts	13.9	-24.8
Pre-FD	Nearest neighbor: no. of analysts & 180-day surprise	13.9	-38.3
Pre-FD	Propensity score	14.1	-41.3

Panel B. Information content results for post-FD sample and matched pre-FD samples

Matching method	Match variable	Pre/post Reg FD	Single or multi-forecast day?	Single-forecast info share	Beta	SD(F - C) (Basis points)	Info content Beta^2*Var (Basis points)
None	None	Post	Multi	26.7%	1.404	19.2	725
			Single		0.476	18.5	77
None	None	Pre	Multi	64.0%	0.979	17.9	307
			Single		0.516	19.2	98
Nearest neighbor	No. analysts	Pre	Multi	62.7%	0.962	16.7	258
			Single		0.494	17.9	78
Nearest neighbor	No. analysts, 180-day surprise	Pre	Multi	62.4%	0.992	19.6	376
			Single		0.520	20.4	112
Propensity score	No. analysts, 180-day surprise	Pre	Multi	62.0%	0.968	22.1	459
			Single		0.572	20.2	134

This tabel attempts to control for two differences between the pre and post-Reg FD samples: 1) the gradual increase in number of analysts covering firms, which resulted in more analysts covering the average firm after Reg FD, and 2) the fact that all three post-FD quarters have been quarters in which average earnings news has been bad. We attempt to control for these differences by using matching techniques to construct a pre-FD sample with the same characteristics as the post-FD sample. Results from both propensity-score and nearest-neighbor matching methods are presented. Number of analysts is the number of unique analysts who issue a forecast for a given firm-quarter combination; 180-day surprise is the difference between the inital consensus forecast at the beginning of the 180-day forecasting window and actual earnings, divided by the share price.

Table 6. Regressions predicting optimism and exaggeration

	Single-forecast days		Multi-forecast days			
Observations	402237		187872		Interaction variable	
Effect of variable on:	Beta	Alpha	Beta	Alpha	summary statistics	
Interaction with	(F - C)	1	(F - C)	1	Mean	SD
Rank(Analyst's Past Beta)	2.37	1.37	1.94	2.42	0.50	0.23
	(0.02)	(0.86)	(0.02)	(0.97)		
Rank(Analyst's Past Alpha)	-0.03	60.83	0.03	52.40	0.50	0.27
	(0.01)	(0.85)	(0.02)	(0.97)		
Sign(F - C)	-0.45	-0.37	-0.36	2.62	0.48	0.50
	(0.01)	(0.50)	(0.01)	(0.60)		
Ln(Number of Analysts at Brokerage)	-0.11	-0.12	-0.02	0.13	3.70	0.93
	(0.00)	(0.24)	(0.01)	(0.30)		
Forecast made 1-2 days after most recent	-0.10	-0.66	-0.08	2.09	0.30	0.46
	(0.01)	(0.54)	(0.01)	(0.53)		
Ln[SD(prior forecasts)/price]	-0.01	-2.71	0.01	-2.28	-10.4	2.8
	(0.00)	(0.08)	(0.00)	(0.11)		

This table reports the results of two regressions, for single and multi-forecast days respective, of A - C on F - C and a constant, interacted with the variables listed. One can interpret the coefficients as the impact of the interaction variable on the expected exaggeration coefficient (beta, lower beta implies more exggeration) and optimism of a forecast (alpha, lower alpha implies more optimism). Regressions also include F - C and the constant directly.

Table 7. Single-forecast share of information, with and without correction for heterogeneity in exaggeration

Methodology	Time period	Single or multi-forecast day?	From a regression of A - C on F - C: Beta	SD(F - C) (basis points)	Beta^2*Var (basis points)	Single-forecast day share of Info
Unadjusted for heterogeneity	1999Q4, 2000Q1-Q2 (pre-Reg FD)	Multi	0.940	12.5	137	65%
			(0.082)	(0.5)	(23)	(4%)
		Single	0.542	14.7	64	
			(0.043)	(0.7)	(9)	
	2000Q4, 2001Q1-Q2 (post-Reg FD)	Multi	1.404	19.2	725	27%
			(0.156)	(1.7)	(208)	(11%)
		Single	0.476	18.5	77	
			(0.111)	(1.5)	(29)	
Adjusted for heterogeneity	1999Q4, 2000Q1-Q2 (pre-Reg FD)	Multi	0.857	12.7	119	59%
			(0.078)	(0.6)	(21)	(4%)
		Single	0.898	7.3	43	
			(0.091)	(0.4)	(7)	
	2000Q4, 2001Q1-Q2 (post-Reg FD)	Multi	1.437	18.8	727	33%
			(0.069)	(1.8)	(74)	(5%)
		Single	1.215	8.5	106	
			(0.140)	(0.6)	(19)	

The analysis in Table 4 is repeated after adjusting forecasts for the analyst's predicted exaggeration and optimism from the model in Table 6. Adjusting for heterogeneity in exaggeration and optimism mainly affects the information content of forecasts on single-forecast days. In the post-Reg FD period, analysts who had historically not exaggerated made especially informative forecasts, so adjusting for predicted exaggeration puts more weight on these forecasts and thus increases the relative information content of forecasts on single-forecast days.

Figure 1. Arrival of information during 180-day forecasting window

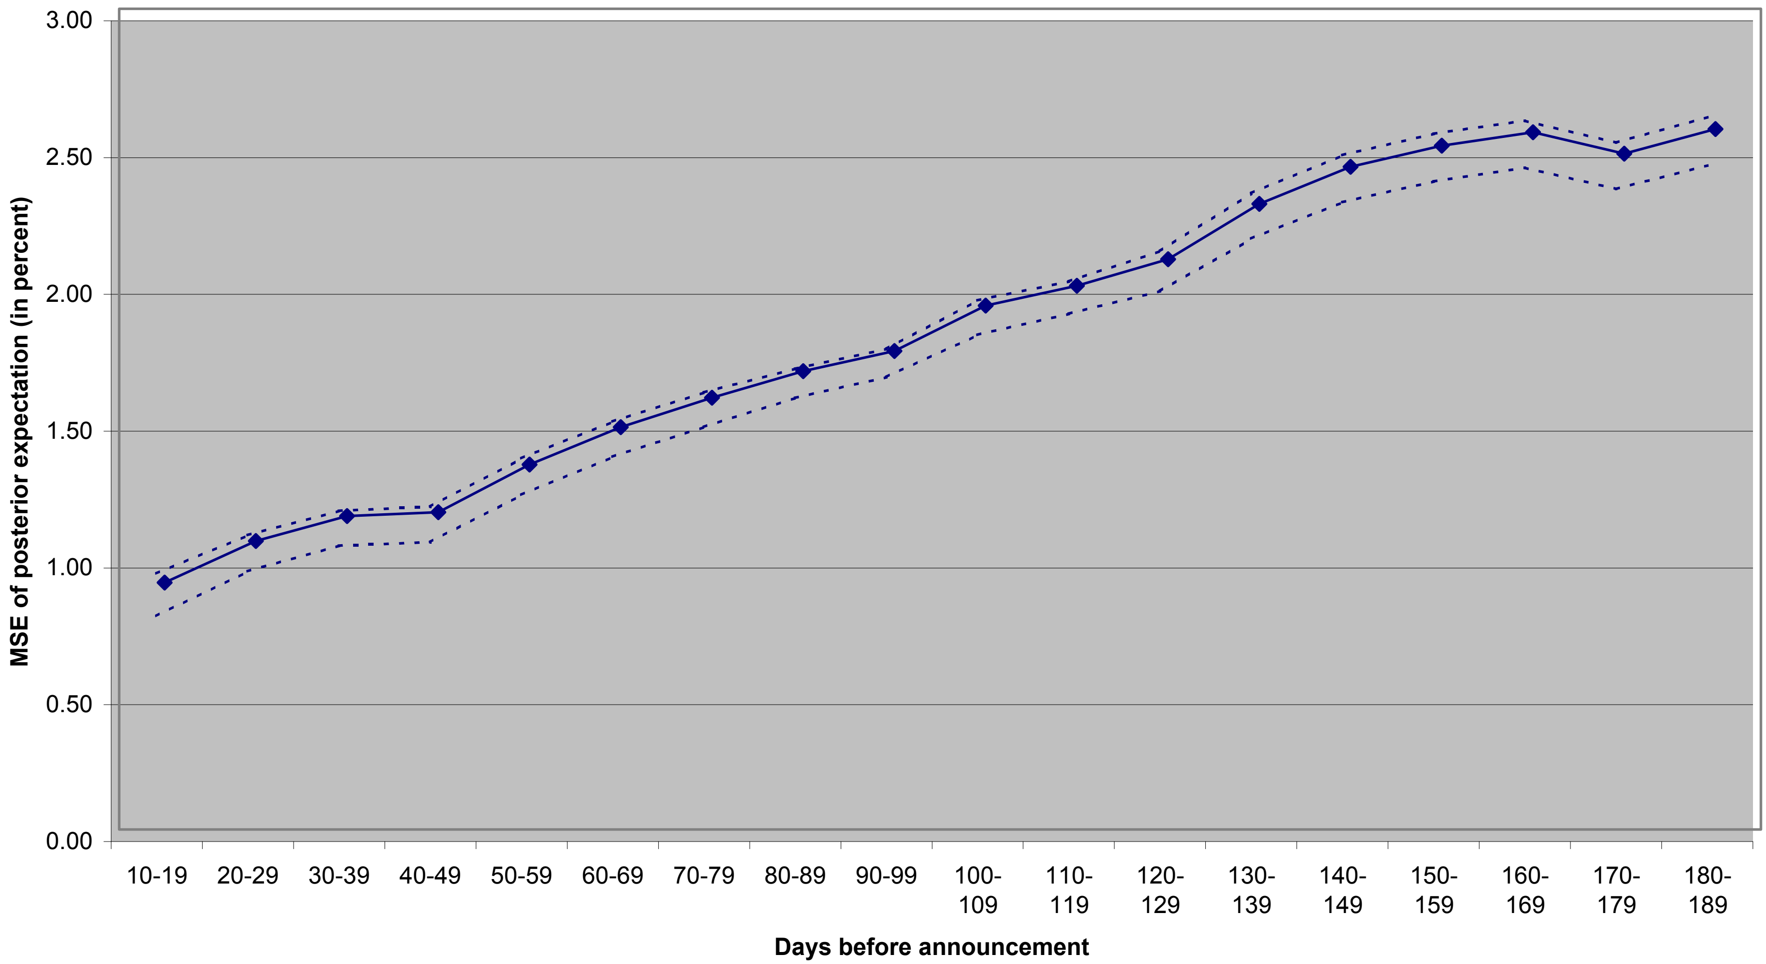

Data points are the coefficients on fixed effects from a regression of squared posterior expectation errors on days before announcement and firm*quarter fixed effects. Dashed lines indicate 95% confidence intervals.

Figure 2. Information arrival, pre and post-reg FD

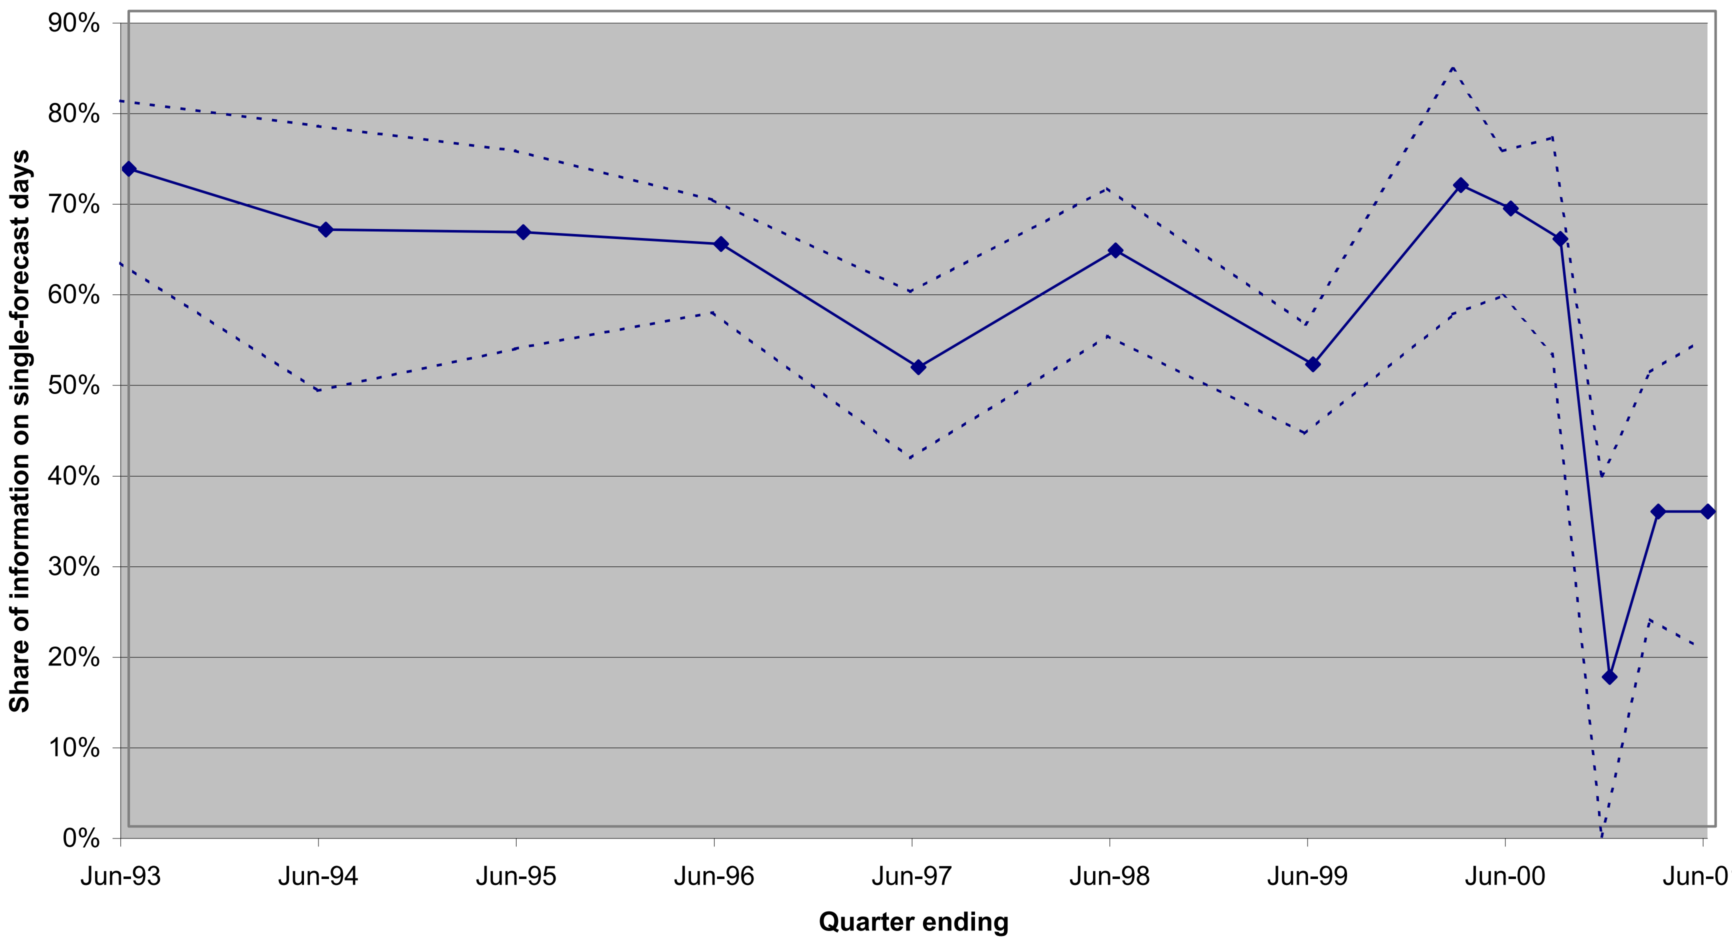

Share of forecast information arriving on single-forecast days is calculated as (# single-forecast days)*(avg. info content of single-forecast day)/(# days with forecast)*(avg. info content on day with forecast). Shares are calculately yearly from 1993-99, quarterly thereafter. Dashed lines indicate 95% confidence intervals, which are bootstrapped, adjusting for clustering on firm-quarter combinations.