

COMPETITION AND BIAS*

HARRISON HONG AND MARCIN KACPERCZYK

We attempt to measure the effect of competition on bias in the context of analyst earnings forecasts, which are known to be excessively optimistic because of conflicts of interest. Our natural experiment for competition is mergers of brokerage houses, which result in the firing of analysts because of redundancy (e.g., one of the two oil stock analysts is let go) and other reasons such as culture clash. We use this decrease in analyst coverage for stocks covered by both merging houses before the merger (the treatment sample) to measure the causal effect of competition on bias. We find that the treatment sample simultaneously experiences a decrease in analyst coverage and an increase in optimism bias the year after the merger relative to a control group of stocks, consistent with competition reducing bias. The implied economic effect from our natural experiment is significantly larger than estimates from OLS regressions that do not correct for the endogeneity of coverage. This effect is much more significant for stocks with little initial analyst coverage or competition.

I. INTRODUCTION

We study the effect of competition on reporting bias. Efficient outcomes in many markets depend on individuals having accurate information. Yet the suppliers who report this information often have other incentives in addition to accuracy. Two notable examples are (1) media outlets that trade off profits from providing informative news to consumers and voters versus printing information favorable to companies or political clients and (2) credit rating agencies such as Moody's and Standard & Poor's (S&P's) that get paid by the corporations they are supposed to evaluate. Will such conflicts of interest in these important economic and political markets lead to consumers, voters, and investors having poor information? Or can the market discipline these supply-side incentives and limit the distortions? The effect of competition from having more suppliers on bias is an important part of answering these questions.

The theoretical literature on the economics of reporting bias yields ambiguous answers when it comes to the potentially

*We thank three anonymous referees, Robert Barro (the editor), Edward Glaeser (the editor), Rick Green, Paul Healy, Jeff Kubik, Kai Li, Marco Ottaviani, Daniel Paravisini, Amit Seru, Kent Womack, Eric Zitzewitz, and seminar participants at Copenhagen Business School, Dartmouth, HBS, INSEAD, Michigan State, the Norwegian School of Management, Princeton, SMU, Texas, UBC, the UBC Summer Conference, the AFA 2009 Meetings, and the NBER Behavioral Finance Conference for a number of helpful suggestions.

disciplining role of competition. One strand of this literature argues that competition from suppliers (e.g., more newspapers or rating agencies) makes it more difficult for a firm (e.g., a political client or a bank) to suppress information (Besley and Prat 2006; Gentzkow and Shapiro 2006). Intuitively, the more suppliers of information are covering the firm, the more costly it will be for the firm to keep unfavorable news suppressed. Another perspective, the catering view, is that competition need not reduce and may increase bias if the end users (voters, consumers, or investors) want to hear reports that conform to their priors (e.g., Mullainathan and Shleifer [2005]).¹

In contrast to the theoretical side, empirical work on reporting bias is limited. There is some evidence that competitive pressures have helped discipline the media market. For example, a number of case studies of political scandals suggest that competition helps subvert attempted suppression of news (Gentzkow and Shapiro 2008). In addition, Gentzkow, Glaeser, and Goldin (2006), in their study of U.S. newspapers in the nineteenth century, point out that the papers of the time were biased tools funded by political clients. Importantly, they show that this situation changed in the late nineteenth century, when cheap newsprint engendered greater competition, which increased the rewards to objective reporting. These case studies notwithstanding, we have virtually no systematic evidence on the relationship between competition and bias.

In this paper, we attempt to measure the effect of competition on bias in the context of the market for security analyst earnings forecasts. This setting is an ideal one for studying this issue for a couple of reasons. First, the trade-offs faced by security analysts in issuing forecasts are similar to those of suppliers of reports in other markets such as media or credit ratings. Namely, analysts' earnings forecasts are optimistically biased because of conflicts of interest—a desire to be objective by producing the accurate forecasts desired by investors versus the prospect of payoffs by the companies that they cover, which want positive reports.² Hence,

1. There is a related literature on professional strategic forecasting in which forecasters, in a rank-order contest based on accuracy, differentiate themselves in their forecasts strategically à la Hotelling in equilibrium (Laster, Bennett, and Geoum 1999; Ottaviani and Sørensen 2005). These models are good at producing dispersion in forecasts due to convex payoffs (associated with publicity, etc.) and competition can lead to more dispersed forecasts. However, they do not necessarily lead to bias on average.

2. Companies naturally like analysts to be optimistic about their stocks, particularly when they are making initial or seasoned equity offerings. They would

the lessons regarding competition and bias obtained from our setting can be applied more broadly to other markets.

Second, at least two of the economic channels laid out in the theory literature predict that competition disciplines supply-side incentive distortions that are likely to apply to the market for analyst forecasts. Both of them point in the direction that an increase in competition among analysts should lead to less bias. The first channel is what Gentzkow and Shapiro (2008) term the independence rationale: Competition means a greater diversity of preferences among suppliers of information (analysts in our context) and hence a greater likelihood of drawing at least one independent supplier or analyst whose preference is such that he or she cannot be bought by the firm. This supplier's independence can have a disciplining effect on other suppliers. For instance, if an independent analyst makes a piece of bad news public, then the other analysts will be forced to do so as well.

A recent example that illustrates this independence-rationale mechanism at work in the market for analyst forecasts is the negative report produced by Meredith Whitney, a then unknown analyst at a lower-tier brokerage house named Oppenheimer, on Citibank on October 31, 2007. Citibank had a large market capitalization and was covered by in excess of twenty analysts. One can view Whitney as the draw of an independent analyst from among many. Whitney argued in the report that Citibank might go bankrupt as a result of their subprime mortgage holdings. Her report is now widely acknowledged as forcing the release of the pent-up bad news regarding financial firms, which had been unreported by other analysts. Her report had a disciplining effect on her peers at other brokerage houses, as they subsequently also reported on the same negative news about Citibank. A similar, though less famous example is Ivy Zelman, a housing analyst for Credit Suisse, who issued negative reports on the housing industry.³

not do business with an investment bank if the analyst were not positive about the stock (e.g., Brown, Foster, and Noreen [1985], Stickel [1990], Abarbanell [1991], Dreman and Berry [1995], and Chopra [1998]). A number of papers find that an analyst from a brokerage house that has an underwriting relationship with a stock tends to issue more positive predictions than analysts from nonaffiliated houses (e.g., Dugar and Nathan [1995], Lin and McNichols [1998], Dechow, Hutton, and Sloan [1999], and Michaely and Womack [1999]). Importantly, analysts' career outcomes depend both on relative accuracy and optimism bias (e.g., Hong and Kubik [2003] and Fang and Yasuda [2009]).

3. Anecdotes also suggest that independent analysts who "blow the whistle" tend to come from lower-tier brokerage houses that have lower investment-banking business revenues.

A second channel whereby competition limits bias that also holds in this market is that a firm's cost of influence increases with the number of suppliers or analysts. In the model of Besley and Prat (2006), if N analysts are all suppressing information in equilibrium and issuing optimistically biased reports, a single deviator who releases a bad forecast gets the same payoff as a monopolist. So the bribe that must be paid to each analyst to suppress information is thus independent of N , and so the total bribe is increasing in N . Increasing the number of analysts makes it more difficult to suppress information for the same reason that it makes it more likely that tacit collusion will break down. An implication of this mechanism is that the rewards are disproportionately high for the deviator if $N - 1$ other analysts are suppressing information.

The Whitney and Zelman examples also seem to bear out this implication. The deviators Whitney and Zelman became famous because few other negative reports were issued by their peers. Their rewards were by all accounts disproportionate. Beyond being offered jobs at better and higher-paying brokerage houses, they ended up being famous enough to start their own advisory businesses with special clients and revenue streams.

And a third reason that the market for analyst forecasts is ideal to study the relationship between competition and bias is that there are plentiful micro data concerning analysts and their performance that are not easily accessible in other markets, as well as opportunities to exploit natural experiments for identification. In particular, we identify the causal effect of competition or coverage on bias by using mergers of brokerage houses as a natural experiment. When brokerage houses merge, they typically fire analysts because of redundancy and potentially lose additional analysts for other reasons, including culture clash and merger turmoil (e.g., Wu and Zang [2009]). For example, if the merging houses each had one analyst covering an oil stock, they would only keep one of the two oil stock analysts after the merger. We use this decrease in analyst coverage for stocks covered by both merging houses before the merger (the treatment sample) to measure the causal effect of competition on bias.

During the period from 1980 to 2005, we identify fifteen mergers of brokerage houses (which took place throughout this twenty-five-year period) that affected 948 stocks (stocks covered by both merging houses) or 1,656 stock observations. We measure the change in analyst coverage and mean bias for the stocks in the treatment sample from one year before the merger to one year

after relative to a control group of stocks. The control group is stocks with the same market capitalization, market-to-book ratio, past return, and analyst coverage features before the merger as the treatment sample.

The exclusion restriction is that the change in the mean bias of the treatment sample across the merger date is not due to any factor other than the merger leading to a decrease in analyst coverage of those stocks. We think this is a good experiment because the merger-related departures of analysts due to redundancy or culture clash ought not *a priori* to be related to anything having to do with the biases of the forecasts of the other analysts, particularly those working for other houses.

As a benchmark, we begin with simple OLS regressions of average bias of earnings forecasts on analyst coverage.⁴ Henceforth, we will refer to the average or median bias of a stock simply as the bias of that stock. We restrict ourselves to stocks in the top 25% of the market capitalization distribution to facilitate comparison with the results from our natural experiment. The mean analyst coverage of these stocks is about 21 analysts and the standard deviation across stocks is about 10 analysts. Depending on the specifications we use, the economic effects are small to none. The largest effect we find is that a decrease in one analyst leads to an increase in bias by 0.0002 (2 basis points). The bias for a typical stock is about 0.03 (3%) with a standard deviation across stocks of about 0.03 (3%). Hence, these estimates obtained from cross-sectional regressions suggest a modest increase in bias by about 60 basis points to none as a fraction of the cross-sectional standard deviation of bias as we decrease coverage by one analyst.

Of course, these regressions are difficult to interpret, because of the endogeneity of analyst coverage. Existing studies suggest a selection bias in coverage in that analysts tend not to cover stocks that they do not issue positive forecasts about (e.g., McNichols and O'Brien [1997]). In this instance, we would then expect to find a larger causal effect from competition if we could randomly allocate analysts to different stocks.

Thus, we evaluate our natural experiment by first verifying the premise of our experiment regarding the change in analyst

4. Lim (2001) also tries to explain analyst bias by arguing that bias helps analysts to get access to a firm and hence to provide more accurate forecasts, and shows a negative correlation between bias and coverage in the cross section. His main variable of interest, however, is stock price volatility, whereas coverage is just another proxy for firm size. As we show, OLS estimates of bias on coverage is specification dependent.

coverage for the treatment sample from the year before the merger to the year after.⁵ We find, as expected, that the average drop in coverage for the treatment sample (using the most conservative control group) is around one analyst, with a *t*-statistic of around 5.7. The effect is economically and statistically significant in the direction predicted.

We then find that the treatment sample simultaneously experiences an increase in optimism bias the year after the merger relative to a control group of stocks. A conservative estimate is that the mean optimism bias increases by fifteen basis points (as a result of reducing coverage by one analyst). As we mentioned earlier, the sample for the natural experiment is similar to that of the OLS by construction. This is a sizable difference and suggests that the OLS estimates are biased downward.

Importantly, we find the same results when we look at the change in bias for analysts covering the same stocks but not employed by the merging firms, so our effect is not due to the selection of an optimistic analyst by the merging firms. We also find that this competition effect is significantly more pronounced for stocks with smaller analyst coverage (less than or equal to five). As we discuss below, these key additional results of our paper are consistent with the competition mechanisms articulated above.

We then conduct a number of analyses to verify the validity of our natural experiment, including showing that mergers are changing bias and not actual earnings or other firm characteristics, that mergers do not predict pretrends, and using nonparametric analysis to verify a shift in the distribution of forecasts. We also conduct a number of robustness exercises, including using an alternative regression framework that controls for brokerage-house fixed effects as well as firm fixed effects. Our results remain after all these additional analyses. Finally, we examine some auxiliary implications of the competitive pressure view, including looking at how implicit incentives for bias vary with analyst coverage.

The rest of the paper proceeds as follows. We describe the data in Section II and estimate the OLS regressions of bias on analyst

5. We expect these stocks to experience a decrease in coverage because one of the redundant analysts is typically let go. The exact number depends on a couple of factors. On one hand, the fired analyst might get a job with another firm and cover the same stock, which means the decrease in coverage might be less than one. On the other hand, a firm might lose or fire both analysts for reasons of culture clash or merger turmoil. In this case, if neither analyst is rehired by another firm, we would see a decrease in coverage of two analysts. What the magnitude turns out to be is an empirical question.

coverage in Section III. In Section IV, we provide background and statistics on the mergers. We discuss the methodology we use to measure the effect of the mergers on analyst coverage and bias in Section V and describe the results in Section VI. We conclude in Section VII.

II. DATA

Our data on security analysts come from the Institutional Brokers Estimates System (IBES) database. Our full sample covers the period 1980–2005. In our study, we focus on annual earnings forecasts because forecasts of these types are most commonly issued. For each year, we take the most recent forecast of the annual earnings. As a result, we have for each year one forecast issued by each analyst covering a stock.

Our data on U.S. firms come from the Center for Research in Security Prices (CRSP) and COMPUSTAT. From the CRSP, we obtain monthly closing stock prices, monthly shares outstanding, and daily and monthly stock returns for NYSE, AMEX, and NASDAQ stocks over the period 1980–2005. From COMPUSTAT, we obtain annual information on corporate earnings, book value of equity, and book value of assets during the same period.⁶ To be included in our sample, a firm must have the requisite financial data from both CRSP and COMPUSTAT. We follow other studies in focusing on companies' ordinary shares, that is, companies with CRSP share codes of 10 or 11.

We use the following variables. Analyst forecast bias is the difference between an analyst's forecast and the actual earnings per share (EPS) divided by the previous year's stock price. Given that the values of EPS reported by IBES tend to suffer from data errors, we follow the literature and use EPS from COMPUSTAT. Because our analysis is conducted at the stock level, we further aggregate forecast biases and consider the consensus bias expressed as a mean or median bias among all analysts covering a particular stock, which is denoted by $BIAS_{it}$. This is our main dependent variable of interest.

We also utilize a number of other independent variables. $COVERAGE_{it}$ is the number of analysts covering stock i in year t . $LNSIZE_{it}$ is the natural logarithm of firm i 's market capitalization

6. Our results are similar if we use IBES earnings numbers as opposed to those from COMPUSTAT.

(price times shares outstanding) at the end of year t . $SIGMA_{it}$ is the variance of daily (simple, raw) returns of stock i during year t . $RETANN_{it}$ is the average monthly return on stock i in year t . $LNBM_{it}$ is the natural logarithm of firm i 's book value divided by its market cap at the end of year t . ROE_{it} is firm i 's return on equity in year t . ROE is calculated as the ratio of earnings during year t to the book value of equity. Earnings are calculated as income before extraordinary items available to common stockholders (Item 237), plus deferred taxes from the income statement (Item 50), plus investment tax credit (Item 51). To measure the volatility of ROE ($VOLROE_{it}$), we estimate an AR(1) model for each stock's ROE using the past ten-year series of the company's valid annual ROEs. We calculate $VOLROE_{it}$ as the variance of the residuals from this regression. $PROFIT_{it}$ is firm profitability, defined as operating income over book value of assets. SP_{it} is an indicator variable equal to one if the stock is included in the S&P 500 index and zero otherwise. As in earlier studies, stocks that do not appear in IBES are assumed to have no analyst estimates.

Following earlier work, we exclude observations (stock-year) in which the stock price is less than five dollars or whose mean bias is at the outer tails—the 2.5% left and right tails. We also exclude analyst forecasts whose absolute difference exceeds ten dollars on the basis that this is likely a coding error.

III. OLS RESULTS

We begin by estimating a pooled OLS regression of the mean and median BIAS on lagged values of COVERAGE and a set of standard control variables, which include LNSIZE, SIGMA, RETANN, LNBM, VOLROE, and PROFIT. We additionally include an S&P 500 index indicator variable (SP500) as well as time and three-digit SIC industry fixed effects and potentially firm and brokerage-house fixed effects. Standard errors are clustered at the industry groupings.

These regressions are based on a sample of stocks in the top 25% of the market capitalization distribution. We restrict ourselves to this sample to facilitate a comparison with the results from our natural experiment.⁷ The summary statistics for these regressions (time-series averages of cross-sectional means,

7. Qualitatively, the same results hold even using the entire universe. We have replicated these results, which are consistent with those in Lim (2001).

TABLE I
SUMMARY STATISTICS ON THE IBES SAMPLE

Variable	Cross-sectional	Cross-sectional	Cross-sectional
	mean (1)	median (2)	st. dev. (3)
COVERAGE _{<i>i,t</i>}	21.45	21	9.57
Mean BIAS _{<i>i,t</i>} (%)	2.70	2.10	3.10
Median BIAS _{<i>i,t</i>} (%)	2.64	2.01	3.17
Mean FERROR _{<i>i,t</i>} (%)	3.31	2.39	2.93
Median FERROR _{<i>i,t</i>} (%)	3.24	2.26	3.00
FDISP _{<i>i,t</i>} (%)	0.75	0.41	1.02
LNSIZE _{<i>i,t</i>}	8.38	8.38	1.62
SIGMA _{<i>i,t</i>} (%)	40.72	35.04	21.03
RETANN _{<i>i,t</i>} (%)	1.73	1.49	4.04
LNBM _{<i>i,t</i>}	-1.02	-0.92	0.88
VOLROE _{<i>i,t</i>} (%)	26.53	10.43	19.79
PROFIT _{<i>i,t</i>} (%)	15.48	15.29	9.38

Notes. We consider a sample of stocks covered by IBES during the period 1980–2005 with valid annual earnings forecast records. COVERAGE_{*it*} is a measure of analyst coverage, defined as the number of analysts covering firm *i* at the end of year *t*. Analyst forecast bias (BIAS_{*jt*}) is the difference between the forecast analyst *j* in year *t* and the actual EPS, expressed as a percentage of the previous year's stock price. The consensus bias is expressed as a mean or median bias among all analysts covering a particular stock. Analyst forecast error (FERROR_{*jt*}) is the absolute difference between the forecast of analyst *j* in year *t* and the actual EPS, expressed as a percentage of the previous year's stock price. The forecast error is expressed as a mean or median bias among all analysts covering a particular stock. FDISP_{*it*} is analyst forecast dispersion, defined as the standard deviation of all analyst forecasts covering firm *i* in year *t*. LNSIZE_{*it*} is the natural logarithm of firm *i*'s market capitalization (price times shares outstanding) at the end of year *t*. SIGMA_{*it*} is the variance of daily (simple, raw) returns of stock *i* in year *t*. RETANN_{*it*} is the average monthly return on stock *i* in year *t*. LNBM_{*it*} is the natural logarithm of firm *i*'s book value divided by its market cap at the end of year *t*. To measure the volatility of ROE (VOLROE), we estimate an AR(1) model for each stock's ROE using a 10-year series of the company's valid annual ROEs. ROE_{*it*} is firm *i*'s return on equity in year *t*. ROE is calculated as the ratio of earnings in year *t* over the book value of equity. We calculate VOLROE as the variance of the residuals from this regression. PROFIT_{*it*} is the profitability of company *i* at the end of year *t*, defined as operating income over book value of assets. We exclude observations that fall to the left of the 25th percentile of the size distribution, observations with stock prices lower than \$5, and those for which the absolute difference between forecast value and the true earnings exceeds \$10.

medians, and standard deviations) are reported in Table I. The cross-sectional mean (median) analyst coverage of these stocks is about 21 (21) analysts and the standard deviation across stocks is about 10 analysts. The cross-sectional mean (median) bias is 0.027 (0.021) with a standard deviation of around 0.03.

The regression results are presented in Table II. We first present the results for the mean bias with just time and industry fixed effects in column (1), with industry, time, and brokerage-house fixed effects in column (2), and additionally with firm fixed effects in column (3). In column (1), the coefficient in front of COVERAGE is -0.0002 and is statistically significant at the 5% level of significance. In column (2), the coefficient is also -0.0002 and it is still statistically significant at the 5% level of significance.

TABLE II
REGRESSION OF CONSENSUS FORECAST BIAS ON COMPANY CHARACTERISTICS

Variables/model	Mean BIAS			Median BIAS		
	(1)	(2)	(3)	(4)	(5)	(6)
COVERAGE _{<i>i,t-1</i>}	-0.0002** (0.0001)	-0.0002** (0.0001)	0.0001 (0.0001)	-0.0002*** (0.0001)	-0.0002** (0.0001)	0.0001 (0.0001)
LNSIZE _{<i>i,t-1</i>}	0.0028*** (0.0009)	0.0026*** (0.0009)	0.0044*** (0.0014)	0.0028*** (0.0008)	0.0026*** (0.0008)	0.0042*** (0.0014)
SIGMA _{<i>i,t-1</i>}	-0.0093 (0.0062)	-0.0031 (0.0058)	0.0108*** (0.0039)	-0.0095 (0.0061)	-0.0061 (0.0059)	0.0108*** (0.0041)
RETANN _{<i>i,t-1</i>}	-0.1000*** (0.0199)	-0.0986*** (0.0199)	-0.0368*** (0.0133)	-0.0986*** (0.0192)	-0.0995*** (0.0198)	-0.0367*** (0.0135)
LNBM _{<i>i,t-1</i>}	0.0121*** (0.0016)	0.0115*** (0.0015)	0.0053*** (0.0013)	0.0118*** (0.0016)	0.0116*** (0.0015)	0.0049*** (0.0013)
VOLROE _{<i>i,t-1</i>}	0.0062*** (0.0019)	0.0061*** (0.0019)	0.0000 (0.0000)	0.0060*** (0.0019)	0.0059*** (0.0019)	0.0000 (0.0000)
PROFIT _{<i>i,t-1</i>}	0.0579*** (0.0095)	0.0571*** (0.0092)	0.0629*** (0.0115)	0.0578*** (0.0098)	0.0574*** (0.0095)	0.0619*** (0.0115)
SP500 _{<i>i,t-1</i>}	-0.0110*** (0.0024)	-0.0110*** (0.0024)	-0.0017 (0.0072)	-0.0110*** (0.0025)	-0.0110*** (0.0024)	-0.0026 (0.0075)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Brokerage fixed effects	No	Yes	Yes	No	Yes	Yes
Firm fixed effects	No	No	Yes	No	No	Yes
Observations	9,313	9,313	9,313	9,313	9,313	9,313

Notes. The dependent variable is BIAS, defined as a consensus forecast bias of all analysts tracking stock *i* in year *t*. Forecast bias is the difference between the forecast of analyst *j* in year *t* and the actual EPS, expressed as a percentage of the previous year's stock price. The consensus is obtained either as a mean or median bias. COVERAGE_{*i,t-1*} is a measure of analyst coverage, defined as the number of analysts covering firm *i* at the end of year *t*. LNSIZE_{*i,t-1*} is the natural logarithm of firm *i*'s market capitalization (price times shares outstanding) at the end of year *t*. SIGMA_{*i,t-1*} is the variance of daily (simple, raw) returns of stock *i* during year *t*. RETANN_{*i,t-1*} is the average monthly return on stock *i* in year *t*. LNBM_{*i,t-1*} is the natural logarithm of firm *i*'s book value divided by its market cap at the end of year *t*. To measure the volatility of ROE (VOLROE), we estimate an AR(1) model for each stock's ROE using a 10-year series of the company's valid annual ROEs. ROE_{*i,t-1*} is firm *i*'s return on equity in year *t*. ROE is calculated as the ratio of earnings in year *t* over the book value of equity. VOLROE is the variance of the residuals from this regression. PROFIT_{*i,t-1*} is the profitability of company *i* at the end of year *t*, defined as operating income over book value of assets. SP500_{*i,t-1*} is an indicator variable equal to one if stock *i* is included in the S&P500 index in year *t*. We exclude all observations that fall to the left of the 25th percentile of the size distribution, observations with stock prices lower than \$5, and those for which the absolute difference between forecast value and the true earnings exceeds \$10. All regressions include three-digit SIC industry fixed effects and year fixed effects. Some specifications also include brokerage house and firm fixed effects. Standard errors (in parentheses) are clustered at the industry level.

***, **, * 1%, 5%, and 10% statistical significance.

However, the coefficient turns positive and statistically nonsignificant when we include firm fixed effects. Because coverage is fairly persistent, it may be that a fixed-effects approach is not picking up the right variation, in contrast to the cross-sectional approach. So depending on the controls used, we find that a decrease in coverage by one analyst leads to an increase in bias of anywhere from 0.0002 (two basis points) to none. The bias for a typical stock is about

0.027 (2.7%) with a standard deviation across stocks of about 0.03 (3%). Hence, these estimates obtained from cross-section regressions suggest only a small increase in bias of about zero to sixty basis points as a fraction of the cross-sectional standard deviation of bias as we decrease coverage by one analyst, though some are very precisely measured. The results using the median bias instead of the mean bias are reported in columns (4), (5), and (6). Again, there is little difference in the coefficient on *COVERAGE*.

The other control variables also come in significantly in these regressions. Bias increases with firms' sizes, book-to-market ratios, volatilities of return on equity, and profits. Bias is lower for firms with high returns and for firms in the S&P 500 index. The sign on stock return volatility is ambiguous, depending on whether firm fixed effects are included.

Of course, as we explained in the Introduction, these OLS regressions are difficult to interpret because of the endogeneity of analyst coverage. If stocks that attract lots of coverage are stocks that analysts are likely to be excited about, then these OLS estimates are biased downward. In contrast, if stocks covered by only a few analysts are likely under-the-radar stocks that analysts have to be very excited about to initiate coverage on, then these OLS estimates of the competition effect are biased upward. Estimating this regression using stock fixed effects is not an adequate solution to the endogeneity critique because analyst coverage tends to be a fairly persistent variable and analysts drop coverage on stocks when the stock is no longer doing well (e.g., McNichols and O'Brien [1997]).

Hence, we rely on a natural experiment to sort out these endogeneity issues. We use mergers of brokerage houses as our experiment, on the premise that mergers typically lead to a reduction in analyst coverage on the stocks that were covered by both the bidder and target firms premerger. If a stock is covered by both firms before the merger, they will get rid of at least one of the analysts, usually the target analyst. It is to this experiment that we now turn.

IV. BACKGROUND ON MERGERS

We begin by providing some background on these mergers. We identify mergers among brokerage houses by relying on information from the SDC Mergers and Acquisition database. We start with the sample of 32,600 mergers of financial institutions. Next,

we choose all the mergers in which the target company belongs to the four-digit SIC code 6211 ("Investment Commodity Firms, Dealers, and Exchanges"). This screen reduces our sample to 696 mergers. Subsequently, we manually match all the mergers with IBES data. This match identifies 43 mergers with both bidder and target covered by IBES. Finally, we select only those mergers where both merging houses analyze at least two of the same stocks—otherwise, there is little scope for our instrumental variables approach below. With this constraint, our search produces fifteen mergers, which we break down to parties involved: bidder and target. We provide further details about these mergers in the Appendix.

Of the fifteen mergers, six are particularly big in the sense that the merging houses both tend to be big firms and have coverage premerger on a large number of similar stocks. The first of these big mergers is Merrill Lynch acquiring on September 10, 1984, a troubled Becker Paribas that was having problems with its own earlier merger to another firm. The second is Paine Webber acquiring Kidder Peabody on December 31, 1994. Kidder was in trouble and had fired a good part of its workforce before the merger in the aftermath of a major trading scandal involving its government bond trader, Joseph Jett. Kidder's owner, General Electric, wanted to sell the company and Paine Webber (a second-tier brokerage house) wanted to buy a top-tier investment bank with a strong research department. The third is Morgan Stanley acquiring Dean Witter Reynolds on May 31, 1997. Morgan Stanley was portrayed as wanting to get in on the more down-market retail brokerage operations of Dean Witter. The fourth is Smith Barney (Travelers) acquiring Salomon Brothers on November 28, 1997. This is viewed as a synergy play led by Sandy Weill.

The fifth and sixth mergers involved Swiss banks trying to diversify their lines of business geographically into the American market. These mergers happened within a few months of each other. Credit Suisse First Boston acquired Donaldson Lufkin and Jenrette on October 15, 2000. A few months later, on December 10, 2000, UBS acquired Paine Webber. The anecdotal descriptions of the motivations for these mergers provide comfort in our proposed experiment, which is that these mergers provide a change in competition that is unrelated to some underlying unobservable of the biases in the stocks.

In Table III, we provide a number of key statistics regarding all fifteen mergers. In Panel A, we summarize the names, the

TABLE III
DESCRIPTIVE STATISTICS FOR MERGERS

Panel A: Mergers used in the analysis and stocks covered					
Brokerage house	Merger number	Merger date	# stocks (bidder)	# stocks (target)	# stocks (bidder and target)
Merrill Lynch	1	9/10/1984	762		173
Becker Paribas				288	
Paine Webber	2	12/31/1994	659		234
Kidder Peabody				545	
Morgan Stanley	3	05/31/1997	739		251
Dean Witter Reynolds				470	
Smith Barney (Travelers)	4	11/28/1997	914		327
Salomon Brothers				721	
Credit Suisse	5	10/15/2000	856		307
First Boston					
Donaldson Lufkin and Jenrette				595	
UBS Warburg	6	12/10/2000	596		213
Dillon Read					
Paine Webber				487	
Chase Manhattan	7	12/31/2000	487		80
JP Morgan				415	
Wheat First Securities	8	10/31/1988	178		8
Butcher & Co., Inc.				66	
EVEREN Capital	9	1/9/1998	178		17
Principal Financial Securities				142	
DA Davidson & Co.	10	2/17/1998	76		8
Jensen Securities				53	
Dain Rauscher	11	4/6/1998	360		26
Wessels Arnold & Henderson				135	
First Union	12	10/1/1999	274		21
EVEREN Capital				204	
Paine Webber	13	6/12/2000	516		28
JC Bradford				182	
Fahnestock	14	9/18/2001	117		5
Josephthal Lyon & Ross				91	
Janney Montgomery Scott	15	3/22/2005	116		10
Parker/Hunter				54	

TABLE III
(CONTINUED)

Brokerage house	Panel B: Career outcomes of analysts after mergers						
	Merger number	# analysts		# analysts after merger			
		Prior	After	Retained in the house	Left to another house	Exited sample (fired)	New analysts
Merrill Lynch	1	90	98	84	0	5	13
Becker Paribas		27	—	1	11	15	—
Paine Webber	2	50	57	42	1	7	6
Kidder Peabody		51	—	9	28	14	—
Morgan Stanley	3	70	92	61	2	7	26
Dean Witter Reynolds		35	—	5	16	14	—
Smith Barney (Travelers)	4	91	140	70	6	15	27
Salomon Brothers		76	—	43	20	13	—
Credit Suisse First Boston	5	120	146	93	5	22	35
Donaldson Lufkin Jenrette		77	—	18	17	42	—
UBS Warburg Dillon Read	6	94	118	80	5	9	0
Paine Webber		64	—	38	8	17	—
Chase Manhattan	7	64	106	48	5	11	24
JP Morgan		50	—	34	1	15	—
Wheat First Securities	8	13	21	13	0	0	8
Butcher & Co., Inc.		13	—	3	3	7	—
EVEREN Capital	9	27	31	21	4	2	8
Principal Financial Securities		18	—	2	6	10	—
DA Davidson & Co.	10	6	8	4	1	1	0
Jensen Securities		4	—	4	0	0	—
Dain Rauscher	11	39	36	19	9	11	6
Wessels Arnold & Henderson		15	—	11	0	4	—
First Union	12	35	54	26	2	7	16
EVEREN Capital		32	—	12	10	10	—
Paine Webber	13	54	55	37	9	8	18
JC Bradford		22	—	0	14	8	—
Fahnestock	14	14	16	7	1	6	9
Josephthal Lyon & Ross		14	—	0	5	9	—
Janney Montgomery Scott	15	13	15	11	1	1	3
Parker/Hunter		5	—	1	0	4	—

TABLE III
(CONTINUED)

Merger	Panel C: Percentage of stocks covered by analysts from bidder and target houses after mergers	
	Percentage of stocks (bidder)	Percentage of stocks (target)
	(1)	(2)
1	85.7	1.1
2	73.7	15.8
3	66.3	5.4
4	50.0	30.7
5	63.7	12.3
6	67.8	32.3
7	45.3	32.1
8	61.9	14.3
9	67.7	6.5
10	50.0	50.0
11	52.8	30.6
12	48.1	22.2
13	67.3	0
14	43.8	0
15	73.3	6.7

Notes. Panel A includes the names of brokerage houses involved in mergers, the date of the merger, and the number of stocks covered by either brokerage house or both of them prior to the merger. Panel B breaks down the merger information at the analyst level. We include number of analysts employed in the merging brokerage houses prior to merger and after the merger as well as the detailed information on the career outcomes of the analysts after the merger. Panel C calculates the percentage of analysts from the merging houses that cover the same stock after the merger. We restrict our sample of stocks to those which were covered by both the bidder and the target house.

dates, and the number of stocks covered by the bidder and target individually and the overlap in the coverage. For instance, in the merger involving Paine Webber and Kidder Peabody, Paine Webber covered 659 stocks and Kidder covered 545 stocks. There was a 234-stock overlap in their coverage. As a result, the merger leads to a potential decrease of around one analyst for a large number of stocks. The size of our treatment sample, the number of firms covered by both merging houses, ranges from a low of five stocks in the merger involving Fahnestock and Josephthal Lyon and Ross to a high of 327 stocks in the Smith Barney and Salomon Brothers deal. Notice that the big six mergers described above give us much of the variation in terms of the number of treatment stocks. In total, we have a significant treatment sample with which to identify our effect.

To better support the premise that mergers lead to less analyst coverage in the treatment sample via job turnover, we

examine career outcomes of analysts employed by merging houses. Panel B presents the results with the breakdown of career outcomes of analysts employed by both the bidder and target houses. A few observations can be noted. First, the big mergers affected a very significant number of analysts. The largest of the mergers—between Credit Suisse First Boston and Donaldson Lufkin and Jenrette—affected almost 200 analysts. The smallest merger in terms of analysts affected is Davidson and Jensen with ten. Given that in our sample the average brokerage house employs approximately fifteen analysts, a number of our mergers constituted important events in the analyst industry.

Second, as expected, mergers generally reduce the number of analysts covering stocks. For example, the two brokerage houses involved in the first merger, Paine Webber and Kidder Peabody, employed a total of 101 analysts prior to merger. After Paine Webber acquired Kidder Peabody, the employment in the joint entity decreased to 57 analysts. Third, the majority of the employment reduction comes from the closure of the target house. In particular, out of 51 analysts employed by Kidder, only nine were retained in the new company, and 28 left to a different house, whereas fourteen exited the sample, which we interpret as a firing decision.

In Panel C, we confirm more precisely that for stocks covered by both houses premerger, it is usually the analyst in the bidding house who remains whereas the target analyst is let go. In the first column of Panel C, we report for the treatment sample, stocks covered by both houses, the fraction that are covered by the bidder analyst after the merger. In the second column, we report the fraction covered by the target analysts after the merger of the treatment sample. In the Paine Webber and Kidder merger, for stocks covered premerger by both houses, it is the target analyst who is indeed the redundant one that gets fired—the corresponding figures are 73.7% for the bidder analysts and only 15.8% for the target analysts. Similarly big gaps exist for most of the other mergers. This gap is much smaller in the Davidson and Jensen merger, 50% for the bidder and 50% for the target. Nonetheless, from Panel B, it still appears that there were fewer analysts working for the merged entity than for the two houses beforehand.

V. EMPIRICAL DESIGN

Our analysis of the effect of competition on analyst forecast bias utilizes a natural experiment involving brokerage house

mergers. The outcome of this process is reduction in the number of analysts employed in the combined entity compared to the total number of analysts employed in bidder and target entities prior to merger. As a result, the number of analysts covering a stock that was covered by both houses before the merger (our treatment sample) should drop as one of the redundant analysts is let go or reallocated to another stock (or maybe even both are let go), and thus the competition in the treatment sample decreases. The questions then are whether there is a decrease in competition among analysts around merger events and whether this decrease is associated with an economically significant effect on average consensus bias.

Our empirical methodology requires that we specify a representative window around the merger events. In choosing the proper estimation window, we face a trade-off, unlike most other event studies, which would have us focus on a very narrow window. As is the case with most event studies, choosing a window that is too long may incorporate information that is not really relevant to the event under consideration. But in our case, choosing too short a window means we may lose observations, because analysts may not issue forecasts on the same date or with the same frequency. We want to keep a long enough window to look at the change in the performance of all analysts before and after the merger.

To this end, we use a two-year window, with one year of data selected for each pre- and postevent period. Most analysts will typically issue at least one forecast within a twelve-month window. Given that in each of the two windows one analyst could issue more than one forecast, we retain only the forecast that has the shortest possible time distance from the merger date. In addition, because we are interested in the effect of merger on various analyst characteristics, we require that each stock be present in both windows around the merger. As a result, for every stock we note only two observations—one in each window of the event.

This one year before and one year after the merger event having been chosen, one then has to factor in the fact that coverage and the average stock bias may vary from one year to the next. In other words, to identify how the merger affected coverage in the stocks covered by both houses premerger and how the bias in these stocks then also changed, one needs to account for the fact that there may be natural changes from year to year in coverage and bias for these stocks.

A standard approach to deal with these time trends is based on the difference-in-differences (DID) methodology. In this approach, the sample of stocks is divided into treatment and control groups. In the context of our paper, the treatment group includes all stocks that were covered by both brokerage houses before the merger. The control group includes all the remaining stocks. If we denote the average observed characteristics in the treatment (T) and control (C) groups in the pre- and postevent periods by $C_{T,1}$, $C_{T,2}$, $C_{C,1}$, and $C_{C,2}$, respectively, the partial effect of change due to merger can be estimated as

$$(1) \quad \text{DID} = (C_{T,2} - C_{T,1}) - (C_{C,2} - C_{C,1}).$$

Here the characteristics might be analyst coverage or bias. By comparing the time changes in the means for the treatment and control groups, we allow for both group-specific and time-specific effects. This estimator is unbiased under the condition that the merger is not systematically related to other factors that affect C .

A potential concern with the above estimator is the possibility that the treatment and control groups may be significantly different from each other and thus the partial effect may additionally capture the differences in the characteristics of the different groups. For example, the average stocks in both groups may differ in terms of their market capitalizations, value characteristics, or past return characteristics. For instance, it might be that companies with good recent returns lead analysts to cover their stocks and to be more optimistic about them. Hence, we want to make sure that past returns of the stocks in the treatment and control samples are similar. We are also worried that higher analyst coverage stocks may simply be different than lower analyst coverage stocks for reasons unrelated to our competition effect. So we will also want to keep the premerger coverage characteristics of our treatment sample similar to those of our control sample.

To account for such systematic differences across the two samples, we use the matching technique similar to that used in the context of IPO event studies or characteristic-based asset pricing. In particular, each stock in the treatment sample is matched with its own benchmark portfolio obtained using the sample of stocks in the control group. We expect our controls typically to do a better job at capturing our true effect by netting out unobserved heterogeneity.

To construct the benchmark, we first sort stocks into tercile portfolios according to their market capitalizations. Next, we sort stocks within each size portfolio according to their book-to-market ratios. This sort results in nine different benchmark portfolios. Further, we sort stocks in each of the nine portfolios into tercile portfolios according to their past returns, which results in 27 different benchmark portfolios. Finally, we sort stocks in each of the 27 portfolios into tercile portfolios according to their analyst coverage. Overall, our benchmark includes 81 portfolios.

Using the above benchmark specification, we then construct the benchmark-adjusted DID estimator (BDID). In particular, for each stock i in the treatment sample, the partial effect of change due to merger is calculated as the difference between two components,

$$(2) \quad \text{BDID}^i = (C_{T,2}^i - C_{T,1}^i) - (BC_{C,2}^i - BC_{C,1}^i),$$

where the first component is the difference in characteristics of stock i in the treatment sample moving from the premerger to the postmerger period. The second component is the difference in the average characteristics of the benchmark portfolios that are matched to stock i along the size/value/momentum/coverage dimensions. In general, the results are comparable if we use benchmarks matched along any subset of the characteristics. To assess the average effect for all stocks in the treatment sample, one can then take the average of all individual BDIDs.

One final issue that we need to account for is that a few of the mergers occurred within several months of each other (e.g., the fifth and sixth mergers occurred on October 15, 2000, and December 10, 2000, respectively). As a result, it might be difficult to separate out the effects of these two mergers individually. As the baseline case, we decided for simplicity to treat each merger separately in our analysis. However, we have also tried robustness checks in which we group mergers occurring close together in time and treat them as one merger. For instance, we consider a one-year window before the third merger on October 15, 2000, as the premerger period and the one-year window after the fourth merger on December 10, 2000, as the postmerger period. As a result, the treatment sample is the union of the 307 stocks jointly covered by Credit Suisse and DLJ and the 213 stocks covered by UBS and Paine Webber. There is potentially some overlap of these two subsets of stocks, and hence it might be the case that some of

TABLE IV
SUMMARY STATISTICS FOR THE TREATMENT SAMPLE

Variable	Cross-sectional mean (1)	Cross-sectional median (2)	Cross-sectional st. dev. (3)
COVERAGE _{<i>i,t</i>}	21.12	20	9.45
Mean BIAS _{<i>i,t</i>} (%)	2.79	2.24	3.10
Median BIAS _{<i>i,t</i>} (%)	2.74	2.21	3.19
Mean FERROR _{<i>i,t</i>} (%)	3.40	2.52	2.90
Median FERROR _{<i>i,t</i>} (%)	3.33	2.43	2.99
FDISP _{<i>i,t</i>} (%)	0.75	0.40	0.94
LNSIZE _{<i>i,t</i>}	8.39	8.37	1.60
SIGMA _{<i>i,t</i>} (%)	41.00	35.86	21.02
RETANN _{<i>i,t</i>} (%)	1.74	1.52	4.13
LNBM _{<i>i,t</i>}	-1.03	-0.92	0.91
VOLROE _{<i>i,t</i>} (%)	25.32	9.89	43.40
PROFIT _{<i>i,t</i>} (%)	15.52	15.25	9.22

Notes. We consider all stocks covered by two merging brokerage houses around the one-year merger event window. COVERAGE_{*it*} is a measure of analyst coverage, defined as the number of analysts covering firm *i* at the end of year *t*. Analyst forecast bias (BIAS_{*jt*}) is the difference between the forecast analyst *j* at time *t* and the actual EPS, expressed as a percentage of the previous year's stock price. The consensus bias is expressed as a mean or median bias among all analysts covering a particular stock. Analyst forecast error (FERROR_{*jt*}) is the absolute difference between the forecast analyst *j* at time *t* and the actual EPS, expressed as a percentage of the previous year's stock price. The forecast error is expressed as a mean or median bias among all analysts covering a particular stock. FDISP_{*it*} is analyst forecast dispersion, defined as the standard deviation of all analyst forecasts covering firm *i* at time *t*. LNSIZE_{*it*} is the natural logarithm of firm *i*'s market capitalization (price times shares outstanding) at the end of year *t*. SIGMA_{*it*} is the variance of daily (simple, raw) returns of stock *i* during year *t*. RETANN_{*it*} is the average monthly return on stock *i* during year *t*. LNBM_{*it*} is the natural logarithm of firm *i*'s book value divided by its market cap at the end of year *t*. To measure the volatility of ROE (VOLROE), we estimate an AR(1) model for each stock's ROE using a 10-year series of the company's valid annual ROEs. ROE_{*it*} is firm *i*'s return on equity in year *t*. ROE is calculated as the ratio of earnings during year *t* over the book value of equity. We calculate VOLROE as the variance of the residuals from this regression. PROFIT_{*it*} is the profitability of company *i* at the end of year *t*, defined as operating income over book value of assets. We exclude observations with stock prices lower than \$5 and those for which the absolute difference between the forecast value and the true earnings exceeds \$10.

these stocks will experience a greater decline in analyst coverage to the extent that they have more than two redundant analysts. However, these alterations do not affect our baseline results.

Table IV presents summary statistics for the treatment sample in the two-year window around the merger. The characteristics of the treatment sample are similar to those reported in Table I for the OLS sample. For instance, the coverage is about 21 analysts for the typical stock. The mean bias is 2.79% with a standard deviation of around 3.10%. These figures, along with those of the control variables, are fairly similar across these two samples. This provides comfort that we can then relate the economic effect of competition obtained from our treatment sample to the OLS estimates presented in Table II.

VI. RESULTS

VI.A. *Analyst Coverage and Optimism Bias*

We first verify the premise of our natural experiment by measuring the change in analyst coverage for the treatment sample from the year before the merger to the year after. We expect these stocks to experience a decrease in coverage.

Panel A of Table V (column (1)) reports the results of this analysis. We present the DID estimator for coverage using our benchmarking technique—size, book-to-market, return, and coverage matched. We observe a discernible drop in coverage due to merger by around 1.02 analysts, using the DID estimator, and the level of the drop of between one and two analysts is in line with our expectations. This effect is significant at the 1% level.

One can think of this finding as essentially the first stage of our estimation. The effect is economically and statistically significant in the direction predicted, and hence confirms the premise of our natural experiment. We will focus on this number in our discussion of the economic effect of competition below.

We next look at how the optimism bias changes for the treatment sample across the mergers. These results are presented in Panel A of Table V. We present the findings in column (2) for the mean BIAS and in column (3) for the median BIAS. Using the DID estimator, we find an increase in optimism bias of 0.0013 for the mean bias (significant at the 10% level) and 0.0016 for the median bias (significant at the 5% level).

Using the estimates obtained above, a conservative estimate is that the mean optimism bias increases by about thirteen basis points (as a result of reducing coverage by one analyst). As we mentioned earlier, the sample for the natural experiment is similar to that of the OLS by construction—the typical stock has a bias of around 2.7% and the standard deviation of the optimism bias is also around 3%. This means that the estimate of the competitive effect from our natural experiment is approximately six to seven times as large as that from the OLS estimates. This is a sizable difference and suggests that the OLS estimates are biased downward, consistent with the documented selection bias that stocks that attract lots of coverage are likely to have more optimistic analysts.

One could argue that our mean bias effect might be driven by selection through which one of the two analysts from the

TABLE V
CHANGE IN STOCK-LEVEL COVERAGE AND BIAS: DID ESTIMATOR

Panel A: Coverage and bias			
	Coverage (1)	Mean BIAS (2)	Median BIAS (3)
SIZE/BM/RET/NOAN-matched	-1.021*** (0.179)	0.0013* (0.0007)	0.0016** (0.0008)
Panel B: Change in forecast bias: DID estimator without analysts from merging houses			
	Mean BIAS (1)	Median BIAS (2)	
SIZE/BM/RET/NOAN-matched	0.0011** (0.0005)	0.0012** (0.0005)	
Panel C: Change in forecast bias: Conditioning on initial coverage			
	Mean BIAS (1)	Median BIAS (2)	
SIZE/BM/RET/NOAN-matched (coverage ≤ 5)	0.0078** (0.0036)	0.0096** (0.0044)	
SIZE/BM/RET/NOAN-matched (coverage > 5 and ≤ 20)	0.0017* (0.0011)	0.0020* (0.0011)	
SIZE/BM/RET/NOAN-matched (coverage > 20)	0.0003 (0.0013)	0.0007 (0.0013)	

Notes. $N = 1,656$ in Panels A and B. We measure analyst coverage as the number of analysts covering firm i at the end of year t . For all mergers, we split the sample of stocks into those covered by both merging brokerage houses (treatment sample) and those not covered by both houses (control sample). We also divide stocks into premerger period and postmerger period (one-year window for each period). For each period we further construct benchmark portfolios using the control sample based on stocks' size (SIZE), book-to-market ratio (BM), and average past year's returns (RET). Our benchmark assignment involves three portfolios in each category. Each stock in the treatment sample is then assigned to its own benchmark SIZE/BM/RET/NOAN-matched). Next, for each period, we calculate the cross-sectional average of the differences in analyst stock coverage across all stocks in the treatment sample and their respective benchmarks. Finally, we calculate the difference in differences between postevent period and pre-event period (DID estimator). Our sample excludes observations with stock prices lower than \$5 and those for which the absolute difference between forecast value and the true earnings exceeds \$10. In Panel B, we exclude from our sample all analysts employed in the merging houses. Panel C presents our results by cuts on initial coverage. There are three groups: lowest coverage (≤ 5), medium coverage (> 5 and ≤ 20) and highest coverage (> 20). Standard errors (in parentheses) are clustered at the merger groupings.

***, **, * 1%, 5%, and 10% statistical significance.

merging firms covering the stock gets fired. It might be that the less optimistic analyst gets fired and hence the bias might be higher as a result. Another possibility could be that analysts employed by the merging houses may compete for the job in the new merged house and thus they may strategically change their reporting behavior.

We deal with these issues in two ways. The first one, because we have turnover data, is simply to check whether the merging brokerage houses selectively fire analysts who are less optimistic. We do not find such a selection bias. The second and more direct way to deal with this concern is to look only at the change in the bias for the analysts covering the same stocks but not employed by the merging firms. The findings are in Panel B of Table V. We report the change in bias for the treatment sample, but now the bias is calculated using only the forecasts of the analysts not employed by the merging houses. The figures are very similar to the main findings—only slightly smaller in some instances by a negligible amount. The mean bias increases by eleven basis points, the median bias increases by twelve basis points, and both are significant at the 5% level. Collectively, these findings provide comfort that our main results are not spuriously driven by outliers or by selection biases.

We next test a key auxiliary prediction that will further buttress our identification strategy. We check whether the competition effect is more pronounced for stocks with smaller analyst coverage. The idea is that the more analysts cover a stock, the less the loss of an additional analyst matters, akin to the Cournot view of competition. For instance, in the independence rationale of Gentzkow and Shapiro, when there are already many analysts, losing one would not change much the likelihood of drawing an independent analyst. In contrast, when there are only a few analysts to begin with, losing one analyst could really affect the likelihood of getting an independent supplier of information.

However, note that if collusion is possible, then we might expect a nonlinear relationship between bias and coverage. Suppose that collusion is easier when there are only a few analysts. Under this scenario, going from one to two analysts may not have an effect because the two can collude. And we might find more of an effect when going from five to six analysts if the sixth analyst does not collude. With collusion, it might be that we expect the biggest effect for stocks covered by a moderate number of analysts—that is, an inverted U-shape with the effect being greatest for medium-coverage stocks.

We examine this issue in Panel C of Table V using the same DID framework as before. We divide initial coverage into three groups: less than or equal to five analysts, between six and twenty analysts, and more than twenty analysts. Column (1) reports the results using mean bias. We expect and find that the effect is

significantly smaller when there are a lot of analysts covering. The effect is greatest for the first group (less than or equal to five analysts). The mean bias increases by 78 basis points and the median bias by 96 basis points, and both are significant at the 5% level. The next largest effect is in the second group (more than five and less than or equal to twenty): The mean bias increases by seventeen basis points and the median bias by twenty basis points. Both are significant at the 10% level. Finally, the effect is much smaller for the highest-coverage group: the mean bias increases by three basis points and the median bias by seven basis points and neither of these point estimates is statistically significant. In sum, the evidence is remarkably comforting, as it conforms well to our priors on competition being more important when there are fewer analysts around. This result reassures us that our estimate is a sensible one.⁸

Next, we delve deeper into our results in Table V by analyzing plots in event time of the change in coverage and bias. This will also allow us to gauge the robustness of the parallel trend assumption required for the difference-in-differences approach. To this end, we consider the event window with three periods before and three periods after the merger. For each event-window date, we calculate the difference in coverage and median bias between the treatment and control groups and plot it against the event time. The left-hand side of Figure I presents the results for coverage and the right-hand side for bias. We report the results separately for the three subgroups of stocks sorted according to coverage (Panels A–C) and then for the entire sample (Panel D).

We do not find support for pretrends driving our results. In particular, both for coverage and for bias, the difference between the treatment and control groups is stable before the event date, and also after the event date. Moreover, we confirm the results from Panel C of Table V. Consider Panel A of Figure I, which shows the results for the low-coverage stocks. Note that the mergers cause a relative drop of about one analyst on the event date (from time -1 to time 0) and an increase in bias of about eighty basis points for low-coverage stocks (from time -1 to time 0). Our matching in terms of premerger coverage characteristics is nearly perfect here, and there is also little difference in the

8. The results are not affected by a particular cutoff level for the number of analysts. The results are generally declining in a nonlinear way with an increase of coverage.




FIGURE I

Trend of Analyst Coverage and Bias in the Treatment Sample (Net of Control)

We show the trend of average analyst coverage and forecast bias in the treatment sample net of the control group (in a given year) up to three years before and after the merger event. Panel A is for stocks with low analyst coverage (less than six analysts); Panel B is for stocks with medium analyst coverage (six to twenty analysts); Panel C is for stocks with high analyst coverage (above twenty analysts). Panel D documents aggregate results. Dotted lines illustrate 95% confidence intervals.




FIGURE I
(CONTINUED)

premerger mean bias between the treatment and control groups. There are no discernible pretrends or posttrends. We also observe a similar degree of statistical significance, as illustrated by two-standard-deviation bands represented by dotted lines.

Panel B shows the results for the medium-coverage stocks. Again, the matching in premerger coverage characteristics is nearly perfect, with the treatment sample having a slightly higher premerger coverage by about one-half analyst than the control sample. Again there is nearly perfect matching in the mean bias premerger between the treatment and control samples. We see a drop of about one analyst on the event date and an increase in bias by about twenty basis points. Importantly, there are no discernible pretrends, though there is a slight posttrend in the coverage drop, which continues to drop a couple of years after the event date. But by and large, the figures in Panel B support the validity of the experiment.

In Panel C, we look at the high-coverage stocks. Notice here that the treatment sample has a much higher coverage than the control sample. Part of the reason for this is that the mergers in our sample involve big brokerage houses, which cover very big stocks. As a result, it is difficult to get a match for these big stocks identical to what we could for the lower-coverage stocks. This might explain why the premerger bias of the treatment sample is slightly lower than that of the control group, which would very much be consistent with the competition mechanism. We believe this difference is not contributing any biases to our estimates related to the merger. Absent this difference, we see roughly the same picture of a one-analyst drop on the event date and a slight increase in the bias, which is not significantly different from zero. The picture that emerges is similar to that of Panel B.

Finally, in Panel D, we draw the same pictures using the entire sample. We see no significant pretrends in either coverage or bias and an observable event day drop in coverage and increase in bias. These figures provide comfort that our results in Table V are not driven by pretrends.

VI.B. Validity of the Natural Experiment

The economic significance of our results strictly depends on the validity of our natural experiment. Although the results in Figure I are a start in the direction of comforting us on the validity of our natural experiment, in this section we report a number of

further tests, which collectively provide strong support for our experiment.

First, we separately estimate our effect using the six biggest mergers. The results are very similar in that the conservative estimates are a one-analyst drop in coverage associated with a 0.0017 increase in bias. Further, we estimate our effect separately for each of the fifteen mergers. Each of the fifteen mergers experienced a decline in coverage using the most conservative DID estimate. Hence, our result is not driven by outliers—there is a distinct coverage drop with mergers. Clearly, the fact that fifteen out of fifteen mergers experienced drops suggests that our effect is robustly significant in a nonparametric sense. Similarly, we find that twelve (thirteen) of the fifteen mergers experienced an increase in mean (median) bias using the most conservative DID estimate. It is important to emphasize that because these mergers occur throughout our entire sample, our effects are not due to any particular macroeconomic event such as a recession or boom.

Second, given that the optimism measure is constructed as a difference between an analyst forecast and actual earnings, one could worry that our results are driven by differences in the actual earnings and not in reported forecasts. To rule out such a possibility, we test whether merger events lead to differential changes in earnings between treatment and control groups. Panel A of Table VI reports the results separately for the mean and median earnings. We find no evidence that competition causes changes in actual earnings.

Third, our experiment relies on the validity of the matching procedure between firms in treatment and control groups. In general, our findings do not raise a major problem with the matching, but to provide further robustness that the differences we observe do not actually capture the *ex ante* differences in various observables, we report similar DID estimators for other response variables—Tobin's *Q*, size, returns, volatility, profitability, and sales. The results in Panel B of Table VI show that none of the important observables is significantly affected by the merger event. These results are comforting, as they confirm the validity of our matching procedure.

Fourth, the nature of our experiment requires that the same company be covered by two merging houses. To ensure that our effects are not due merely to the fact that the selection of the companies to brokerage houses is not random, we reexamine our evidence by focusing on stocks that are covered by one of the

TABLE VI
VALIDITY OF THE NATURAL EXPERIMENT

Panel A: Change in earnings			
	Mean EARN (1)	Median EARN (2)	
DID estimator (SIZE/BM/RET/NOAN-matched)	−0.0002 (0.0005)	−0.0002 (0.0005)	
Panel B: Change in firm characteristics			
Stock characteristic	SIZE/BM/RET/NOAN-matched		
Tobin's Q	0.0397 (0.1138)		
Size	30.52 (20.63)		
Returns	0.0015 (0.0012)		
Volatility	−0.0069 (0.0054)		
Profitability	−0.0593 (0.0597)		
Log(sales)	0.0046 (0.0090)		
Panel C: Change in forecast bias for non-overlapping stocks			
	Coverage (1)	Mean BIAS (2)	Median BIAS (3)
SIZE/BM/RET/NOAN-matched	0.080 (0.106)	0.0001 (0.0004)	0.0001 (0.0004)
Panel D: Change in forecast bias for non-overlapping stocks as a control			
	Mean BIAS (1)	Median BIAS (2)	
SIZE/BM/RET/NOAN-matched	0.0017*** (0.0006)	0.0018*** (0.0006)	

Notes. $N = 1,656$ in Panels A, B, and D. In Panel A, we measure analyst earnings ($EARN_{it}$) as the actual EPS expressed as a percentage of the previous year's stock price. For all mergers, we split the sample of stocks into those covered by both merging brokerage houses (treatment sample) and those not covered by both houses (control sample). We also divide stocks into premerger period and postmerger period (one-year window for each period). For each period we further construct benchmark portfolios using the control sample based on stocks' size (SIZE), book-to-market ratio (BM), average past year's returns (RET), and analyst coverage (NOAN). Our benchmark assignment involves three portfolios in each category. Each stock in the treatment sample is then assigned to its own benchmark portfolio (SIZE/BM/RET/NOAN-matched). Next, for each period, we calculate the cross-sectional mean and median of the differences in earnings across all stocks in the treatment sample and their respective benchmarks. Finally, we calculate the difference in differences between postevent period and pre-event period (DID estimator). In Panel B, we provide the DID estimator for various corporate characteristics, including Tobin's Q, asset size, stock returns, volatility, profitability, and log sales. In Panel C, the treatment sample is constructed based on the stocks that are covered by one but not both merging houses. In Panel D, the control sample is constructed using the stocks that are covered by one but not both merging houses. Our sample excludes observations with stock prices lower than \$5 and those for which the absolute difference between forecast value and the true earnings exceeds \$10. Standard errors (in parentheses) are clustered at the merger groupings.

***, **, * 1%, 5%, and 10% statistical significance.




FIGURE II

Kernel Densities of Differences between Treatment and Control before and after Merger

We show Epanechnikov kernel densities of differences in forecast bias between treatment and control groups for the period before and after the merger. The bandwidth for the density estimation is selected using the plug-in formula of Sheather and Jones (1991). The rightward shift of the distribution after the merger is significant because the hypothesis of equality of distributions is rejected at 1% level using the Kolmogorov–Smirnov test.

merging houses, but not by both. We show in Panel C of Table VI that the average stock coverage does not change significantly on the event date across these treatment and control groups, and the change in the bias is statistically not different from zero. We further apply this setting to validate the quality of our control group. Specifically, in Panel D of Table VI, we show that using stocks covered by only one of the two merging houses as a control group does not change the nature of our results. In fact, the results become slightly stronger than those in our baseline specification.

Fifth, we examine whether competition changes the entire distribution of forecasts. To this end, we plot Epanechnikov kernel densities of bias in the treatment group relative to the control group before and after the merger. The bandwidth for the density estimation is selected using the plug-in formula of Sheather and Jones (1991). Figure II presents the results. We observe a significant rightward shift in the entire distribution of bias in the

postmerger period. The rightward shift of the distribution after the merger is significant because the hypothesis of equality of distributions is rejected at the 1% level using the Kolmogorov–Smirnov test. Moreover, the average relative bias becomes strictly positive, consistent with our earlier findings. These results suggest that the findings of our experiment are not driven by outliers and further indicate that the merger mainly causes previously unbiased analysts to become biased.

VI.C. Robustness

In this section, we report a number of tests that confirm the robustness of our results.

An alternative econometric approach to capturing the effect of change in the bias due to merger that we consider is to estimate the following regression model:

$$(3) \quad C_i = \alpha + \beta_1 \text{Merge}_i + \beta_2 \text{Affected}_i + \beta_3 \text{Merge}_i \times \text{Affected}_i + \beta_4 \text{Controls}_i + \varepsilon_i,$$

where C is the characteristic that may be subject to merger; Merge is an indicator variable, equal to one for observations after the merger, and zero otherwise; Affected is an indicator variable equal to one if stock i is affected by the merger, and zero otherwise; and Controls is a vector of stock-specific covariates affecting C . In this specification, the coefficient of primary interest is β_3 , which captures the partial effect of change due to merger; in the version with additional controls its value is similar in spirit to the DID estimator in equation (2). By including additional controls, we account for any systematic differences in stocks that may affect the partial effect of change due to merger.

Importantly, the regressions include merger fixed effects and industry fixed effects, which ensures comparability of the samples across various time-invariant characteristics. We also include brokerage fixed effects and firm fixed effects, which help us understand whether the observed effects in the data are driven by any systematic time-invariant differences between brokerage houses covering particular companies and the companies themselves. These regressions ought to provide similar answers as the DID approach except that we can control for additional sources of heterogeneity.

We estimate our regression model using a pooled (panel) regression and calculating standard errors by clustering at the

merger level. This approach addresses the concern that the errors, conditional on the independent variables, are correlated within merger groupings (e.g., Moulton [1986]). One reason this may occur is that the bias occurring in one company may also naturally arise in another company covered by the same house because the broker tends to cover stocks with similar bias pressures.⁹

The results for the effect on bias obtained using an alternative regression approach outlined in equation (3) are presented in Table VII. The first column shows the result using mean bias and the second column shows the results for median bias. In the first column, the coefficient of interest in front of $\text{MERGE} \times \text{AFFECTED}$ is 0.0021, which is significant at the 10% level. The coefficient of interest increases slightly to 0.0024 for median bias and the statistical significance level is 5%. Hence, the results in this table are consistent with those using the DID estimator, though the estimates are a bit bigger.

Further, we account for the fact that the bias change we capture may result from the difference in the timeliness of the forecasts issued premerger compared to postmerger. In particular, empirical evidence suggests that analyst bias is more pronounced the farther out is the forecast. Indeed, there is a tendency for an analyst to undershoot the earnings number for forecasts issued near the earnings date. To this end, we first document that there is no difference in the timeliness of the forecasts issued premerger as compared to postmerger. Further, in our regression model, we include an additional control variable—recency (Rec)—that measures the average distance of the forecast from the date for which the forecast is obtained. The results, presented in columns (3) and (4) of Table VII, show that controlling for forecast timing does not qualitatively affect our results.

In this paper, we focus on annual earnings forecasts, because these are the key numbers that the market looks to, and every analyst has to submit such a forecast. For completeness, we also look at how long-term growth forecasts and stock recommendations change for the treatment sample in comparison to the control sample around these mergers. One downside is that data in this case are more sparse as analysts do not issue as many timely growth forecasts or recommendations. Moreover, we cannot measure bias

9. We have also considered other dimensions of clustering: clustering by industry, by stock, by time, and by time and industry. All of them produced standard errors that were lower than the ones we report.

TABLE VII
CHANGE IN FORECAST BIAS: REGRESSION EVIDENCE

	Mean BIAS (1)	Median BIAS (2)	Mean BIAS (3)	Median BIAS (4)
MERGE _{<i>i</i>}	0.0005 (0.0008)	0.0005 (0.0008)	0.0005 (0.0008)	0.0006 (0.0008)
AFFECTED _{<i>i</i>}	-0.0019** (0.0006)	-0.0019** (0.0007)	-0.0019*** (0.0006)	-0.0019*** (0.0006)
MERGE _{<i>i</i>} × AFFECTED _{<i>i</i>}	0.0021* (0.0012)	0.0024** (0.0012)	0.0021* (0.0012)	0.0024** (0.0012)
LNSIZE _{<i>i,t-1</i>}	0.0038*** (0.0010)	0.0037*** (0.0010)	0.0038*** (0.0009)	0.0037*** (0.0010)
RETANN _{<i>i,t-1</i>}	0.0005 (0.0017)	0.0004 (0.0017)	0.0005 (0.0017)	0.0004 (0.0017)
LNBM _{<i>i,t-1</i>}	-0.0037 (0.0073)	0.0001 (0.0074)	-0.0037 (0.0073)	0.0001 (0.0074)
COVERAGE _{<i>i,t-1</i>}	0.0001 (0.0001)	0.0001 (0.0001)	0.0001 (0.0001)	0.0001 (0.0001)
SIGMA _{<i>i,t-1</i>}	0.0001** (0.0000)	0.0000 (0.0000)	0.0001** (0.0000)	0.0000 (0.0000)
VOLROE _{<i>i,t-1</i>}	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)
PROFIT _{<i>i,t-1</i>}	0.0629*** (0.0052)	0.0620*** (0.0052)	0.0630*** (0.0052)	0.0621*** (0.0052)
SP500 _{<i>i,t-1</i>}	0.0032 (0.0031)	0.0039 (0.0037)	0.0032 (0.0031)	0.0039 (0.0037)
REC _{<i>i,t-1</i>}			-0.0000 (0.0000)	-0.0000 (0.0000)
Merger fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Brokerage fixed effects	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
Observations	57,005	57,005	57,005	57,005

Notes. The dependent variable is forecast bias (BIAS), defined as the difference between forecasted earnings and actual earnings, adjusted for the past year's stock price. For each merger, we consider a one-year window prior to merger (pre-event window) and a one-year window after the merger (postevent window). We construct an indicator variable (MERGE) equal to one for the postevent period and zero for the pre-event period. For each merger window, we assign an indicator variable (AFFECTED) equal to one for each stock covered by both merging brokerage houses (treatment sample) and zero otherwise. LNSIZE is a natural logarithm of the market cap of the stock; SIGMA_{*it*} is the variance of daily (simple, raw) returns of stock *i* during year *t*; RETANN is annual return on the stock; LNBM is a natural logarithm of the book to market ratio; COVERAGE denotes the number of analysts tracking the stock. To measure the volatility of ROE (VOLROE), we estimate an AR(1) model for each stock's ROE using a ten-year series of the company's valid annual ROEs. ROE_{*it*} is firm *i*'s return on equity in year *t*. ROE is calculated as the ratio of earnings in year *t* over the book value of equity. VOLROE is the variance of the residuals from this regression. PROFIT_{*it*} is the profitability of company *i* at the end of year *t*, defined as operating income over book value of assets. SP500 is an indicator variable equal to one if a stock is included in the S&P500 index. REC_{*it*} is the recency measure of the forecast, measured as an average distance between the analyst forecast and the earnings' report. All regressions include three-digit SIC industry fixed effects, merger fixed effects, brokerage fixed effects, and firm fixed effects. We report results based on both mean and median bias. Standard errors (in parentheses) are clustered at the merger groupings.

***, **, * 1%, 5%, and 10% statistical significance.

TABLE VIII
CHANGE IN ALTERNATIVE FORECAST BIAS MEASURES: DID ESTIMATOR

	Mean BIAS (1)	Median BIAS (2)
Panel A: Long-term growth		
DID estimator (SIZE/BM/RET/NOAN-matched)	0.553*** (0.202)	0.352* (0.212)
Panel B: Analyst recommendations		
DID estimator (SIZE/BM/RET/NOAN-matched)	0.0501 (0.0412)	0.0902* (0.0556)

Notes. We measure analyst forecast bias ($BIAS_{jt}$) using two different measures: the forecast of long-term growth of analyst j at time t (Panel A), and the analyst's j stock recommendation at time t (Panel B). For each analyst, the recommendation variable is ranked from 1 to 5, where 1 is strong sell, 2 is sell, 3 is hold, 4 is buy, and 5 is strong buy. The consensus bias is expressed as a mean or median bias among all analysts covering a particular stock. For all mergers, we split the sample of stocks into those covered by both merging brokerage houses (treatment sample) and those not covered by both houses (control sample). We also divide stocks into premerger period and postmerger period (one-year window for each period). For each period we further construct benchmark portfolios using the control sample based on stocks' size (SIZE), book-to-market ratio (BM), average past year's returns (RET), and analyst coverage (NOAN). Our benchmark assignment involves three portfolios in each category. Each stock in the treatment sample is then assigned to its own benchmark portfolio (SIZE/BM/RET/NOAN-matched). Next, for each period, we calculate the cross-sectional average of the differences in analyst forecast bias across all stocks in the treatment sample and their respective benchmarks. Finally, we calculate the difference in differences between postevent period and pre-event period (DID estimator). Our sample excludes observations with stock prices lower than \$5 and those for which the absolute difference between forecast value and the true earnings exceeds \$10. Standard errors (in parentheses) are clustered at the merger groupings.
***, **, * 1%, 5%, and 10% statistical significance.

in the same way, because there are no actual earnings forecasts to make the comparison to. However, we can gauge the extent to which the average long-term forecast or recommendation changes across the merger date for our treatment sample (provided data are available) compared to the control group. To the extent that there is less competition as a result of these mergers, we expect forecasts for percentage growth to be higher after the merger and for them to have more positive recommendations.

The results for the long-term growth forecasts and recommendations are in Table VIII. Panel A reports the results for long-term growth forecasts (which is the percentage long-term growth in earnings). Using the most conservative benchmark, we see that long-term growth forecasts increase by 55 bp's after the merger, using mean forecasts, and by 35 bp's, using median forecasts. The mean long-term growth forecast in the treatment sample is 14%, with a standard deviation of 6%. So a one-analyst drop in coverage in our treatment sample results in an increase in the mean long-term growth forecast that is about 9% of a standard deviation of these forecasts. This is both an economically and a statistically significant effect.

Panel B reports the results using recommendations. Recommendations are given in terms of the following five rankings: strong sell, sell, hold, buy, and strong buy. We convert these into a score of 1 for strong sell, 2 for sell, 3 for hold, 4 for buy, and 5 for strong buy. We then take the mean and median of these recommendation scores and look at how they vary for the treatment sample and the control group across the merger date. Using again the most conservative benchmark, the merger event is associated with an increase in the average recommendation score for the treatment sample of 0.05 using the mean score and 0.09 using the median score. The result using the mean score is not statistically significant, but the result using the median score is statistically significant at the 10% level. However, both estimates imply quite significant economic effects. The mean score for the treatment sample is 3.87 with a standard deviation of 0.44. Hence, we find that a one-analyst drop in coverage leads to about a 20% (10%) increase in the median (mean) recommendation score as a fraction of the standard deviation of these recommendations.

In sum, we conclude that our baseline results based on annual forecasts are robust to different measures of bias. Moreover, in economic magnitude, they are half as large as the alternative effects we document above, and thus they constitute a lower bound in estimating the effect of competition on bias. One explanation of this fact is that long-term and recommendation forecasts might be more difficult to verify and thus they are subject to stronger competition effects.

VI.D. Additional Analysis

In this section, we bring to bear additional evidence on one of the mechanisms behind the competition effect in the data.¹⁰ Consider the independence-rationale mechanism in Gentzkow and Shapiro in which one independent or honest analyst can discipline her peers and subvert attempts to suppress information. Now imagine a firm covered by many analysts. This firm is less likely to try to bribe analysts to suppress information because the likelihood of drawing an independent or honest analyst is so high. In contrast, a firm with only a couple of analysts covering it is more likely to attempt to influence their few analysts because the payoffs for doing so are higher. One can measure attempts

10. We would like to thank an anonymous referee for several suggestions on addressing this issue.

on the firm to influence analysts by comparing the incentives for analysts for optimism bias for analysts who cover stocks with low analyst coverage versus those for analysts who cover stocks with high analyst coverage—by our logic, we expect higher incentives for analyst optimism bias for analysts covering stocks with little initial coverage or competition to begin with.

Hence, our focus in this section is to measure how the incentives for optimism bias faced by analysts change depending on how much competition there is. Here we build on the work of Hong and Kubik (2003), who measure the implicit incentives of analysts for bias by looking at how career outcomes depend on optimism bias. They document strong evidence that optimism is correlated with subsequent career outcomes: Optimism increases chances of being promoted, whereas pessimism increases chances of being demoted. Our twist, building on their framework, is to examine whether, in the cross section, the subsequent career outcomes of analysts that cover stocks followed by fewer analysts are more strongly related to the degree of their bias. We interpret this as firms with little analyst coverage attempting to influence analysts by providing them incentives through their brokerage houses.

To this end, using analysts as a unit of observation, we estimate the following linear probability models:

$$(4a) \quad \text{Promotion}_{it+1} = \alpha + \beta_1 \text{Bias}_{it} + \beta_2 \text{Coverage}_{it} + \beta_3 \text{Bias}_{it} \\ \times \text{Coverage}_{it} + \beta_4 \text{Controls}_{it} + \varepsilon_{it+1},$$

$$(4b) \quad \text{Demotion}_{it+1} = \alpha + \beta_1 \text{Bias}_{it} + \beta_2 \text{Coverage}_{it} + \beta_3 \text{Bias}_{it} \\ \times \text{Coverage}_{it} + \beta_4 \text{Controls}_{it} + \varepsilon_{it+1}.$$

Our coefficient of interest is β_3 . Following Hong and Kubik (2003), Promotion_{it+1} equals one if an analyst i moves to a brokerage house with more analysts, and zero otherwise; Demotion_{it+1} equals one if an analyst i moves to a brokerage house with fewer analysts, and zero otherwise; Controls is a vector of controls including forecast accuracy, natural logarithm of an analyst's experience, the size of the brokerage house. We also include year fixed effects, broker fixed effects, and analyst fixed effects. We estimate our regression model using a pooled (panel) regression and calculating standard errors by clustering at the analyst level.

An important control in the regression model is forecast accuracy. To construct our measure of accuracy, we first calculate an analyst forecast error, defined as the absolute difference between his or her forecast and the actual EPS of firm i at time t .

We express the difference as a percentage of the previous year's stock price. Subsequently, we follow the methodology in Hong and Kubik (2003) and construct a measure of relative analyst forecast accuracy. To this end, we first sort the analysts who cover a particular stock in a year based on their forecast error. We then assign a ranking based on this sorting; the best analyst (the one with the lowest forecast error) receives the first rank for that stock, the second best analyst receives the second rank, and onward until the worst analyst receives the highest rank. If more than one analyst is equally accurate, we assign all those analysts the midpoint value of the ranks they take up. Finally, we scale an analyst's rank for a firm by the number of analysts who cover that firm.

Following Hong and Kubik (2003), who have shown that analyst accuracy predicts career outcomes in a nonlinear fashion, we define two measures of accuracy: *High Accuracy* is an indicator variable that equals one if an analyst's error falls into the highest decile of the distribution of accuracy measure, and zero otherwise; *Low Accuracy* is an indicator variable that equals one if the analyst's error falls into the lowest decile of the distribution of accuracy measure, and zero otherwise. Note that in the regression with promotions as the dependent variable, the indicator for high accuracy is used as a control, whereas the indicator for low accuracy is used as a control when demotions is the dependent variable. We can have a more symmetric specification in which we include set of dummies for different accuracy deciles, along with the respective interaction terms with coverage, as control variables for both promotions and demotions and the results would be identical (these are available from the authors).

We present results from estimation of equation (4) in Table IX. Columns (1) and (2) show the results for promotion and columns (3) and (4) for demotion. In columns (1) and (3), we replicate the results in Hong and Kubik (2003) and show that bias positively affects the probability of being promoted and negatively affects the probability of being demoted. More important, consistent with the hypothesis that competition affects incentives for bias, we find that the probability of being promoted is positively correlated with bias in an environment with lower competition. The result is statistically significant at the 1% level. Likewise, we find a qualitatively similar result for the probability of being demoted at the 10% level of significance. These results provide support for the independence rationale mechanism behind the disciplining effect of competition on bias.

TABLE IX
INCENTIVES AND CAREER OUTCOMES

	Promotion		Demotion	
	(1)	(2)	(3)	(4)
BIAS	0.0106* (0.0067)	0.0448*** (0.0136)	-0.0036 (0.0074)	-0.0065 (0.0092)
HIGH ACCURACY	0.0082** (0.0040)	0.0183** (0.0080)		
LOW ACCURACY			0.0262*** (0.0062)	0.0233** (0.0098)
BIAS × COVERAGE		-0.0019*** (0.0006)		0.0005* (0.0003)
HIGH ACCURACY × COVERAGE		-0.0006 (0.0004)		
LOW ACCURACY × COVERAGE				0.0001 (0.0005)
COVERAGE	-0.0004* (0.0002)	0.0002 (0.0003)	0.0004 (0.0003)	0.0002 (0.0003)
EXPERIENCE	0.0262*** (0.0052)	0.0265*** (0.0052)	0.0067 (0.0043)	0.0065 (0.0043)
BROKERAGE SIZE	-0.0016*** (0.0003)	-0.0016*** (0.0003)	0.0008*** (0.0002)	0.0008*** (0.0002)
Year fixed effects	Yes	Yes	Yes	Yes
Brokerage fixed effects	Yes	Yes	Yes	Yes
Analyst fixed effects	Yes	Yes	Yes	Yes
Observations	45,770	45,770	45,770	45,770

Notes. The dependent variables are promotion (PROMOTION), defined as an indicator variable equal to one when an analyst moves to a larger brokerage house, and zero otherwise; and demotion (DEMOTION), defined as an indicator variable equal to one when an analyst moves to a smaller brokerage house, and zero otherwise. BIAS is forecast bias, defined as the difference between forecasted earnings and actual earnings, adjusted for the past year's stock price. COVERAGE denotes the number of analysts tracking the stock. HIGH ACCURACY is an indicator variable that equals one if an analyst's error falls into the highest decile of the distribution of accuracy measure, and zero otherwise; LOW ACCURACY is an indicator variable that equals one if the analyst's error falls into the highest decile of the distribution of accuracy measure, and zero otherwise. EXPERIENCE is the natural logarithm of the number of years that the analyst is employed in the brokerage house. BROKERAGE SIZE is the number of analysts employed by the brokerage house. This table includes an interaction term between BIAS and COVERAGE. All regressions include year fixed effects, brokerage fixed effects, and analyst fixed effects. Standard errors (in parentheses) are clustered at the analyst groupings.
***, **, * 1%, 5%, and 10% statistical significance.

Another key part of the independence-rationale mechanism through which we think competition affects bias is that independent analysts keep other analysts honest. To this end, one can assess the incentives for being biased by comparing the penalty that competition imposes for an analyst who is contradicted by other analysts. A reasonable hypothesis is that the impact of bias on the probability of being promoted for such an analyst is lower the more contradicted she is. To evaluate this hypothesis, one can

define contradiction by other analysts as the average across stocks of an absolute difference between an analyst's bias and the average bias of all other analysts calculated for each stock. We find some mixed evidence in support of this hypothesis. We find that the importance of bias for an analyst's promotion decreases with the degree of her being contradicted. The result is significant at the 5% level. Unfortunately, being contradicted by other analysts does not significantly change the probability of being demoted. We omit these results for brevity. Nonetheless, this set of findings, along with the ones on how incentives for analyst optimism bias vary with competition, provide some comforting support for the independence-rationale mechanism.

Finally, even though it might seem ideal to implement the above ideas using the context of our merger experiment, we note that this task proves to be quite difficult empirically. The main problem is that of statistical power. Because we do not observe many career changes for each analyst around merger events, it is very difficult to estimate the incentives for bias separately before and after the merger. Hence, we decided to use the approach that is most appealing statistically.

VII. CONCLUSIONS

We attempt to measure the effect of competition on bias in the context of analyst earnings forecasts, which are known to be excessively optimistic due to conflicts of interest. Using cross-sectional regressions, we find that stocks with more analyst coverage, and presumably competition, have less biased forecasts on the average. However, these OLS estimates are biased, because analyst coverage is endogenous. We propose a natural experiment for competition—namely, mergers of brokerage houses, which result in the firing of analysts because of redundancy and other reasons including culture clash and general merger turmoil. We use this decrease in analyst coverage for stocks covered by both merging houses before the merger (the treatment sample) to measure the causal effect of competition on bias. We find that the treatment sample simultaneously experiences a decrease in analyst coverage and an increase in optimism bias the year after the merger relative to a control group of stocks. Our findings suggest that competition reduces analyst optimism bias. Moreover, the economic effect from our estimates is much larger than that from the OLS estimates.

Our findings have important welfare implications. Notably, a number of studies find that retail investors in contrast to institutional investors cannot adjust for the optimism bias (i.e., debias) of analysts and hence these optimistic recommendations have an effect on stock prices (e.g., Michaely and Womack [1999], Mal-mendier and Shanthikumar [2007]). One conclusion of our findings is that more competition can help protect retail investors because it tends to lower the optimism bias of analysts.

Finally, our natural experiment for analyst coverage can also be useful for thinking about the determinants of stock prices. There is a large literature in finance and accounting that has tried to pin down whether analyst coverage increases stock prices. These studies are typically biased because of endogeneity, as analysts tend to cover high-priced, high-performing, or large stocks, for a variety of reasons. In other words, the causality might be reversed. Our natural experiment can hence be used to identify the causal effect of coverage on stock prices. Recent interesting research in the spirit of our experiment is that of Kelly and Ljungqvist (2007), who use closures of brokerage houses as a source of exogenous variation in coverage. We anticipate more exciting work will be done along this vein.

APPENDIX: MERGERS INCLUDED IN THE SAMPLE (SORTED BY DATE)

Merger number	Merger date	Target house	Target's industry	IBES no.	Industry code	Bidder house	IBES no.	Bidder's industry	Industry code
1	9/10/1984	Becker Paribas	Brokerage firm	299	6211	Merrill Lynch & Co., Inc.	183	Pvd investment, financial advisory services	6211
8	10/31/1988	Butcher & Co., Inc.	Securities dealer; RE broker	44	6211	Wheat First Inc (WFI)	282	Investment bank, brokerage firm	6211
2	12/31/1994	Kidder Peabody & Co.	Investment bank	150	6211	PaineWebber Group, Inc.	189	Investment bank	6211
3	5/31/1997	Dean Witter Discover & Co.	Pvd sec brokerage services	232	6211	Morgan Stanley Group, Inc.	192	Investment bank	6211
4	11/28/1997	Salomon Brothers	Investment bank	242	6211	Smith Barney	254	Investment bank	6211
9	1/9/1998	Principal Financial Securities	Investment bank; securities firm	495	6211	EVEREN Capital Corp.	829	Securities brokerage firm	6211
10	2/17/1998	Jensen Securities Co.	Securities brokerage firm	932	6211	DA Davidson & Co.	79	Investment company	6799
11	4/6/1998	Wessels Arnold & Henderson LLC	Investment bank	280	6211	Dain Rauscher Corp.	76	Investment bank	6211
12	10/1/1999	EVEREN Capital Corp.	Securities brokerage firm	829	6211	First Union Corp., Charlotte, NC	282	Commercial bank; holding company	6021
13	6/12/2000	JC Bradford & Co.	Securities brokerage firm	34	6211	PaineWebber Group, Inc.	189	Investment bank	6211
5	10/15/2000	Donaldson Lufkin & Jenrette	Investment bank	86	6211	CSFB	100	Investment bank	6211
6	12/10/2000	Paine Webber	Investment bank	189	6211	UBS Warburg Dillon Read	85	Investment bank	6211
7	12/31/2000	JP Morgan	Investment bank	873	6211	Chase Manhattan	125	Investment bank	6211
14	9/18/2001	Josephthal Lyon & Ross	Security brokers and dealers	933	6211	Fahnestock & Co.	98	Securities brokerage firm	6211
15	3/22/2005	Parker/Hunter Inc	Pvd investment, investment banking services	860	6211	Janney Montgomery Scott LLC	142	Pvd sec brokerage services	6211

PRINCETON UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH
 NEW YORK UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH

REFERENCES

- Abarbanell, Jeffery S., "Do Analysts' Earnings Forecasts Incorporate Information in Prior Stock Price Changes?" *Journal of Accounting and Economics*, 14 (1991), 147–165.
- Besley, Timothy, and Andrea Prat, "Handcuffs for the Grabbing Hand? The Role of the Media in Political Accountability," *American Economic Review*, 96 (2006), 720–736.
- Brown, Phillip, George Foster, and Eric Noreen, *Security Analyst Multi-year Earnings Forecasts and the Capital Market* (Sarasota, FL: American Accounting Association, 1985).
- Chopra, Vijay K., "Why So Much Error in Analysts' Earnings Forecasts?" *Financial Analysts Journal*, 54 (1998), 30–37.
- Dechow, Patricia, Amy Hutton, and Richard G. Sloan, "The Relation between Affiliated Analysts' Long-Term Earnings Forecasts and Stock Price Performance Following Equity Offerings," *Contemporary Accounting Research*, 17 (1999), 1–32.
- Dreman, David, and Michael Berry, "Analyst Forecasting Errors and Their Implications for Security Analysis," *Financial Analysts Journal*, 51 (1995), 30–42.
- Dugar, Abhijeet, and Siva Nathan, "The Effect of Investment Banking Relationships on Financial Analysts' Earnings Forecasts and Investment Recommendations," *Contemporary Accounting Research*, 12 (1995), 131–160.
- Fang, Lily, and Ayako Yasuda, "The Effectiveness of Reputation as a Disciplinary Mechanism in Sell-Side Research," *Review of Financial Studies*, 22 (2009), 3735–3777.
- Gentzkow, Matthew, Edward L. Glaeser, and Claudia Goldin, "The Rise of the Fourth Estate: How Newspapers Became Informative and Why It Mattered," in *Corruption and Reform: Lessons from America's History*, Edward L. Glaeser and Claudia Goldin, eds. (Cambridge, MA: National Bureau of Economic Research, 2006).
- Gentzkow, Matthew, and Jesse M. Shapiro, "Media Bias and Reputation," *Journal of Political Economy* 114 (2006), 280–316.
- , "Competition and Truth in the Market for News," *Journal of Economic Perspectives*, 22 (2008), 133–154.
- Hong, Harrison, and Jeffrey D. Kubik, "Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts," *Journal of Finance*, 58 (2003), 313–351.
- Kelly, Bryan, and Alexander Ljungqvist, "The Value of Research," Stern NYU Working Paper, 2007.
- Laster, David, Paul Bennett, and In Sun Geoum, "Rational Bias in Macroeconomic Forecasts," *Quarterly Journal of Economics*, 114 (1999), 293–318.
- Lim, Terence, "Rationality and Analysts' Forecast Bias," *Journal of Finance*, 56 (2001), 369–385.
- Lin, Hsiou-wei, and Maureen F. McNichols, "Underwriting Relationships, Analysts' Earnings Forecasts and Investment Recommendations," *Journal of Accounting and Economics*, 25 (1998), 101–127.
- Malmendier, Ulrike, and Devin Shanthikumar, "Are Small Investors Naïve about Incentives?" *Journal of Financial Economics*, 85 (2007), 457–489.
- McNichols, Maureen, and Patricia C. O'Brien, "Self-Selection and Analyst Coverage," *Journal of Accounting Research*, 35 (1997), Supplement, 167–199.
- Michaely, Roni, and Kent L. Womack, "Conflict of Interest and the Credibility of Underwriter Analyst Recommendations," *Review of Financial Studies*, 12 (1999), 653–686.
- Moulton, Brent, "Random Group Effects and the Precision of Regression Estimates," *Journal of Econometrics*, 32 (1986), 385–397.
- Mullainathan, Sendhil, and Andrei Shleifer, "The Market for News," *American Economic Review*, 95 (2005), 1031–1053.

- Ottaviani, Marco, and Peter N. Sørensen, "Forecasting and Rank-Order Contests," Kellogg Working Paper, 2005.
- Sheather, Simon J., and Chris M. Jones, "A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation," *Journal of the Royal Statistical Association B*, 53 (1991), 683–690.
- Stickel, Scott E., "Predicting Individual Analyst Earnings Forecasts," *Journal of Accounting Research*, 28 (1990), 409–417.
- Wu, Joanna S., and Amy Zang, "What Determines Financial Analysts' Career Outcomes during Mergers?" *Journal of Accounting and Economics*, 47 (2009), 59–86.