The B.E. Journal of Economic Analysis & Policy

Advances

Volume 7, *Issue* 1 2007 *Article* 60

Detecting Information Pooling: Evidence from Earnings Forecasts after Brokerage Mergers

Serena Ng* Matt Shum†

*University of Michigan, serena.ng@umich.edu
†Johns Hopkins University, mshum@jhu.edu

Recommended Citation
Serena Ng and Matt Shum (2007) "Detecting Information Pooling: Evidence from Earnings Forecasts after Brokerage Mergers," *The B.E. Journal of Economic Analysis & Policy*: Vol. 7: Iss. 1 (Advances), Article 60.
Available at: http://www.bepress.com/bejeap/vol7/iss1/art60

Detecting Information Pooling: Evidence from Earnings Forecasts after Brokerage Mergers*

Serena Ng and Matt Shum

Abstract

Forecast improvements can be expected if the two partners involved in a brokerage merger pool information and expertise. We examine four large mergers of brokerage firms in the last decade to study the incidence of and explanations for forecast improvements after the mergers. At the brokerage-level, we find that for two of the four mergers, forecast improvements appear more pronounced in subsamples of stocks for which both of the pre-merger analysts were retained in the merged brokerage. At the analyst-level, we find only weak evidence of forecast improvements after the merger. However, we find evidence that after a merger, a stock is more likely to be assigned to an analyst with overall better forecasting performance before the merger. This suggests that analyst selection can be a mechanism generating the post-merger forecasting improvements.

KEYWORDS: information pooling, earnings forecasts, brokerage mergers

---

*The authors may be contacted via e-mail at serena.ng@columbia.edu (Ng) and mshum@jhu.edu (Shum). We thank Liran Einav, Andreas Lehnert, Pai-ling Yin, Nicholai Iskrev, and participants at the Federal Reserve Board, the 2005 Winter Econometric Society Meetings, and the Tuck IO conference for suggestions. The editor and two anonymous referees provided extremely useful suggestions. Yonatan Ben Shalom, Eric Millis and Migiwa Tanaka provided outstanding research assistance. Two anonymous referees and the editor provided suggestions which greatly improved the paper. We gratefully acknowledge the contribution of Thomson Financial for providing earnings per share forecast data, available through the Institutional Brokers Estimate System.

1 Introduction

In the last two decades, there have been a very large number of mergers in the financial sector. This merger wave can perhaps be attributed to two large shifts in government policy during this period. First, starting in the mid-1980s, restrictions on interstate banking were loosened, triggering a consolidation wave among former state-level commercial banks. Then in 1999, the Gramm-Leach-Bliley Act took effect and effectively repealed the Glass-Steagall Act, which for decades restricted commercial banks from offering investment banking services, and vice versa.

Not only have the causes of mergers in the financial sector been varied, so have the effects of these mergers. Mergers are large complicated events, which affect the merging firms in many aspects, including stock market valuations, hierarchical and organizational structure, and leading to substantial employee turnover. However, empirical work on the effects of these mergers have, for the most part, only focused on the effects of mergers on stock market valuations of the merged entities, or on the pricing of financial products offered by the merging firms (such as loan or deposit interest rates quoted by commercial banks).

A well-accepted view of financial companies — whether commercial or investment banks, or insurance companies – is that they provide intermediation services in markets in which there are information asymmetries between the suppliers and demanders of credit (eg. Diamond (1984)). Given the importance of information in the activities of financial companies, it is surprising that little is known about the informational effects of mergers among financial companies at the empirical level.

In this paper, we seek to understand the informational effects of a merger between financial firms by focusing on a specific service offered by financial companies: earnings forecasting. We focus on one particular type of informational effect of a merger, namely, the *pooling* of information and informational resources which, prior to the merger, were privately owned by each of the merging firms. The forecasting enterprise is fundamentally information based. The accuracy of a particular forecast depends on the resources that a brokerage firm devotes to collecting information, and on the assigned analyst's ability to synthesize that information. What makes a forecast especially good or bad often depends on the access to and interpretation of an analyst's private information.

These features of the earnings forecast enterprise make it well-suited for investigating information pooling. Prior to the merger, each of the merging brokerages would have assigned an analyst to cover a given stock, eg. Apple Inc. After the merger, both of these analysts could potentially be retained in the merged brokerage. By comparing the post- vs. pre-merger changes in forecast accuracy across stocks for which both of the pre-merger analysts were retained, versus those for which only one of the analysts was retained, we can measure the importance of information pooling, which should only be present for those stocks where both of the pre-merger analysts were retained. We also distinguish the effects of information pooling from analyst selection, which is the possibility that better-abilitied analysts are more likely to be retained following the merger.

Our empirical analysis is based on the comparison of the accuracy of earnings forecasts before and after four large mergers of brokerage firms in the IBES database. The IBES dataset contains detailed analyst-level information for each forecast and allows us to track performance at the stock- and analyst-level, both before and after the merger, which is ideal for addressing the presence of information pooling via the exercise described above.

At the brokerage-level, we find some evidence consistent with information pooling for two of the four mergers. For these two mergers, we find that forecast improvements appear more pronounced in subsamples of stocks where information pooling should be strongest. These subsamples include the stocks which were covered by both of the merging brokerages before the merger, as well as the stocks where both of the pre-merger analysts were retained in the merged brokerage. This evidence persists even after controlling for changes in the timing of forecast released after the mergers. These two mergers were also the ones where the merging firms were most equal in forecasting ability before the mergers, which perhaps made information pooling more likely.

At the analyst-level, our evidence is more mixed. For one of the four mergers, we find that while the post-merger forecasts of analysts employed before the merger at the acquired (target) brokerage benefit from the presence of the analyst who covered the same stock at the acquiring (bidder) brokerage, the bidder analysts do not benefit as much from having the target analysts around. For the other three mergers, however, we have no robust evidence of information pooling.

Finally, we also consider whether the post-merger forecast improvements can be attributed to analyst selection. We find no evidence that better analysts are more likely to be retained in the merged brokerage following the merger. This confirms anecdotal evidence that in the wake of job uncertainty due to the mergers, many of the best analysts at the merging firms were poached away by competing brokerages, so that the analysts remaining at the merged brokerage following the merger are not the best analysts working at the two brokerages before the merger. However, in the cases where both of a stock's pre-merger analysts were retained in the merged brokerage, we find strong evidence (for three of the four mergers), that the stock is likely to be assigned after the merger to the analyst with the better overall pre-merger forecasting performance. This suggests that analyst selection can be a mechanism generating the post-merger forecasting improvements.

Two bodies of empirical work are related indirectly to this paper. First, a number of papers have focused on the informational aspects of the earnings forecasting enterprise. Hong, Kubik, and Solomon (2000) focus on analysts' career concerns and the tendency towards herding, and Kandel and Pearson (1995) look for evidence that competing analysts have differential interpretations of public information regarding a stock. Bernhardt and Kutsoati (1999) look for evidence in the data consistent with a model in which analysts are compensated for the relative (rather than absolute) accuracy of their forecasts. However, none of these papers have focused on information pooling, which is the topic of this paper.

Second, there is a literature on measuring the effects on mergers among financial institutions, especially commercial banks (including Prager and Hannan (1993), Sapienza (2002), Focarelli and Panetta (2003), Panetta, Schivardi, and Shum (2006), Calomiris and Pornrojnangkool (2005)). All of these studies examine the effect of mergers on the interest rates subsequently offered and charged by the merged banks. Panetta, Schivardi, and Shum (2006) examined whether merging banks pooled information regarding borrowers which they had in common (before the merger) and found no evidence for information pooling. To our knowledge, however, this is the first paper on mergers among financial institutions which explores the micro-effects of the mergers – in our case, the employment turnover caused by the mergers – on earnings forecasts.

Finally, while there is a large theoretical literature in industrial organization on information pooling (cf. Vives (1999), ch. 8), the empirical work on information pooling has been limited. There are several studies of the effects of information sharing within trade organizations. Genesove and Mullin (1999) presents a case study of the sugar producers in the early part of the twentieth century, and Snyder and Doyle (1999) examined how automobile manufacturers' announcements of their future production plans in a prominent trade journal affected the manufacturers' actual production levels.

In the next section, we present a simple framework for looking at the effects of information pooling. In Section 3, we introduce the data, and discuss the four mergers which we focus on in this paper. Sections 4 and 5 contain the empirical results, and Section 6 concludes.

2 Information pooling, analyst selection, and forecast improvements

In this section, we explore the relationship between forecast improvements, and how they can be explained by information pooling and analyst selection. We also present several definitions of forecast improvements which we will use in our empirical work. Hereafter, we will use the word "firm" to denote a brokerage firm, which forms earnings forecasts, while reserving the word "stock" to denote publicly-traded companies about which forecasts are being made.

Consider two competing brokerage firms, $j = 1, 2$, who have assigned, respectively, analysts A_1 and A_2 to forecast a variable v_i, which is the quarterly earnings per share of stock i. Prior to the merger, brokerage 1's forecast of v_i is $x_{i,1}$, and brokerage 2's forecast of v_i is $x_{i,2}$.[1] Let $z_{i,j}$, $j = 1, 2$ denote the private information of brokerage j, which is used to form brokerage j's forecast $x_{i,j}$ using the forecasting function $h_j(z_{i,j}), j = 1, 2$. The forecast is related to the true value by

$$\begin{aligned} h_1(z_{i,1}) &= x_{i,1} = v_i + e_{i,1} \\ h_2(z_{i,2}) &= x_{i,2} = v_i + e_{i,2} \end{aligned}$$

where $e_{i,1}$ and $e_{i,2}$ are forecast errors. For simplicity, assume that $x_{i,1}$ and $x_{i,2}$ are unbiased for v_i, so the pre-merger mean squared forecast error is

$$\text{MSE}_{i,j}^{pre} = \sigma_{i,j}^2 = var(e_{i,j}),\ j = 1, 2.$$

Now suppose a merger occurs between firms 1 and 2. Consider the case when both analysts A_1 and A_2 continue to work at the merged brokerage. If $z_{i,1}$ and $z_{i,2}$ are still available, the post merger forecast is $x_i^{post} = h(z_{i,1}, z_{i,2})$ for some function $h(\cdot)$, which is not necessarily the same as h_1 or h_2. Let MSE_i^{post} denote the post-merger forecast error. We look for two types of forecasting improvements.

1. Brokerage-level improvements To define brokerage-level improvements, we need to compare the merged brokerage post-merger performance to a benchmark for the two brokerages' individuals performances before the merger. Given $x_{i,1}^{pre}$ and $x_{i,2}^{pre}$, we use a benchmark equal to $w_1\text{MSE}(x_{i,1}^{pre}) + w_2\text{MSE}(x_{i,2}^{pre})$, a weighted average of the two brokerages' individual pre-merger forecast accuracies, for weights $0 \leq w_1 \leq 1$ and $w_2 = 1 - w_1$. Then brokerage-level improvements for stock i are defined as the event that

$$\Delta\text{MSE}_i = \text{MSE}_i^{post} - \left[w_1\text{MSE}_{i,1}^{pre} + w_2\text{MSE}_{i,2}^{pre} \right] \leq 0. \tag{1}$$

In the empirical work, we will, for the most part, weigh the two brokerages' pre-merger forecast errors equally (ie. $w_1 = w_2 = \frac{1}{2}$).

One premise of information pooling is that the merging brokerages share information and expertise regarding stock i. Therefore, forecast improvements for stocks which were not covered by both brokerages before the merger is not evidence of information pooling. Furthermore, to the extent that information and expertise is

[1] While the stock subscript i is not required for the discussion in this section, we include it to facilitate comparison of the equations in this section with those in subsequent section, in which the stock subscript will be important.

analyst-specific, information pooling should imply that the improvements are more prominent for stocks where both of the pre-merger analysts continued to work at the merged brokerage following the merger. These considerations will guide our empirical work below.

2. Analyst-level improvements While it is possible for both analysts to continue forecasting stock i after the merger, it also seems reasonable for the merged brokerage to consolidate resources and release one forecast instead. Provided that at least one of the pre-merger analysts covering stock i in the two merging brokerages continues to forecast stock i after the merger, we can also compare the forecast performance of this analyst on stock i before and after the merger. For a stock which is forecast by analyst j after the merger, we say that analyst j's forecasting accuracy of stock i improved relative to his pre-merger performance if

$$\Delta\text{MSE}_{i,j} \equiv \text{MSE}_{i,j}^{post} - \text{MSE}_{i,j}^{pre} \leq 0. \tag{2}$$

The main difference between the analyst-level and brokerage-level forecast change measures is that we can only compute the analyst-level change $\Delta\text{MSE}_{i,j}$ if analyst j covers the stocks both before and after the merger, but the brokerage-level change ΔMSE_i can be computed even if the analyst who covers stock i after the merger did not cover it before the merger.

Again, a precondition for information pooling is that the analyst chosen to produce the forecast after the merger has access to the skills, information, and expertise of the analyst who covered stock i for the other brokerage prior to the merger. For this reason, information pooling should imply that analyst-level forecasting improvements are more likely for those stocks where both of the pre-merger analysts are retained in the merged brokerage. Of course, it will be rare for both pre-merger analysts who covered stock i before the merger to be assigned to cover the same stock after the merger. It is more likely that one analyst will be assigned to other stocks. However, the mere presence of both analysts who had experience with stock i in the merged brokerage means that the post-merger analyst has access to knowledge and information about stock i not available before the merger.

To understand the way that information pooling leads to forecast improvements, suppose for simplicity that the merger leads to a linear aggregation of information. The pooled post-merger forecast is

$$x_i^{post} = \psi_1 h_1(z_{i1}) + \psi_2 h_2(z_{i2}) = \psi_1 x_{i1}^{pre} + \psi_2 x_{i2}^{pre}, \quad 0 \leq \psi_1 \leq 1 = 1 - \psi_2. \tag{3}$$

The MSE (which is also the variance) of the post-merger forecast is

$$\text{MSE}_i^{post} = \psi_1^2\sigma_1^2 + \psi_2^2\sigma_2^2 + 2\psi_1\psi_2\rho\sigma_1\sigma_2 \tag{4}$$

where $\rho = Corr(e_1, e_2)$. For fixed values of ψ_1, and σ_1, MSE_i^{post} is increasing in ρ and σ_2. Intuitively, under linear forecast aggregation, the primary benefit of information pooling is variance reduction. Forecast improvements are thus larger when σ_2 is smaller. Also, averaging the forecasts will lead to a larger reduction in variance when the errors of the individual forecasts are negatively correlated ($\rho < 0$) because the errors offset. This implies that the forecast improvements, at both the brokerage and analyst-level, will depend on the relative ability of the analysts involved and are more likely when ρ is small or negative. Indeed, for some configuration of the parameters and depending on the benchmark, information pooling may not even lead to better forecasts.

2.1 Analyst selection

Typically, mergers of financial institutions lead to a great deal of employment turnover, and the mergers studied in this paper are no exception. The possibility therefore arises that forecasting improvements can also be due to analyst selection. Analyst turnover following a merger can lead to two types of analyst selection. First, if the analysts who remained in the merged firm were systematically better than those who left, a brokerage-level improvement might result. Second, after a merger, the better of the bidder and target analyst can be chosen to forecast a given stock.

In the context of the simple model from the previous section, analyst selection amounts to letting ψ_1 in Eq. (3) be a binary indicator chosen using the criterion

$$\psi_1 = \begin{cases} 1 & \text{if } \text{MSE}_{i,1}^{pre} < \text{MSE}_{i,2}^{pre} \\ 0 & \text{otherwise} \end{cases}$$

so that the better analyst in the pre-merger period is chosen to cover each stock in the post-merger period.

Observationally, both information pooling and analyst selection can appear very similar, because both imply that post-merger forecasts should be more accurate than pre-merger forecasts, and the second type of analyst selection implies that having both of the pre-merger analysts around should lead to better post-merger forecasts, because the firm is able to choose the better analyst to forecast the stock after the merger. However, because our dataset contains detailed information on analyst turnover and assignment to stocks, we can directly measure the importance of both types of analyst selection after the mergers, and hence distinguish information pooling from analyst selection.

3 Data

Our dataset of analyst forecasts is derived from the IBES (Institutional Brokers Estimate System) database, which is a comprehensive database containing every forecast and forecast revision formed by analysts for a near-complete sample of brokerage firms and securities. For a given stock (e.g. IBM) and forecast period (a quarter, e.g. 92III), we observe every earnings per share (EPS) forecast and revision which was submitted by analysts working at brokerage firms surveyed in the dataset. The dataset contains forecasts from the beginning of 1983 to the middle of 2002. In this paper, we focus only on quarterly EPS forecasts, since these are the most common forecasts in the database.[2] We also observe the actual realized earnings for each stock in each quarter for which forecasts are available.

For the remainder of this paper, we refer to the acquiring brokerage as the **bidder firm**, and the acquired brokerage as the **target firm**. We use the terms **bidder analyst** to refer to the analyst who covered a given stock at the bidder firm before the merger, and **target analyst** to refer to the analyst who covered this stock at the target firm before the merger. Finally, for a given stock i and analyst j, we use the term **rival analyst** to denote the analyst covering stock i during the pre-merger period at the brokerage other than the one where analyst j works. For example, for a bidder analyst, her rival analyst is the analyst covering the same stock at the target firm before the merger.

In the IBES dataset, we are able to track a particular analyst across different employers. Particularly, for each stock covered by both the bidder and target firms before the merger, we are able to tell whether the particular analyst who covered this stock at the target firm remained employed at the bidder firm after the merger. This will be a crucial component for our tests for the presence of information pooling.

3.1 Four mergers of brokerage firms

In order to identify mergers among brokerage firms in the IBES data, we used the SDC Mergers and Acquisition database to obtain information on all mergers within SIC four-digit sector 6311 ("Investment and Commodity Firms, Dealers, and Exchanges"). From this list, we identified four sizable mergers among brokerage firms. Since the number of mergers is small, we will do our analysis on a merger-by-merger basis, and rely on the variation across time, across stocks, and across analysts to identify the information pooling effects. The four mergers are listed in Table 1. Generally, all four of these mergers represented attempts by the bidder

[2] The second-most common forecasts are annual earnings forecasts but, for a given year, they are derived simply as the sum of the quarterly earnings forecasts for the four quarters which make up that year

Table 1: List of Mergers Used in the analysis

Merger	A	B	C	D
Bidder Brokerage	Paine Weber	Morgan Stanley	Credit Suisse First Boston	UBS Warburg Dillon Read
Merger Date	12-94	05-97	11-00	11-00
Earliest EPS Latest EPS	2-10-82 11-27-00	5 19-82 5-16-02	7-13-81 5-16-02	4-18-84 5-16-02
Target Brokerage	Kidder Peabody	Dean Witter Reynolds	Donaldson Lufkin and Jenrette	Paine Webber
Earliest EPS Latest EPS	2-17-82 12-19-94	7-30-81 4-28-97	5-19-82 10-10-00	2-10-82 11-27-00

Notes:

1. Earliest EPS is the date for which we have an earnings forecast from this brokerage firm.
2. Last EPS is the date for which we have an earnings forecast from this brokerage firm.

firms to expand the scope of their retail business by purchasing another brokerage. Merger A, between Paine Webber and Kidder Peabody, was portrayed in the press as a "company in trouble" deal, in which the second-tier brokerage (Paine Webber) bought a top-tier investment bank with a strong research department (KP) at an opportune time. Just prior to the merger, KP was reeling in the aftermath of a trading scandal involving its chief government bond trader, Joseph Jett, and had already laid off 10% of its workforce. Subsequently, KP's owner, General Electric, was looking to sell the company.

Merger B was another diversifying merger, in which high-end investment bank Morgan-Stanley was portrayed as wanting to get in on the more down-market retail brokerage operations of Dean Witter.[3] As an indicator of the differences in operations between the merging parties, we note that in 1996, the year before the merger, Morgan-Stanley was the chief underwriter in 43 IPOs, with a combined offer amount of over $7 billion, while Dean Witter underwrote only 4, with a combined offer amount of just under $1 billion.[4]

Mergers C and D occurred only within a few months of each other, and both were perceived to be attempts by Swiss banks to geographically diversify their lines of business (into the American market). Merger C was a merger between two top-of-the-line investment banks (CSFB and DLJ underwrote, respectively, 57 and 36

[3]For a time in the 1980s, Dean Witter operated service desks in Sears department stores.

[4]These figures, as well as those in the following paragraph, are drawn from `www.ipodata.com`.

Table 2: Analyst employment before and after mergers

Merger	Pre-merger[a]	Post-merger[b]
Merger A:		
Paine Webber	45	34
Kidder Peabody	54	9
New		13
Total	99	56
Merger B:		
Morgan Stanley	77	69
Dean Witter	41	5
New		13
Total	118	102
Merger C:		
CS-FB	130	104
DLJ	86	17
New		39
Total	216	160
Merger D:		
UBS	98	71
Paine Webber	70	40
New		24
Total	168	135

[a]Defined as number of analysts who provided forecasts at brokerage within one year before the merger

[b]Defined as number of analysts who provided forecasts at brokerage within one year after the merger

IPOs in 1999), and concerns were raised about whether CSFB would be able to retain many of DLJ's brokers and analysts. Merger D was characterized similarly as a geographically diversifying merger, with the difference being that both Paine Webber's and UBS's American investment banking operations were smaller than those of, respectively, DLJ and CSFB. A common effect of all four mergers is that they precipitated a large degree of turnover. This turnover will be an important source of variation for detecting information pooling, because an important exercise that we do is to compare changes in forecast accuracy for stocks where both pre-merger analyst were retained, versus stocks where only one (or none) of the pre-merger analysts were retained. Table 2 shows the number of analysts employed by the merging units before and after their respective mergers varied substantially.[5] After all four mergers, the number of analysts grew in all four post-merger brokerages (relative to the pre-merger number of analysts in the bidder firms). In percentage terms, the larger increases in the number of analysts occurred after Merger B, where the number of analysts increased by 32.5% (from 77 to 102 analysts), and Merger D, where the increase was 37.8% (from 98 to 135).

The retainment percentages also depend on whether an analyst worked at the bidder or target firm before the merger. Clearly, a higher percentage of analysts from the bidder firm than target firm were retained. For Merger A, 34 out of 45 Paine Webber analysts were retained, but only 9 out of 54 Kidder-Peabody analysts. This pattern holds across all four mergers. Indeed, only for Merger D were more than half of the analysts from the target firm retained, while more than half of the analysts from the bidder firm were retained in all four mergers. Furthermore, across all the mergers, a substantial percentage of the post-merger analysts were new hires, which make up from 20-25% of the post-merger analyst workforce.

Clearly, these four mergers feature very different brokerages, but we note that the improvement of the research group was not a stated objective for merger in any of these four mergers. Hence, it would not be surprising to find that these mergers had no substantial impact on forecast performance. However, if information pooling is important, then these mergers may have provided opportunities for the merged brokerage to experience incidental forecasting improvements via the sharing of private information and expertise, even when these improvements were not a fundamental reason for the mergers.

3.2 Measuring forecast accuracy

The main empirical exercise in this paper is to examine whether forecast accuracy was improved in the merged brokerage following a merger. We utilize a standardized forecast error,

$$FE_{ijt} = \frac{f_{ijt} - a_{it}}{p_{it}} \tag{5}$$

where f_{ijt} denotes broker j's forecast of the earnings per share (EPS) of stock i, for the period t, and a_{it} the actual realized EPS. For each stock i and quarter t, we only consider analyst j's final forecast, and do not focus on the forecast revision process.

The error $f_{ijt} - a_{it}$ is standardized by dividing by p_{it}, the price per share of stock

[5] Because analyst turnover is common without or without mergers, we look at analysts employed at the brokerages in the year before the merger to isolate the turnover due to the mergers.

i on the first trading day of quarter t.[6] FE, as defined in this way, can be positive or negative, depending on whether or not $f_{ijt} > a_{it}$. Additionally, we also follow Lim (2001) by deleting observations when $|f_{ijt} - a_{it}| > 10$, and also only consider stocks i and quarters t where $p_{it} \geq 1$.[7]

Because FE can be both positive and negative, and it is still an open questions as to the unbiasedness of analyst forecasts, it is not enough to compare averages of FE across different time periods or stocks. Hence, we focus on the mean-squared error (hereafter MSE) of FE.[8]

We focus on how the MSE's changed across different stocks before and after the merger. Accordingly, we calculate the MSE of FE_{ijt} for each stock i, brokerage j, over a range of pre-merger and post-merger quarters. Specifically, define the pre- and post-merger MSE for a given stock i and brokerage j as

$$\begin{aligned} \text{MSE}_{ij}^{pre} &= 10^6 \times \frac{1}{K} \sum_{t=merg-K}^{merg-1} FE_{ijt}^2, \quad j = bidder,\ target \\ \text{MSE}_{i}^{post} &= 10^6 \times \frac{1}{K} \sum_{t=merg+1}^{merg+K} FE_{ijt}^2, \end{aligned} \tag{6}$$

where $merg$ denotes the quarter of the merger. To ensure that we isolate the effects of the mergers, we only consider forecasts within the K quarters before and after the merger. In this paper, we use a value $K = 8$ for our empirical results.[9] Furthermore, because earnings are defined on a per-share basis, the standardized forecast errors are usually very small, so that we scale up by a factor of 10^6 in computing the MSE.

3.3 Summary statistics: all stocks

In Table 3, we present summary statistics of forecast accuracy for the four mergers. For each merger, we report the median and mean, as well as the 10-th and 90-th quantiles, of $\text{MSE}_{i,bidder}^{pre}$, $\text{MSE}_{i,target}^{pre}$, and MSE_{i}^{post}, across all stocks which were forecast at least twice in the two years preceding the merger (for the pre-merger MSE measures), and the stocks which were forecast at least twice following the merger (for the post-merger MSE measure). First, note that the distribution of

[6]In normalizing by p_{it}, we follow many of the empirical studies which utilize the IBES data, including Rajan and Servaes (1997), Keane and Runkle (1998), and Lim (2001).

[7]The results are qualitatively robust to using alternative cutoff thresholds.

[8]Note that in the illustrative model if the previous section, forecasts are always unbiased, in which case the MSE simplifies to the variance of FE.

[9]Some stocks i were not forecast by brokerage j in each of the K quarters before and after the merger. In these cases, we compute the MSE as the average of the squared forecast errors only for those quarters in which the stock was forecast.

Table 3: Pre- and Post-Merger Mean-squared Errors in Merging Brokerages

		(a)	(b)		(c)		
Merger	Statistic	Pre-merger			Post-merger		
		$\text{MSE}^{pre}_{bidder}$	$\text{MSE}^{pre}_{target}$	(a)=(b)?	MSE^{post}	(c)=(a)?	(c)=(b)?
A	median	3.93	2.32	**	4.11	–	**
	mean	1873.8	1431.2		4808.3		
	10%	0.06	0.05		0.06		
	90%	262.47	164.01		235.59		
	#stocks	440	381		504		
B	median	5.03	2.43	***	4.81	–	***
	mean	2259.7	317.1		2128.2		
	10%	0.05	0.07		0.10		
	90%	290.00	405.74		305.32		
	#stocks	852	418		764		
C	median	6.86	6.90	–	3.69	***	***
	mean	9173.1	2717.5		10207.6		
	10%	0.19	0.09		0.12		
	90%	538.54	497.98		309.43		
	#stocks	1238	749		967		
D	median	5.91	6.74	–	3.92	**	***
	mean	651.14	3706.2		462.69		
	10%	0.14	0.01		0.09		
	90%	302.89	640.03		210.98		
	#stocks	948	494		797		

***: reject equality at 1%
**: reject equality at 5%
*: reject equality at 10%

MSE's is highly skewed to the right. Across all the mergers, the mean MSE generally exceeds the 90-th quantile of the MSE distribution, both before and after the merger. For this reason, in this paper, we employ median (quantile) regressions because, for such a skewed distribution, the median is a better measure of the central tendency of the MSE distribution than the mean.

Table 3 shows that Mergers A and B were quite different from Mergers C and D. In Mergers A and B, the target firms appeared to be better than the bidder firms, in terms of median MSE before the merger. For Merger A, the median MSE for target firm Kidder Peabody was 2.32, while for bidder firm Paine Webber it was

3.93. Column 5 of Table 3 shows that these differences in medians were statistically different from zero at a 5% significance level. For these two mergers, however, the post-merger median MSE was virtually the same as the median of the bidder firm's pre-merger MSE, and substantially higher than the target firm's pre-merger MSE. For example, the median MSE after Merger A was 4.11, which is just slightly higher than Paine Webber's pre-merger median of 3.93. Hence, for these two mergers, we have evidence that worse-performing bidder firms acquired better-performing target firms, and that forecasting accuracy actually deteriorated after the merger, relative to the target firms' pre-merger forecasting accuracy.[10]

The numbers for Mergers C and D tell a different story. The two mergers involved partners which, in terms of their pre-merger forecast accuracy, were rough equals. For example, the median pre-merger MSE's for Credit Suisse–First Boston and Donaldson, Lufkin, and Jenrette were, respectively, 6.86 and 6.90 (and statistically not different from each other). However, there were clear improvements in forecasting accuracy, as the post-merger median MSE in both of these mergers was lower, and statistically different (at the 1% level) from the pre-merger median MSE for both the bidder and target firms.

The simple model of forecasting improvements in the previous section assumes that analysts' forecasts are always unbiased, whereas the MSE (our measure of forecast performance) summarizes both the bias and variance of the forecasts. In Table 4, we report the bias and standard deviation of the forecast errors for each merger, and also before and after the merger.

Across all the results, the magnitude of the standard deviation is much larger than that of the bias. For example, for Merger A, the median pre-merger bias of the forecast errors for the bidder brokerage is -0.0526, but the corresponding standard deviation is 1.7728. Hence, even though the theoretical discussion in the previous section assumed a model where analysts' forecasts are unbiased, the results here suggest that this may not be a bad approximation, because in the data the variance component of the MLE far exceeds the bias component.

3.4 Summary statistics: affected stocks

An important subset of stocks which we focus on in this paper are those which were covered by both the bidder and target firms prior to the merger, and continued to be covered by the merged brokerage following the merger. This particular subset of stocks will be referred to as the **affected** stocks in the rest of this paper. Comparisons of affected and non-affected stocks play an important role in our tests for

[10]This is somewhat surprising for Merger B, because the bidder firm in this merger (Morgan Stanley) is widely considered a better research brokerage than the target firm in that merger (Dean Witter).

Table 4: Pre- and Post-Merger Mean-squared Errors in Merging Brokerages
Mean-squared errors broken down into Bias and Standard Deviation Components

		(a)	(b)	(c)	(d)	(e)	(f)
Merger	Statistic	Pre-merger				Post-merger	
		Bias_{bid}	Stdev_{bid}	Bias_{bid}	Stdev_{bid}	Bias_{bid}	Stdev_{bid}
A	median	-0.0526	1.7728	-0.0547	1.2589	-0.0507	1.8335
	mean	3.5671	9.3330	3.0644	7.4749	4.7202	9.9199
	#stocks	439		380		504	
B	median	-0.0163	2.1013	-0.0435	1.3353	-0.0203	2.0598
	mean	2.7800	9.6790	1.4225	4.7287	3.2028	9.6862
	#stocks	852		418		764	
C	median	-0.2681	2.3221	-0.0775	2.3132	-0.3311	1.7190
	mean	-0.8429	14.0994	2.1434	11.7908	-1.1986	12.5897
	#stocks	1238		749		967	
D	median	-0.1873	2.1069	-0.0326	2.4658	-0.2810	1.7178
	mean	0.9947	8.3386	3.7317	12.7572	0.3163	20.7542
	#stocks	948		494		797	

information pooling.

In Table 5, we report the same statistics as in Table 3, but only for the affected stocks. Across all four mergers, the bidder firm tends to produce more accurate forecasts of the affected stocks than the target firm before the merger, even though this difference is statistically significant only for Mergers C and D. For Mergers C and D, the median post-merger MSE is significantly lower than the pre-merger MSE for the target firm. For Merger B, the evidence here indicates a deterioration in forecast accuracy, relative to the pre-merger performance of both bidder and target firms. For Merger A, we find no evidence of changes in forecasting accuracy after the merger. These numbers seem to suggest that brokerage-level forecast improvements are driven by the pre-merger forecast performance of the bidder firm. Specifically, if the bidder and target firms are roughly equal-abilitied before the merger, then forecast improvements obtain; if the bidder firm is worse than the target firm, then there are no forecasting improvements.

4 Empirical results

In this section, we examine whether the changes in forecasting accuracy documented in Tables (3) and (5) can be attributed to information pooling by seeing whether forecast improvements appear more pronounced in subsamples of stocks

Table 5: Mean Squared Errors of Forecasts: Affected Stocks

		(a)	(b)		(c)		
Merger	Statistic	Pre-merger			Post-merger		
		$\text{MSE}^{pre}_{bidder}$	$\text{MSE}^{pre}_{target}$	(a)=(b)?	MSE^{post}	(c)=(a)?	(c)=(b)?
A	median	0.89	1.00	–	1.29	–	–
	mean	45.81	139.50		129.35		
	10%	0.02	0.03		0.02		
	90%	32.46	74.37		53.17		
	#stocks	137					
B	median	0.79	1.30	–	3.25	***	**
	mean	70.78	69.69		138.74		
	10%	0.01	0.02		0.05		
	90%	31.58	65.26		132.34		
	#stocks	197					
C	median	1.51	3.03	***	1.76	–	**
	mean	154.62	338.78		1278.7		
	10%	0.04	0.05		0.08		
	90%	54.58	112.45		92.54		
	#stocks	383					
D	median	1.04	3.59	***	1.78	–	**
	mean	27.10	83.64		172.52		
	10%	0.03	0.07		0.07		
	90%	45.29	166.79		144.71		
	#stocks	224					

***: reject equality at 1%
**: reject equality at 5%
*: reject equality at 10%

where information pooling should be stronger, such as the affected stocks, and the stocks for which both of the pre-merger analysts were retained in the merged brokerage.

4.1 Brokerage-level forecast improvements

We start by documenting the brokerage-level forecast improvements. Define

- AFFECTED$_i = 1$ if stock i was an affected stock and hence covered by both the bidder and target firms prior to the merger.

We also define two more dummy variables to isolate subsamples of the affected stocks where information pooling should be even stronger:

- BOTHSTAY$_i = 1$ if both the analysts who covered stock i at the bidder and target firms before the merger were retained in the merged brokerage.

- BOTHCOVER$_i$=1 if both analysts cover stock i within two years after the merger after the merger.

Note that the subsample of stocks with BOTHCOVER$_i = 1$ is included in the subsample with BOTHSTAY$_i = 1$, which is in turn included in the subsample of affected stocks (AFFECTED$_i = 1$).

Information pooling is fundamentally about the sharing of private information and expertise, so that forecast improvements should be more prominent for the affected stocks, when presumably both brokerages possess some information and expertise. If the information or expertise required in the forecasting enterprise is analyst-specific, then information pooling should be more pronounced when BOTHSTAY$_i = 1$. Information pooling should be even more pronounced when BOTHCOVER$_i = 1$, especially if information were very time-sensitive and changes quickly, so that pooling occurs only when both analysts are still actively covering the stock in the post-merger period.

In this section, brokerage-level forecast improvements are defined as

$$\Delta\text{MSE}_i \equiv \begin{cases} \text{MSE}_i^{post} - \frac{1}{2}[(\text{MSE}_{i,bidder}^{pre} + \text{MSE}_{i,target}^{pre})] & \text{if AFFECTED}_i = 1 \\ \text{MSE}_i^{post} - \text{MSE}_{i,bidder}^{pre} & \text{if only bidder covers stock } i \\ \text{MSE}_i^{post} - \text{MSE}_{i,target}^{pre} & \text{if only target covers stock } i \end{cases}$$

with ΔMSE$_i < 0$ indicating forecast improvements. Notably, ΔMSE$_i$ is defined differently depending on whether stock i is an affected stock. This is because for non-affected stocks, only one of the brokerages – usually the bidder firm – covers the stock before the merger. For the affected stocks, our base case is to compare the post-merger forecast to a equally weighted pre-merger forecast. Robustness to this definition of forecast improvement will be considered below.

Under information pooling, ΔMSE$_i$ should be more negative when AFFECTED$_i =$ 1 than when AFFECTED$_i = 0$, and even more negative when BOTHSTAY$_i = 1$ and BOTHCOVER$_i = 1$. Some insight can be obtained from Figure (1), which presents the empirical cumulative distribution functions of ΔMSE$_i$, for the various subsamples of interest. Across all the mergers, the CDF for AFFECTED$_i = 1$ (in the solid lines) tends to lie above and to the left of the CDF for the AFFECTED$_i = 0$ subsample of stocks (in the dashed lines), especially for values of ΔMSE greater than zero. This suggests that the values of ΔMSE$_i$ are smaller (in a distributional

Figure 1: Empirical Cumulative Distribution Functions for Brokerage-level Forecast Improvements

X-Axis: $\Delta\text{MSE}_i = \text{MSE}_i^{post} - (0.5 * \text{MSE}_{i,bid}^{pre} + 0.5 * \text{MSE}_{i,tar}^{pre})$
Y-Axis: empirical cumulative distribution function

Solid line: stocks for which $\text{AFFECTED}_i = 0$
Dashed line: stocks for which $\text{AFFECTED}_i = 1$
Dotted line: stocks for which $\text{AFFECTED}_i = 1$ and $\text{BOTHSTAY}_i = 1$

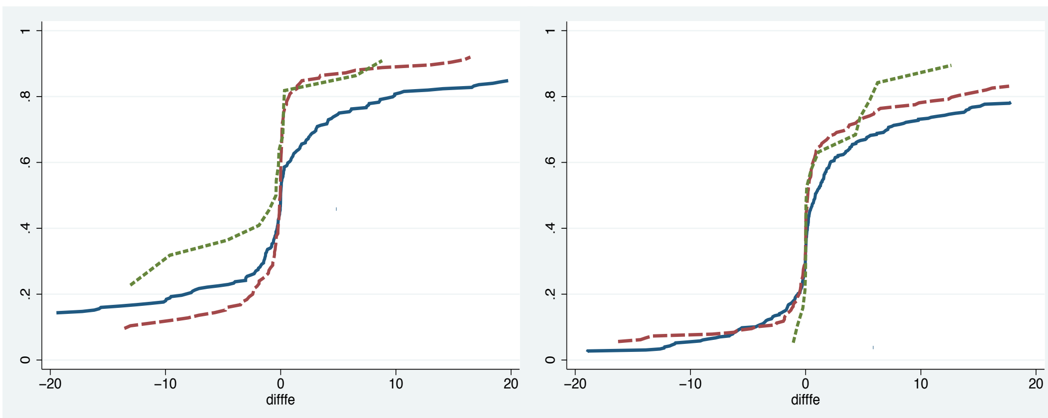

Merger A

Merger B

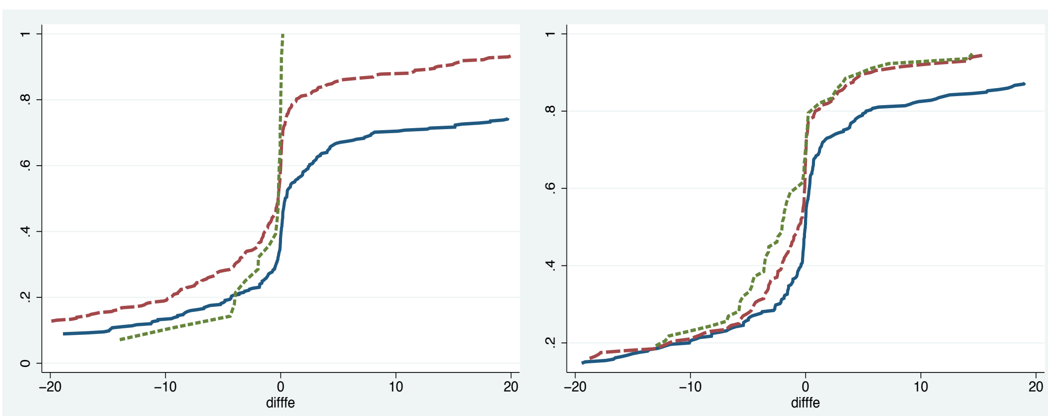

Merger C

Merger D

sense) when both brokerages forecast the stock before the merger, which is consistent with information pooling. Conditional on $\text{AFFECTED}_i = 1$, we then single out those stocks with $\text{BOTHSTAY}_i = 1$. The empirical CDF's for this subsample shows that while the differences are not as sharp as between the $\text{AFFECTED}_i = 0$ and $\text{AFFECTED}_i = 1$, in all four mergers, there are substantial ranges of quantiles where the dotted CDF lies above and to the left of the other two CDFs. This suggests that forecast improvements are larger when $\text{AFFECTED}_i = 1$ and even larger when both AFFECTED$= 1$ and BOTHSTAY$= 1$.

To provide a more formal analysis, consider LAD (least absolute deviation) estimation of the model:[11]

$$\Delta\text{MSE}_i = \alpha + \beta \cdot \text{AFFECTED}_i + \gamma \cdot Z_i + \text{error}_i \tag{7}$$

where a negative β would be consistent with information pooling. The results from this regression are reported in Columns A1, B1, C1, and D1 of Table 6. Since each observation in this regression is a stock, we also include stock-level covariates Z_i to control for additional variation across observations. In the reported specifications, these covariates are AVGMCAP_i and SDEVMCAP_i which measure, respectively, the average and standard deviation of market capitalization of stock i during the eight quarters preceding each of the four mergers studied in this paper. We use AVGMCAP_i to proxy for stock i's size, and SDEVMCAP_i to measure its volatility.[12]

The existing literature on analyst forecasts (eg. Zitzewitz (2001), Gallo, Granger, and Joon (2002)) has stressed the relationship between forecast timing and accuracy. Particularly, later forecasts are usually more accurate because they contain the information revealed in earlier forecasts, so that forecast improvements after the merger could arise simply from the merged firm choosing to release forecasts later, and not from information pooling. To control for this possibility, we create a stock-level variable, DIFFTIMING_i, defined as

$$\begin{aligned} \text{DIFFTIMING}_i =& \text{Avg(Days bef EOQ)}_i^{post} \\ & - [\frac{1}{2}\text{Avg(Days bef EOQ)}_i^{pre,bid} + \frac{1}{2}\text{Avg(Days bef EOQ)}_i^{pre,targ}] \end{aligned} \tag{8}$$

where $\text{Avg(Days bef EOQ)}_i$ is the average number of days before the end-of-quarter for which a forecast for stock i was released. DIFFTIMING_i measures changes in

[11]As we remarked before, we employ LAD regressions here because the results appeared very sensitive to outliers in OLS regressions.

[12]In other specifications (not reported for brevity), we have used shares outstanding, and pre-merger share prices as covariates. The results reported here are robust.

Table 6: Brokerage-level Forecast Improvements

Results from median (quantile) regression; Dependent variable: ΔMSE_i

	Merger A			Merger B			Merger C			Merger D		
Variable	(A1) Est (Stder)	(A2) Est (Stder)	(A3) Est (Stder)	(B1) Est (Stder)	(B2) Est (Stder)	(B3) Est (Stder)	(C1) Est (Stder)	(C2) Est (Stder)	(C3) Est (Stder)	(D1) Est (Stder)	(D2) Est (Stder)	(D3) Est (Stder)
AFFECTED	-0.0632	-0.0533	—[a]	-0.1296	-0.1702	—[a]	-1.9489***	-1.9861***	-1.9211***	-1.0251***	-0.3195	-0.3195
	0.1169	0.1284		0.4938	0.4920		0.4233	0.3754	0.3678	0.3377	0.2455	0.3982
BOTHSTAY		-0.2943			0.6485			0.2022	0.6836		-1.5029***	-2.4190***
		0.2373			1.1351			0.6345	0.9391		0.5683	0.5454
BOTHCOVER									-6.3008***			0.9161
									2.2577			0.9825
Stock controls:												
DIFFTIMING	0.0025***	0.0025***		0.0117***	0.0116***		0.0219***	0.0219***	0.0209***	0.0147***	0.0151***	-0.0151***
	0.0009	0.0009		0.0032	0.0031		0.0036	0.0032	0.0031	0.0028	0.0032	0.0028
CONSTANT	-0.0262	-0.0262		0.7270**	0.7218**		1.5552***	1.5525***	1.4767***	0.2819	0.2727	0.2727
	0.0671	0.0680		0.2974	0.2888		0.3043	0.2637	0.2589	0.2208	0.2503	0.2240
N	407	407	407	561	561	561	744	744	744	539	539	539
med(ΔMSE_i)	-0.0126			0.4747			0.0148			-0.1550		
med(DIFFTIMING)	5.4			28.9			-24.5			-28.4		
#(AFFECTED=1)	137			197			383			224		
#(BOTHSTAY=1)	25			21			31			87		
#(BOTHCOVER=1)	2			1			4			17		

***: statistically significant at 1%; **: statistically significant at 5%; *: statistically significant at 10%

Coefficients for stock-level controls AVGMCAP and SDEVMCAP are not reported for convenience.

[a]: Results were not reliable, due to small number of observations with $BOTHCOVER = 1$.

forecast timing after the merger, with positive values indicating that forecasts were released earlier, on average, after the merger. Since more negative values of the LHS variable ΔMSE_i indicate more forecast improvements, we expect the regression coefficient on DIFFTIMING to be positive, implying that earlier forecasts lead to less forecast improvements after the merger.

For Mergers C and D, the coefficient on AFFECTED is negative and significant (the coefficients are, respectively, -1.9489 and -1.0251). Moreover, the magnitudes of these coefficients are economically nontrivial, in that they are large in comparison to the magnitudes of the unconditional median of the dependent variable, reported at the bottom of the table. Hence, for these two mergers, the regression results confirm the graphical evidence from Figure (1) that information pooling might be present.

Among the stock-level controls, the coefficient on DIFFTIMING is positive and significant across all four mergers, and all specifications of the regression. This is in the expected direction, and indicates that for stocks where the forecast was released sooner following the merger, the forecast improvements were smaller. At the bottom of the table, we present the unconditional median of the DIFFTIMING variable across the four mergers, which shows that forecasts tended to be released sooner after Mergers A and B, but later following Mergers C and D.

Next, we narrow our focus to smaller subsets of stocks for which information pooling should be stronger. In Columns A2, B2, C2, and D2 of Table 6, we report results for the regression when BOTHSTAY is added as a right-hand side variable. Only for Merger D is the coefficient on BOTHSTAY_i negative and significant (and equal to -1.5029), indicating larger forecasting improvements for the stocks for which both pre-merger analysts were retained. This is evidence of a stronger notion of information pooling.[13]

Finally, we further narrow the analysis to those stocks where both analysts continue to produce forecasts in the post-merger period. The results from the regression with the BOTHCOVER_i dummy included are reported in Columns C3 and D3 of Table 6 (we were only able to run this regression for Mergers C and D, but to the small number of observations where BOTHCOVER_i=1). The coefficient on BOTHCOVER_i is negative and significant only for Merger C, but not in Merger D. Hence, for Mergers C and D, we obtain evidence indicating that a stronger notion of information pooling (as captured by the negative coefficients on either the BOTH-

[13]The finding that information pooling is more prominent when both of the pre-merger analysts were retained also provides support against an alternative explanation for improved analyst performance following a merger, namely that retained analysts may work harder following a merger because of increased job security concerns following the merger. This alternative story does not provide an explanation for why a retained analyst's performance improves more after the merger when her former rival is also retained.

STAY or BOTHCOVER variables) may be an explanation for the post-merger forecast improvements. From Table 3, we see that these two mergers were the ones where the merging firms were most equal in forecasting ability before the mergers, which perhaps made information pooling more likely.

4.1.1 Robustness

One worry with the above regressions is that the changes in forecasting accuracy after the merger could simply reflect changes in forecasting accuracy around the time of the mergers due, for instance, to unanticipated business-cycle movements, but not directly related to the mergers. As a robustness check, we expand the sample to include the changes in MSE's for brokerages which did not participate in any merger, as a control group. The idea is that any time-specific factors affecting forecasting accuracy should impact on both the merging and non-merging brokerages, whereas the effects of the merger (such as information pooling) should predominantly affect the merging brokerages only. The modified regression is

$$\begin{aligned}\Delta\text{MSE}_{i,k} =& \alpha + \alpha_1 \cdot \text{MERGE}_k + \beta \cdot \text{AFFECTED}_i + \beta_1 \cdot \text{AFFECTED}_i * \text{MERGE}_k \\ &+ \gamma \cdot \text{BOTHSTAY}_i + \gamma_1 \cdot \text{BOTHSTAY}_i * \text{MERGE}_k + \alpha Z_i + \text{error}_{i,k}\end{aligned} \tag{9}$$

where the k subscript denotes different brokerage firms, and MERGE_i is a binary indicator for whether brokerage k is the merged brokerage. The sample includes all brokerages k and stocks i for which forecasts were submitted for at least two quarters before and after the merger. The main benefit from including the observations from the non-merging firms is that we can estimate the coefficient on MERGE_i, which measures the part of the forecasting changes due specifically to the mergers, and not due to changes across time in forecasting abilities which are common across all brokerages. The interaction of AFFECTED and BOTHSTAY with MERGE are now used to capture the incremental effects of the AFFECTED and BOTHSTAY indicators on the forecasting changes of the merged brokerage. A finding that β_1 and β_2 is negative would suggest that forecast improvements result from the merger, and are not due to time specific effects.[14]

The results are reported in Table (7). The coefficient on the interaction terms are negative and significant in most of the regressions, the sole exception being the negative but insignificant coefficient on AFFECTED*MERGE in Merger D. This furnishes strong evidence that the mergers had distinctive effects on the forecasting performance of the merged brokerage, relative to others non-merging brokerages. Moreover, comparing these results with the corresponding results in Table

[14] In these regressions, BOTHSTAY and MERGE are always equal to zero for the observations of the non-merging firms.

Table 7: Brokerage-level Forecast Improvements: Including Observations from Non-merging Brokerages

Results from median (quantile) regression; Dependent variable: ΔMSE_i

	Merger A	Merger B	Merger C	Merger D
Variable	Est (Stder)	Est (Stder)	Est (Stder)	Est (Stder)
MERGE	0.5025*** 0.0500	0.1583 0.1823	0.2844 0.2936	-0.1946 0.3020
AFFECTED	-0.0741*** 0.0208	-0.2453*** 0.0618	-0.3296*** 0.0630	0.1059** 0.0525
AFFECTED*MERGE	-0.5257*** 0.0790	-0.0214*** 0.2506	-0.6999** 0.3245	-0.1543 0.3356
BOTHSTAY	0.3633*** 0.0422	2.8797*** 0.1393	0.2855*** 0.1479	0.0467 0.1646
BOTHSTAY*MERGE	-0.6916*** 0.1511	-12.5833*** 0.5213	-12.8695*** 0.4813	-3.3433*** 0.5436
Stock controls:				
DIFFTIMING	0.0017 0.0001	0.0073*** 0.0004	0.0058*** 0.0004	0.0047*** 0.0004
CONSTANT	0.0426*** 0.0132	0.4863*** 0.0412	0.5847*** 0.0508	0.1816*** 0.0388
N	5908	5312	4219	4149

***: statistically significant at 1%; **: statistically significant at 5%; *: statistically significant at 10%

Coefficients for stock-level controls AVGMCAP and SDEVMCAP are not reported for convenience.

(6), the most striking change is that the effect of BOTHSTAY on the merged brokerage (which is equal to the sum of the coefficients on BOTHSTAY and BOTHSTAY*MERGE) is now more negative, and significant, across all four mergers.[15] Thus, these results demonstrate even clearer evidence of information pooling.

For the second set of specification checks, we consider alternative definitions of

[15]From inspections of the data, this appears to be driven by the fact that the average value of ΔMSE across all observations in this regression (from both the merging and non-merging firms) is positive, indicating worsening forecasts. On the other hand, the values of ΔMSE in the merging firms are, on average, either negative or slightly positive. Hence, compared with the overall sample, the values of ΔMSE in the merging firms are smaller, which explains the negative and significant coefficient on BOTHSTAY.

the dependent variable ΔMSE_i. Because improving forecast precision may not be and is likely not the goal of the mergers, what is deemed an improvement from a statistical perspective need not be an improvement from the brokerages' perspectives. In the regressions in Table (6), we measured brokerage-level forecast improvement as the difference between the post-merger MSE and the simple average (ie. taking $w_1 = w_2 = \frac{1}{2}$ in Eq. (1) of the pre-merger MSEs of the bidder and the target firm. This was used because simple averaging favors neither the bidder nor the target firm, nor does not it weigh the better performing brokerage pre-merger more or less than the weaker brokerage.[16]

Nevertheless, to assess the robustness of the results, we re-ran the regression (7) for alternative values of the weights w_1 and w_2, including the special cases of putting all the weight on the bidder firm and none on the target firm, and vice versa. The robustness check consists of reconsidering the brokerage-level regressions reported in Table (6) for alternative definitions of ΔMSE_i. We do not report the results for the sake of brevity, but summarize them here. The regression results are similar to those reported in Table (6) when we put larger weight on the pre-merger MSE of the target firm. It is only when we put increasingly heavy weights on the MSE of the bidder firm that the negative coefficients on AFFECTED and BOTHSTAY become less significant.[17] For Mergers A and C, however, we find that the results reported in Table (6) are robust across a wide range of alternative weighting schemes. Overall, we find that although the results are not uniform over all alternatives values of w_1 and w_2, the evidence for information pooling at the brokerage level occurring in the form of forecast improvements holds up for most values of the weights.

4.2 Analyst-level forecast improvements

While the evidence above suggests that brokerage-level forecast improvements occurred after three of the four mergers, we do not know if these improvements extend to the analyst level. Indeed, the benefits of information pooling may be larger at the analyst level especially for the analysts who were substandard prior to the mergers. To focus on analyst-level forecast changes, define

$$\Delta\text{MSE}_{i,j} \equiv \text{MSE}_{i,j}^{post} - \text{MSE}_{i,j}^{pre}$$

as a measure of the change in analyst j's forecast accuracy for stock i. Again, $\Delta\text{MSE}_{i,j} < 0$ indicates forecast improvements.

[16] Furthermore, the forecast combination literature finds that simple averaging often outperforms more sophisticated forms of averaging. See, for example, Timmermann (2005). Thus in a sense, simple averaging forms a harder to beat benchmark.

[17] However, the results reported in Table (6) are robust even when we set $w_1 = 0.85$, where w_1 denotes the weight on the bidder firm's MSE in the pre-merger benchmark.

In order to measure analyst-level forecasting changes, we have to restrict our sample to stocks which were covered, in the post-merger period, by either the bidder or target analyst. Let j denote the analyst who covers stock i after the merger. As in the previous section, we define dummy variables which indicate subsets of stocks where information pooling should be strongest. We now define

- $\text{RIVALSTAY}_{i,j} = 1$ if analyst j's former rival (ie. the analyst who covered stock i in the other brokerage before the merger) was retained in the merged brokerage.

The dummy variable equals one in two circumstances. First, if analyst j, covering stock i, worked at the bidder firm before the merger, then $\text{RIVALSTAY}_{ij} = 1$ if the analyst who covered stock i at the target firm before the merger was retained at the merged brokerage after the merger. Second, if analyst j, covering stock i, worked at the target firm before the merger, then $\text{RIVALSTAY}_{ij} = 1$ if the analyst who covered stock i at the bidder firm before the merger was retained at the merged brokerage after the merger.

Figure (2) plots the CDFs of $\Delta\text{MSE}_{i,j}$ separately for $\text{RIVALSTAY}_{i,j} = 0$ (in solid lines) and $\text{RIVALSTAY}_{i,j} = 1$ (in dashed lines). Compared to the brokerage-level graphs in Figure 1, forecast improvements are less apparent here. For Mergers A, B, and D, substantial portions of the $\text{RIVALSTAY}_{i,j} = 1$ graphs lie to the right of the $\text{RIVALSTAY}_{i,j} = 0$ graphs, indicating a deterioration in forecasting accuracy after the mergers. Only for Merger C is there evidence of forecast improvements.

Hence, from the graphs, the evidence for analyst-level forecast improvements is more mixed, which is confirmed in regression results. We run the LAD regression of:

$$\Delta\text{MSE}_{i,j} = \alpha + \beta \cdot \text{RIVALSTAY}_{ij} + \gamma Z_i + \delta W_j + \text{error}_i$$

separately for (i) the subsample of stocks covered in the post-merger period by the bidder analyst, which we call the "bidder stocks"; and (ii) the subsample covered by the target analyst, which we call the "target stocks". A finding that $\beta < 0$ would be consistent with the presence of information pooling.

In addition to the stock-level control variables Z_i , we also include analyst-specific covariates to control for possible analyst heterogeneity. These covariates W_j are: (1) PREMSE_j, analyst j's pre-merger mean-squared forecast error, taken across all the stocks covered by analyst j in the two years prior to the merger; and (2) DIFFNUM_j, the difference between the total number of stocks covered by analyst j in the year after the merger, versus the year before the merger. The first covariate controls for analyst-specific forecasting ability, while the second covariate controls for the "attention" that analyst j pays to each stock that she covers. Because of the small number of stocks where RIVALSTAY=1, only five of the eight regressions had reliable results, and are reported in Table 8.

Figure 2: Empirical Cumulative Distribution Functions for Analyst-level Forecast Improvements

X-Axis: $\Delta \text{MSE}_{i,j} = \text{MSE}_{i,j}^{post} - \text{MSE}_{i,j}^{pre}$, for stocks i and analyst j [a]

Y-Axis: empirical cumulative distribution function

Solid line: stocks i for which $\text{RIVALSTAY}_{i,j} = 0$

Dashed line: stocks i for which $\text{RIVALSTAY}_{i,j} = 1$

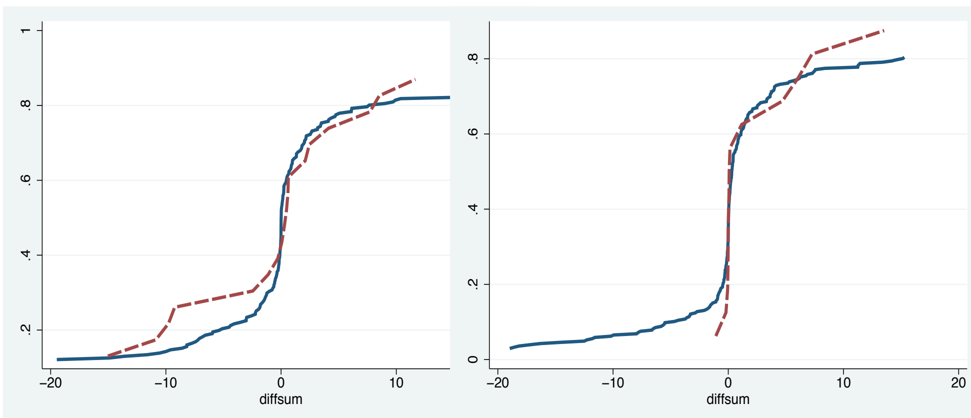

Merger A　　　　Merger B

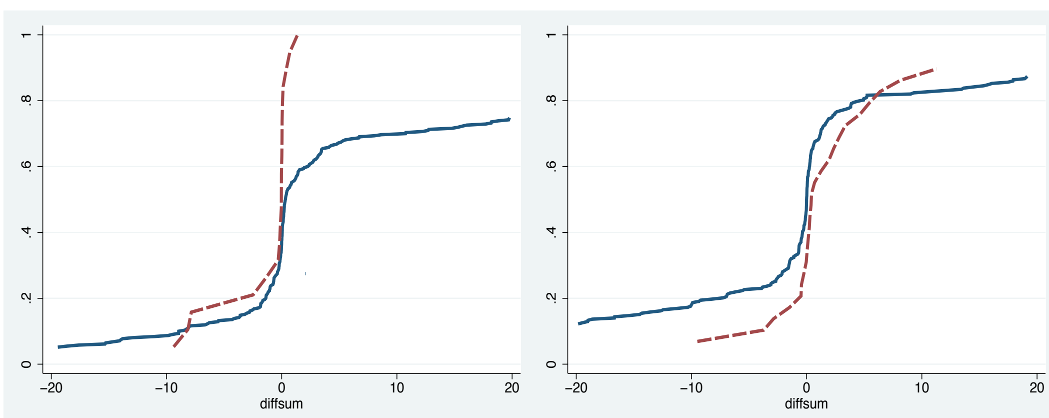

Merger C　　　　Merger D

[a] : $j \in \{bid, tar\}$ denotes analyst who forecast stocks after the merger. That is, $j = bid$ if stock was forecast by bidder analyst after the merger, and $j = tar$ if stock was forecast by target after the merger.

Table 8: Analyst-level forecast improvements

Results from median (quantile) regression; Dependent variable: $\Delta\text{MSE}_{i,j}$

	Merger A		Merger B		Merger C		Merger D	
	(A1) Bidder stocks[a]	(A2) Target stocks[b]	(B1) Bidder stocks	(B2) Target stocks	(C1) Bidder stocks	(C2) Target stocks	(D1) Bidder stocks	(D2) Target stocks
Variable	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)
RIVALSTAY	—[c]	-17.1619***	0.5514	—[c]	-0.7928	—[c]	0.2634	-9.3631***
		6.1988	0.9530		1.2877		0.5882	2.6488
DIFFTIMING		0.0580	0.0145***		0.0113*		0.0060*	0.0585***
		0.0446	0.0034		0.0064		0.0035	0.0150
PREMSE		-0.0180***	-0.0001		0.0000		-0.0015***	-0.092**
		0.0037	0.0004		0.0000		0.0004	0.0036
DIFFNUM		2.8192***	-0.1250*		0.0854		-0.0274	-0.1013
		0.3728	0.0667		0.1461		0.0810	0.1688
CONSTANT		17.5660***	0.3992		1.0538***		0.5067	1.7853
		5.1054	0.2925		0.3790		0.2450	1.1748
N	222	60	330	22	292	75	187	166
med(ΔMSE_i)	0.0031	0.3911	0.3451	0.0360	0.3303	-0.0074	0.1472	-0.5195
med(DIFFTIMING)	19.6	-7.4	43.8	-67.4	-5.8	-41.2	-6.2	-32.0
#(RIVALSTAY=1)	3	24	16	1	19	4	27	18

***: statistically significant at 1%; **: statistically significant at *5%; *: statistically significant at 10%

Coefficients for stock-level controls AVGMCAP and SDEVMCAP are not reported for convenience.

[a]: stocks which were covered in post-merger period by analyst who worked at bidder brokerage before merger

[b]: stocks which were covered in post-merger period by analyst who worked at target brokerage before merger

[c]: There were not enough stocks with RIVALSTAY= 1 to obtain reliable estimates for this regression.

For the bidder stocks, we find no evidence for information pooling. For Mergers B, C, and D, the coefficient on RIVALSTAY is statistically indistinguishable from zero. For the target stocks, however, there is some evidence of information pooling. For the two mergers where we had enough data to run this regression, we find a negative and significant coefficient on RIVALSTAY: -17.16 for Merger A, and -9.36 for Merger D. These results suggest that information pooling occurs primarily for the target analysts covering a stock in the merged brokerage when the bidder analyst is also retained, but assigned to other stocks.

For the other control variables, we find that the coefficient on the timing variable DIFFTIMING continues to be positive, and significant in four of the five regressions.

4.2.1 Robustness and additional results

As with the brokerage-level regressions, we also considered analyst-level regressions (analogous to Eq. (7) that include the changes in MSE of the non-merging brokerages as a control group in the sample, to isolate the effects due specifically to the mergers. because the results in Table 8 The regression is

$$\begin{aligned}\Delta\text{MSE}_{i,j} =& \alpha + \alpha_1 \cdot \text{MERGE}_j + \beta \cdot \text{RIVALSTAY}_i + \beta_1 \cdot \text{RIVALSTAY}_i * \text{MERGE}_j \\ & + \gamma Z_i + \delta W_j + \text{error}_{i,j}\end{aligned} \tag{10}$$

where the j subscript denote different analysts, and MERGE_j is a binary indicator for whether analyst j works at the merged brokerage after the merger. A finding that β_1 is negative would suggest that there are changes in forecasting accuracy due to the merger, and not just to time effects. [18] Table (9) report the results, separately for the subsamples of bidder and target stocks in each merger. The results weaken the earlier results in Table 8. For the target stocks, the coefficient RIVALSTAY*MERGE remains negative and significant for Merger D, but no longer for Merger A. For the bidder stocks, the coefficient on this interaction continues to be small and insignificant. This regression shows that only for Merger D do we have robust evidence of information pooling, and only for the target stocks.

As our strongest test of information pooling, we also investigate an implication of information pooling motivated in section 2, whereby post-merger forecast improvements should be inversely related to the pre-merger correlation in the forecast errors of the analysts. Let ρ_i be the correlation between the errors in the bidder and target firms' pre-merger forecasts of stock i. In order to compute this correlation,

[18] In these regressions, RIVALSTAY and MERGE are always equal to zero for the observations of the non-merging firms.

Table 9: Analyst-level forecast improvements: Including forecasts of analysts at non-merging brokerages

Results from median (quantile) regression; Dependent variable: $\Delta \text{MSE}_{i,j}$

	Merger A		Merger B		Merger C		Merger D	
	(A1) Bidder stocks[a]	(A2) Target stocks[b]	(B1) Bidder stocks	(B2) Target stocks	(C1) Bidder stocks	(C2) Target stocks	(D1) Bidder stocks	(D2) Target stocks
Variable	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)
MERGE	—[c]	-1.5495	0.1102	—[c]	-0.1462	—[c]	0.0454	-0.1107
		3.2081	0.1387		0.1769		0.2209	0.3168
RIVALSTAY		0.5547	-0.0427		-0.5051***		-0.3079***	0.0124
		1.6502	0.1041		0.1321		0.0975	0.1368
RIVALSTAY*MERGE		1.3354	-0.5257		0.2274		0.3834	-2.5030***
		3.8299	0.3885		0.5127		0.3995	0.4296
PREMSE		-0.0020**	0.0000***		-0.0000		-0.0000	0.0000
		0.0008	0.0000		0.0000		0.0000	0.0000
DIFFNUM		-0.2229	-0.0151		-0.0110		-0.0133	0.0027***
		0.1678	0.0096		0.0100		0.0117	0.0007
DIFFTIMING		0.0376***	0.0060***		0.0097***		0.0038***	0.0125***
		0.0089	0.0005		0.0007		0.0008	0.0012
CONSTANT		-0.1246	0.2877***		0.5176***		0.3326***	0.6658***
		1.7813	0.0510		0.0542		0.0624	0.1053
N	1257	346	2022	108	2529	574	1300	1463

***: statistically significant at 1%; **: statistically significant at *5%; *: statistically significant at 10%

Coefficients for stock-level controls AVGMCAP and SDEVMCAP are not reported for convenience.

[a]: stocks which were covered in post-merger period by analyst who worked at bidder firm before merger

[b]: stocks which were covered in post-merger period by analyst who worked at target firm before merger

[c]: There were not enough stocks with RIVALSTAY= 1 to obtain reliable estimates for this regression.

we restrict attention to the stocks which were forecast by both bidder and target firms during at least two quarters before the merger, and for which both analysts were retained in the merged firm. In Table 10, we consider the regression

$$\Delta \text{MSE}_{i,j} = \alpha + \gamma \cdot \text{NEGCOR}_i + \text{error}_{i,j} \tag{11}$$

where NEGCOR_i is a dummy variable that is equal to one if $\rho_i < 0$. The discussion in section 2 suggests that the coefficient on NEGCOR should be negative. This is because a negative correlation will lead to larger forecast improvements, and more negative values for $\Delta\text{MSE}_{i,j}$. From the bottom of Table 10, we see that the number of stocks when NEGCOR_i=1 is a small fraction of the sample, so that we pooled both the bidder and target stocks observations to run the regression.[19]

The results show that the coefficient on NEGCOR is *positive* and significant for Mergers A,B,C, which rejects the strong implication of information pooling which we are testing. This may not be surprising because the results in section 2 were derived from a very stylized modeling framework.

5 Analyst selection

As shown earlier, all four mergers led to a great deal of employment turnover. The possibility therefore arises that forecasting improvements can also be due to analyst selection. As we explained in Section 2 above, analyst turnover following a merger can lead to two types of analyst selection. First, if the analysts who remained were systematically better than those who left, a brokerage-level improvement might result. Second, after a merger, the better of the bidder and target analyst can be chosen to forecast a given stock. In both of these cases, the forecast improvement would not be due to information pooling.

While information pooling and analyst selection can be observationally very similar, our data permit us to look for direct evidence of analyst selection by examining patterns in analyst retention and the post-merger assignment of analysts to stocks in our data. We examine the two types of analyst selection in turn.

1. Analyst selection in retention We start with the first type of analyst selection, and examine whether the better-abilitied analysts were retained in the merged firm, after the merger. Table 11 compares the pre-merger MSEs of the forecast errors, for the analysts who were retained following the mergers, and those who were not retained. The top of Table 11 shows that the difference in the median pre-merger MSE between the retained and non-retained analysts is insignificant across three

[19]The brokerage-level covariates AVGMCAP and STDMCAP are also included in the regressions, but the coefficients are not reported for the sake of brevity.

Table 10: Regressions of Change in Mean Squared Error of Forecast on Pre-Merger Correlation

	Merger A	Merger B	Merger C	Merger D
Variable	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)
NEGCOR	58.0008***	12.5190	7.9645***	-0.1487
	7.6633	11.6276	1.1970	3.9765
DIFFTIMING	-0.0447	0.2972***	-0.0105	-0.0005
	0.0817	0.1169	0.0099	0.0343
CONSTANT	4.1991	-11.9197	-1.5000	-1.3910
	5.4377	10.2791	0.6856	1.6906
#(NEGCOR=1)	2	3	2	12
med(ρ_i)	0.511	0.398	0.500	0.513
N	24	14	25	76

These regressions also included the stock-level controls AVGMCAP and STDMCAP.

***: statistically significant at 1%

**: statistically significant at 5%

*: statistically significant at 10%

[a]: stocks which were covered in post-merger period by analyst who worked at bidder firm before merger

[b]: stocks which were covered in post-merger period by analyst who worked at target firm before merger

[c]: There were not enough target stock observations with NEGCOR$=1$ to permit estimation of this regression

of the mergers. However, the difference is negative and significant for Merger C (-0.1312) which is evidence of analyst selection.

Once we break down the numbers among the bidder and target firms, the evidence for analyst selection becomes even more mixed. For bidder analysts, the median MSE amongst retained analysts is less than that of non-retained analysts across all four mergers, which is evidence of positive selection. However, this difference is significant only in Mergers A and C. For example, for the bidder analysts in Merger C, the median MSE for the non-retained analysts is 0.3074, but lower (0.0309) for the retained analysts. For the target analysts, the selection is generally in the *adverse* direction, with the median MSE of retained analysts higher than that of non-retained analysts in three of the four mergers (except Merger C). Only

Table 11: Are retained analysts adversely or positively selected?
Pre-Merger MSE Among Retained and Non-retained Analysts

Status[a]	Employer[b]		Merger A	Merger B	Merger C	Merger D
Retained	Both	Med. MSE	0.0387	0.0173	0.0354	0.0550
		N[c]	40	60	110	100
Not Retained	Both	Med. MSE	0.0538	0.0294	0.1666	0.0658
		N	47	24	93	48
Δ Med MSE	Both		-0.0151	-0.0121	-0.1312**	-0.0108
Retained	Bidder	Med. MSE	0.0366	0.0100	0.0309	0.0410
		N[d]	31	55	93	61
Not Retained	Bidder	Med. MSE	0.0898	0.0144	0.3074	0.0774
		N	8	8	31	22
Δ Med MSE	Bidder		-0.0532**	-0.0044	-0.2765*	-0.0364
Retained	Target	Med. MSE	0.1770	0.1222	0.1409	0.0742
		N	9	5	17	39
Not Retained	Target	Med. MSE	0.0307	0.0335	0.1586	0.0621
		N	39	34	62	26
Δ Med MSE	Target		0.1463	0.0887*	-0.0177	0.0121

***: significant at 1%
**: significant at 5%
*: significant at 10%

[a] Was analyst retained or not retained, in merged brokerage?

[b] Analyst's pre-merger employer (either bidder or target firm)

[c] The total number of analysts in each column may be fewer than the number of pre-merger analysts given in Table 2, because in this table, we eliminate analysts who do not provide at least four quarterly forecasts for a given stock.

[d] The total number of analysts in each column may be fewer than the number of pre-merger analysts given in Table 2, because in this table, we eliminate analysts who do not provide at least four quarterly forecasts for a given stock.

in Merger B, where the median MSE for retained analysts is 0.1222 and for non-retained analysts is 0.0335, is the difference significant.

2. Analyst selection in post-merger stock assignment To examine the second type of analyst selection, arising if the better of the bidder and the target analyst might be chosen to forecast a given stock, we restrict our attention to the subset of the affected stocks which were (i) forecast by either the bidder or target firm

Table 12: Determinants of post-merger stock assignment

Sample is restricted to affected stocks which were (i) forecast by either the bidder or target firm analyst, after the merger; and (ii) *both* the bidder and target analysts were retained in the merged brokerage.

		Merger A	Merger B	Merger C	Merger D
	Total N	23	20	27	73
of which:					
(A):	#(analyst w/lower stock MSE chosen):	14	9	14	27**
(B):	#(analyst w/lower overall MSE chosen):	20***	18***	20**	35
(C):	#(analyst w/longer tenure chosen):	2***	14*	7**	40
(D):	#(analyst from bidder brokerage chosen):	1***	19***	24***	27**

Stars denote significance of a test for whether the probability is equal to $\frac{1}{2}$, with 3/2/1 stars denoting a p-value below 1/5/10% under the null hypothesis that $p = \frac{1}{2}$. The t-test statistic is $\frac{\hat{p}_N - \frac{1}{2}}{\sqrt{\frac{1}{N}\frac{1}{2}^2}}$, where $\hat{p}_N$ denote the fraction of stocks which satisfy each criterion. Asymptotically, this statistic is distributed standard normal under the null that $p = \frac{1}{2}$.

analyst, after the merger; and for which (ii) *both* the bidder and target analysts were retained in the merged brokerage. In Table 12, we investigate whether stocks were systematically assigned to analysts with better abilities, based on pre-merger forecasting performance.

In row (A), we consider whether the stock was assigned to the analyst who was better at forecasting this particular stock before the merger. We find no evidence of this. For Mergers A, B, and C, we could not reject that the stock was assigned randomly. For Merger D, our evidence indicates that the stock is more likely to be assigned to the analyst who was worse at forecasting this particular stock before the merger.

However, row (B) provides a partial explanation of this. Here, we consider whether the stock was assigned to the analyst with better *overall* performance (across all stocks that he/she forecast) before the merger. We find strong evidence in favor of this hypothesis. For Mergers A, B, and C, the overwhelming majority of stocks are assigned to the analyst with the better overall pre-merger performance, strong evidence of analyst selection in stock assignment. Only for Merger D do we find no evidence of this type of analyst selection.

Table 13: Determinants of post-merger assignment of analysts to stocks

Results from probit regressions where left-hand side variable is
$\mathbf{1}(\text{analyst } j \text{ assigned to stock } i \text{ after merger})$.
Sample is restricted to affected stocks which were (i) forecast by either the bidder or target firm analyst, after the merger; and (ii) *both* the bidder and target analysts were retained in the merged brokerage.
Each observation is a (stock i, analyst j).

	Merger A	Merger B	Merger C	Merger D
	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)	Estimate (Stder)
STOCK_BEST_{ij}	–[a]	-1.0690	-0.3063	-0.0973
		0.8992	0.4279	0.2234
OVERALL_BEST_j		1.7817***	1.1130***	0.5783**
		0.8432	0.4092	0.2365
TENURE_j		0.5309	-0.4728	-0.1592
		0.6643	0.2983	0.0986
BIDDER_j		2.4935**	0.9021*	-1.0958***
		1.0261	0.5158	0.2323
Constant		-2.9731*	-0.2274	0.4729*
		1.6782	0.8540	0.2888
N		46	63	170

[a]: could not estimate due to collinearity of the regressors.

These findings are echoed in Table 13, which contains results from probit regressions where the dependent variable is an indicator for whether an analyst j is assigned to cover stock i after the merger. As right-hand side variables, we included the two indicators of pre-merger forecasting superiority used in Table 12 (STOCK_BEST and OVERALL_BEST), as well as two additional analyst-level controls which likely influence stock assignment after the merger. These two controls are TENURE, measured as the number of years than the analyst was employed before the mergers, and BIDDER, an indicator for whether the analyst worked at the bidder firm before the merger. The one consistent result across the three mergers for which we were able to run the regression is the positive and significant coefficient on OVERALL_BEST, an indicator for whether analyst j had the better overall forecasting performance before the merger. The signs and magnitudes for the other variables varied across the mergers.

Summing up, our analysis of turnover and coverage patterns in this section yields no evidence for the first type of analyst selection, that better analysts are more likely to be retained in the merged brokerage following the merger. This confirms anecdotal evidence that in the wake of job uncertainty due to the mergers, many of the best analysts at the merging firms were poached away by competing brokerages, so that the analysts remaining at the merged brokerage following the merger are not the best analysts working at the two brokerages before the merger. However, in the cases where both of a stock's pre-merger analysts were retained in the merged brokerage, we find strong evidence (for three of the four mergers), that the stock is likely to be assigned after the merger to the analyst with the better overall pre-merger forecasting performance. This suggests that analyst selection can be a mechanism generating the post-merger forecasting improvements.

6 Conclusions

We exploit four large mergers of brokerage firms in the last decade to examine whether the patterns of changes in forecasting accuracy following the mergers can be attributed to information pooling. Given the large differences between the brokerages involved in the mergers, and the motives for the merger, it is not surprising that our results varied across mergers. However, several conclusions can be drawn. First, at the brokerage-level, we find some evidence of information pooling across two of the four mergers (Mergers C and D), in that forecast improvements were larger in the subsample of stocks which were covered by both of the merging brokerages before the merger, and the subsample where both of the pre-merger analysts were retained in the merged brokerage. These are indeed situations where information pooling should be strongest. Furthermore, the merging brokerages in these two mergers were of roughly equal forecasting ability before the mergers, which perhaps made information pooling more likely.

Second, at the analyst-level, we find no general evidence of forecast improvements, except for Merger D. For this merger we found that the post-merger forecasts of analysts from the target firm benefit more from the presence of the analyst who covered the same stock at the bidder firm around than vice versa. We also find evidence that after a merger, a stock is more likely to be assigned to an analyst with overall better forecasting performance before the merger. This suggests that analyst selection can be a mechanism generating the post-merger forecasting improvements.

Finally, the results for Merger D yield the strongest evidence consistent with information pooling. Only for this merger do we find evidence consistent with information pooling in both the brokerage-level regressions and the analyst-level regressions (for the target stocks only). Moreover, only for this merger do we find no evidence of analyst selection. This finding of information pooling after Merger D also corroborates to some extent the coverage of this merger in the business press. Particularly, the general perception is that the analysts from Paine Webber (the target brokerage) were absorbed into the merged brokerage without much problem (cf. *Crain Communications, Inc*. (2000)). Perhaps this post-merger collegiality between the Paine Webber and UBS analysts explains the evidence of information pooling that we found for this merger.

References

BERNHARDT, D., AND E. KUTSOATI (1999): "Can Relative Performance Compensation Explain Analysts' Forecasts of Earnings?," Tufts University Dept. of Economics discussion paper 99-09.

CALOMIRIS, C., AND T. PORNROJNANGKOOL (2005): "Monopoly-Creating Bank Consolidation? The Merger of Fleet and BankBoston," mimeo, Columbia Business School.

Crain Communications, Inc. (2000): "One on One: UBS didn't buy us to change (us). They recognize the value of a financial advisor (Interview with Paine Webber's Mark Sutton)," September 11, pg. 58.

DIAMOND, D. (1984): "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 51, 393–414.

FOCARELLI, D., AND F. PANETTA (2003): "Are mergers beneficial to consumers? Evidence from the market for bank deposits," *American Economic Review*, 93, 1152–1172.

GALLO, G., C. GRANGER, AND Y. JOON (2002): "Copycats and Common Swings: The Impact of the Use of Forecasts in Information Sets," IMF Staff Reports, vol. 49, no. 1.

GENESOVE, D., AND W. MULLIN (1999): "The Sugar Institute Learns to Organize Information Exchange," in *Learning by Doing in Markets, Firms, and Countries*, ed. by N. Lamoreaux, D. Raff, and P. Temin. University of Chicago Press.

HONG, H., J. KUBIK, AND A. SOLOMON (2000): "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *RAND Journal of Economics*, 31, 121–144.

KANDEL, E., AND N. PEARSON (1995): "Differential Interpretation of Public Signals and Trade in Speculative Markets," *Journal of Political Economy*, 103, 831–872.

KEANE, M., AND D. RUNKLE (1998): "Are Financial Analysts' Forecasts of Corporate Profits Rational?," *Journal of Political Economy*, 106, 768–805.

LIM, T. (2001): "Rationality and Analysts' Forecast Bias," *Journal of Finance*, 56, 369–385.

PANETTA, F., F. SCHIVARDI, AND M. SHUM (2006): "Do Mergers Improve Information? Evidence from the Loan Market," Manuscript, Bank of Italy and Johns Hopkins University.

PRAGER, R., AND T. HANNAN (1993): "Do Substantial Horizontal Mergers Generate Significant Price Effects? Edvidence from the Banking Industry," *Journal of Industrial Economics*, 46, 433–452.

RAJAN, R., AND H. SERVAES (1997): "Analyst Following of Initial Public Offerings," *Journal of Finance*, 52, 507–529.

SAPIENZA, P. (2002): "The effects of banking mergers on loan contracts," *Journal of Finance*, 57, 329–266.

SNYDER, C., AND M. DOYLE (1999): "Information Sharing and Competition in the Motor Vehicle Industry," *Journal of Political Economy*, 107, 1326–1364.

TIMMERMANN, A. (2005): "Forecast Combinations," forthcoming in *Handbook of Econometrics*.

VIVES, X. (1999): *Oligopoly Pricing*. MIT Press.

ZITZEWITZ, E. (2001): "Measuring Herding and Exaggeration by Equity Analysts and Other Opinion Sellers," Manuscript, Stanford University GSB.