

The Value of Research ^{*} [†]

Bryan Kelly
Stern School of Business
New York University

Alexander Ljungqvist
Stern School of Business
New York University
and *CEPR*

December 16, 2007

^{*} We thank Xavier Gabaix, Joel Hasbrouck, Dirk Jenter, Ralph Koijen, Holger Mueller, Mitchell Petersen, Ilya Strebulaev, Stijn van Nieuwerburgh, and Bernard Yeung for useful suggestions and seminar audiences at American University, Stanford GSB, NYU, Melbourne, and the University of Queensland for helpful comments. We are grateful to Stephen Brown of Emory University for sharing his PIN data and to Ashwani Kaul and Rajiv Kumar at Reuters Estimates for generous help in assembling our dataset. All errors are our own.

[†] Address for correspondence: Stern School of Business, New York University, 44 West Fourth Street, New York NY 10012-1126. Kelly: Phone 212-998-0368; e-mail bkelly@stern.nyu.edu. Ljungqvist: Phone 212-998-0304; fax 212-995-4220; e-mail al75@nyu.edu.

The Value of Research

Abstract

We estimate the value added by sell-side equity research analysts and explore the links between analyst research, informational efficiency, and asset prices. We identify the value of research from exogenous changes in analyst coverage. On announcement that a stock has lost all coverage, share prices fall by around 110 basis points or \$8.4 million on average. The share price reaction is attenuated the more analysts continue to cover the stock, suggesting that there are diminishing returns to coverage at the margin. The adverse effect of coverage terminations is proportional to the analyst's reputation and experience and to the size of the broker's retail sales force. Exogenous reductions in coverage are followed by: less efficient pricing and lower liquidity; greater earnings surprises and more volatile trading around subsequent earnings announcements; increases in required returns; and lower idiosyncratic volatility. Simulations suggest investors can trade profitably on the changes in volatility. Finally, retail investors sell and large institutional investors buy around coverage terminations, suggesting that different investor clienteles have different demands for analyst research.

Key words: Sell-side research; Coverage terminations; Informational efficiency; Trading strategies; Global Settlement.

JEL classification: G12, G14, G18, G24

Information is the lifeblood of financial markets. To learn about firm-specific or common factors affecting stock prices, investors buy analysis from information producers, ranging from the *Wall Street Journal* to investment newsletters and sell-side equity research departments at brokerage houses. In this paper, we estimate the value added by equity research analysts and explore the links between analyst research, informational efficiency, and asset prices.

Whether analysts add much value is controversial. Some view analyst research as an important channel through which information is impounded in stock prices (see, for instance, Grossman and Stiglitz (1980) and Gleason and Lee (2003)). Others view analysts mainly as cheerleaders for companies, who produce biased research in the hope of currying favor with executives (see, for instance, James and Karceski (2006)). The empirical evidence suggests analyst research does convey information, in the sense that recommendation changes can move prices (Womack (1996), Barber et al. (2001), and Jegadeesh et al. (2004)). But this leaves open the question whether analyst research is a zero-sum game (one investor's gain is another investor's loss), or whether analysts create value over and above trading gains.

How might analyst research create value? Suppose the assumptions of the Gordon growth model hold. Then, $P = D/(r - g)$, where P is the stock price, D is the future dividend, r is the discount rate, and g is the dividend growth rate. For analyst research to affect stock prices, it must affect either r or g . One channel through which it may affect r is informational efficiency: If analyst research makes a stock more informationally efficient and thus more liquid, trading costs and hence r may fall (Amihud and Mendelson (1986)). Analyst research may also have real consequences. Greater information production may, for instance, lead to better monitoring of managers. This, in turn, could increase g and so firm value. Similarly, it may affect the cost of external funds and hence the firm's ability to optimally invest in positive NPV projects.

In this paper, we focus on the asset pricing consequences of analyst research, leaving an examination of the potential real consequences to future work. We identify the value and function

of analyst research from exogenous variation in the supply of research caused by recent structural changes in the U.S. equity research industry. Our tests show that as brokerage firms reduce analyst coverage, share prices fall; that these falls are permanent; that informational efficiency and liquidity decrease; that the required return r increases; and that firm-specific volatility decreases.

Stocks are constantly added to and dropped from brokers' coverage lists. As McNichols and O'Brien (1997) observe, coverage changes are usually endogenous. Initiations may reflect an analyst's positive opinion of a company's future performance, resulting in large abnormal returns (e.g., Demiroglu and Ryngaert (2005)). Terminations, in contrast, are often viewed as implicit sell recommendations. The resulting share price fall may hence reflect the revelation of an analyst's negative view of a firm's prospects rather than the effects of reduced research coverage. Similarly, an analyst may drop a stock because institutional investors have lost interest in the stock (Xu (2006)). If institutional interest correlates with informational efficiency, we may find a spurious correlation between coverage terminations and decreases in efficiency.

To avoid these biases, we construct a set of coverage terminations that are the result of brokerage firms downsizing their research operations in response to adverse changes in the economics of producing research in the early 2000s. At least 20 firms quit the research business altogether. FleetBoston, for example, closed its Robertson Stephens unit in July 2002, with a loss of more than 40 analyst positions and coverage terminations on nearly 600 stocks. Other brokers reduced headcount, often drastically. To illustrate, on May 23, 2003, Citigroup dropped coverage of eight of the 43 sectors its analysts had covered.¹ We identify more than a thousand instances of brokerage firms terminating sector coverage between 2000 and 2005.

Sell-side research was downsized because market developments and new regulations had undermined its business model. While producing research is expensive, it does not generate any direct revenue, as few brokers charge investors directly for their research. Instead, research is

¹ The eight affected sectors were metals/mining, life sciences, utilities, healthcare services, airlines, industrials, specialty chemicals, and telecom equipment/wireless.

funded through cross-subsidies from trading, market-making, or investment banking. Each of these revenue streams diminished in the early 2000s. The prolonged decline in trading volumes that accompanied the bear market of 2000-2003 reduced brokerage revenue and income from market-making activities. And since the 2003 Global Settlement, investment banking has been prohibited from cross-subsidizing research, in an effort to protect investors from conflicts of interest.²

As sell-side research declined, buy-side research expanded. Yet buy-side research is proprietary; it is not shared with other investors. So while the net amount of information produced may not have changed, information is now more unequally distributed, to the possible detriment of those whose marginal cost of producing substitute research is high (such as small institutions or retail investors).

Our sample contains 16,253 coverage terminations over the period 2000-2005. Are they plausibly exogenous? First, we note that all results (with one minor exception) are robust to restricting the sample to the 1,792 terminations caused by brokerage closures. Such closures are very unlikely to correlate with variables of interest, such as performance. Sector terminations, on the other hand, might, but we find no evidence that brokers' choice of which sectors to drop reflect either prior or future performance. Finally, we show that sample terminations appear to carry no information about future operating performance, unlike a control group of endogenous terminations.

We find that reductions in the supply of information have a direct, adverse effect on asset prices. Cumulative abnormal returns average minus 47 basis points on the day of an exogenous coverage termination, increasing to minus 64 basis points by day +3. This effect appears to be permanent: There is no bounce back over the first month of trading, and over the next six to 24 months, terminated stocks perform in line with risk and style benchmarks. Thus, we find no evidence that the market over- or underreacts to coverage terminations.

The market reaction is strongest among the 838 sample stocks that lose all analyst coverage. On

² Analysts may come under pressure to publish biased research to please investment banking clients (Michaely and Womack (1999), Dugar and Nathan (1995), Lin and McNichols (1998)), to boost commissions (Irvine (2004), Jackson (2005), Agrawal and Chen (2004), Cowen, Groysberg, and Healy (2006)), or to curry favor with company managers (Francis and Philbrick (1993), Das, Levine, and Sivaramakrishnan (1998), Lim (2001)).

average, orphaning a stock reduces its share price by around 110 basis points on announcement, equivalent to an \$8.4 million reduction in market value. This provides an estimate of the capitalized value that sell-side research can add. The share price reaction around terminations is attenuated the more analysts continue to cover the stock, suggesting there are diminishing returns to producing information at the margin. It is more negative the more experienced the analyst, which hints at heterogeneity in the quality of information produced. Finally, the larger the brokerage firm's retail sales force, the more negative is the market's reaction.

Coverage terminations are accompanied by abnormally high trading volumes as investors rebalance their portfolios. Moreover, using CDA/Spectrum data, we show that institutional investors are unusually large net buyers in the quarter of a termination announcement, and especially when coverage drops to zero. By implication, the net sellers are individual investors and institutions without a duty to file 13f reports. One interpretation of these patterns is that individual investors and smaller institutions are more dependent on sell-side analyst research, so a coverage termination may reduce their valuation and hence their demand for the stock.

Changes in informational efficiency following coverage drops support this interpretation. Using difference-in-difference tests to control for secular and unrelated trends over the sample period, we show that three popular measures of informational efficiency, the probability of informed trading (PIN), Amihud's (2002) illiquidity measure, and Lo and MacKinlay's (1988) variance ratios, deteriorate following exogenous coverage terminations. Moreover, these reductions in informational efficiency are associated with more volatile returns around subsequent earnings announcements and with greater earnings surprises. As coverage is reduced, the average magnitude of earnings surprises increases by 16.9% compared to style-matched control firms whose coverage remains unchanged.

Deterioration in informational efficiency and liquidity is associated with changes in investors' required returns. Using standard factor models, we show that factor loadings change to produce an

increase of seven basis points in average required returns. The direction of the change is consistent with share prices falling around exogenous coverage terminations. Moreover, the magnitude of the change is commensurate with our point estimates of the announcement-day share price falls.

Roll (1988) views high firm-specific volatility as evidence of active trading by informed arbitrageurs who ensure that stock prices closely reflect fundamental value. If this is correct, we expect volatility to fall following exogenous coverage terminations, and it does. Controlling for the fact that volatility has trended down over the sample period, our difference-in-difference estimates range from minus 1.39 to minus 1.85 volatility points. When we simulate trading strategies (using zero-beta straddles or butterflies) designed to capitalize on these changes in volatility, we find that they are profitable, both economically and statistically. Trading at the bid and ask to allow for trading costs, our simulations suggest that shorting straddles following coverage terminations yields average daily returns of 1.36%, or 0.27% more per day than an equivalent position in options on style-matched control firms whose coverage remains unchanged.

Our paper is related to four contemporaneous studies that also analyze changes in analyst coverage. The paper closest in spirit to ours is Hong and Kacperczyk (2007). Like us, these authors exploit exogenous variation in coverage (in their case, variation due to mergers among brokerage firms with overlapping coverage universes), though their focus is different. Their main finding is that reductions in coverage are followed by an increase in forecast optimism bias. They interpret this finding as evidence of competition among analysts mitigating a tendency to produce biased research. Kecskes and Womack (2007) focus on investor overreactions to (presumably mostly endogenous) changes in the number of analysts covering a stock. Increases (decreases) in covering intensity in their sample are associated with lower (higher) future returns. Khorana, Mola, and Rau (2007) examine the causes of complete loss of analyst coverage. By design, they focus on endogenous coverage terminations, finding that small, poorly performing, highly leveraged

companies at high risk of bankruptcy are more likely to lose coverage. Loss of coverage is followed by a higher incidence of involuntary delisting, suggesting that analysts in their sample dropped coverage in part due to foreknowledge of impending distress. Finally, Scherbina (2007) reports evidence that analysts with negative signals about future earnings terminate coverage rather than lowering their earnings forecasts, in line with the arguments in McNichols and O'Brien (1997).

Our contribution is threefold. First, we identify a source of variation that gives rise to exogenous changes in the supply of information in the stock market. We believe this can serve as a useful instrument for empirical work in either asset pricing or corporate finance that examines questions of asymmetric information. Second, we provide a quantitative estimate of the value of sell-side research. At a time when the sell-side research business model is in flux, estimates of the value added by analysts may provide a basis for business model innovation. Third, we quantify the relation between analyst coverage on the one hand and liquidity and informational efficiency on the other, both of which are important considerations in asset pricing and portfolio formation.

I. Sample and Data

A. Sample Construction

We use several sources to construct a dataset of coverage terminations, and then impose filters to arrive at a set of terminations that can plausibly be treated as exogenous. From Reuters Estimates, we obtain a coverage table with precise coverage start and end dates for each broker/analyst/ticker combination. Defining a coverage termination as the absence of a reinitiation within two months³ of an end date, there are 64,596 terminations in the Reuters table over the period January 1, 2000 to December 31, 2005. However, Reuters does not identify why a broker dropped a stock. Most drops are likely firm-specific, i.e., endogenous. Tellingly, the 64,596 Reuters terminations typically affect only a minority of an analyst's coverage universe, suggesting that they are the result of a realignment of coverage in response to firm-specific information.

³ This is conservative. Brokers tend to reinitiate within ten days (usually following black-out periods) or not at all.

A.1 Sector Terminations and Brokerage Closures

We conjecture that sector terminations (as opposed to changes in firm coverage within a sector) are likely exogenous. We test this conjecture in Section II. As sector definitions vary by broker, we define sectors using six-digit GICS codes.⁴ Between 2000 and 2005, Reuters has 770 instances of a brokerage firm terminating coverage of an entire sector on a single day. This will understate sector terminations if GICS do not coincide precisely with a given broker's sector definition.⁵ To allow for this, we define a sector termination as any instance of a brokerage firm dropping coverage of at least 75% of the stocks it covered in a GICS sector. This yields a further 376 sector drops.

A subset of these 1,146 sector terminations relate to closures of research departments at 20 brokerage firms, ranging from large firms like Robertson Stephens and Wells Fargo to small outfits like IRG Research and J.B. Hanauer & Co. Collectively, the 1,146 sector terminations account for 11,504 terminations of stock coverage; 1,792 of these are due to brokerage closures.⁶

A.2 Termination Announcements and Termination Notices

We augment the 11,504 terminations from Reuters with 11,592 cases identified from other sources in which brokers disclose a reason for the termination.⁷ First, we search in Factiva and Bloomberg for announcements of sector terminations not already identified in the Reuters data. This yields a further four dropped sectors, taking the total number of sector terminations to 1,150. Second, since 2003, brokerage firms have had to explain their coverage termination decisions by issuing a termination notice.⁸ We search for termination notices in three text archives of analyst

⁴ Bhojraj, Lee, and Oler (2003) compare four industry classifications, concluding that “GICS ... are significantly better at explaining stock return co-movements, as well as cross-sectional variations in ... key financial ratios.”

⁵ Suppose some oil & gas stocks (GICS code 101020) are covered by the energy team. If the broker fires the oil & gas team but not the energy team, we will observe less than 100% of the oil & gas GICS coverage being terminated.

⁶ We treat the $64,596 - 11,504 = 53,092$ remaining Reuters terminations as likely endogenous.

⁷ Many terminations appear in more than one source. The numbers reported in the text are based on eliminating duplicates. In the rare instances where the sources disagree on announcement dates, we choose the earliest one.

⁸ NYSE Rule 472(f)(5) and NASD Rule 2711(f)(6) require that “a member provide notice to customers that it is terminating research coverage of an issuer that is the subject of a research report. ... [I]f the research analyst covering the subject company has left the member, or where the member has terminated coverage on an industry or sector ... the rationale for termination will be required.” See <http://www.sec.gov/rules/sro/34-48252.htm>.

reports (Investext, Thomson One, and the client-intranet of a large brokerage house).⁹ Termination reasons which we categorize as exogenous include: “This decision [...] relates to a disruption in analyst coverage and is unrelated to any awareness on our part of a change in the fundamental condition of the companies.” We ignore termination notices that give endogenous reasons, such as “The company no longer fits with our core investment objectives and coverage universe;” “owing to potential bankruptcy proceedings;” or “because of the company’s relatively low institutional ownership, low retail interest, and low trading volume.”

Finally, we identify stocks dropped for exogenous reasons from the “upgrades/downgrades” service on Dow Jones’s marketwatch.com service, which provides brief summaries of analyst reports (including termination notices) along with relevant announcement dates.

Combining the Reuters, newswire, text archive, and Dow Jones samples, we have 23,096 unique candidate events. For each of these, we record the date of the termination announcement, the reason for the termination, and the identity of the brokerage house and analyst. We identify the analyst covering the stock from the name on the termination notice or the Reuters coverage table where available, or else from the I/B/E/S earnings forecast database.

A.3 Filters

We filter out 6,843 of the 23,096 candidate events as follows. First, using CRSP delisting codes, we remove 434 stocks that were involuntarily delisted from an exchange within 60 days of a termination notice. This allows for the possibility that the analyst dropped the stock based on foreknowledge of impending distress (Khorana, Mola, and Rau (2007)), and so screens out potentially endogenous terminations. Second, to isolate the effects of reduced information production, we remove 5,396 terminations that are followed by a reinitiation by the same analyst at a new employer, or by a different analyst at the original brokerage firm, within three months of the drop. In addition, we exclude 528 REITs, 326 ADRs, 80 non-common stocks and closed-end funds

⁹ Keywords used include: “Terminate/discontinue/withdraw/suspend/cease/drop/stop coverage”, “Not rated/covered”.

(i.e., CRSP share codes > 12), and 79 companies without valid CRSP permno identifiers.

Our final sample consists of 16,253 coverage terminations, though the sample size will vary depending on the availability of variables of interest. It includes 4,150 unique stocks and spans all Fama-French industries, 179 brokerage houses, and 1,755 analysts. There were 1,955 terminations in 2000, 1,783 in 2001, 3,864 in 2002, 2,976 in 2003, 2,736 in 2004, and 2,939 in 2005.

How frequent are exogenous terminations relative to endogenous ones? The Reuters coverage table described earlier has 64,596 terminations. Not every broker contributes data to Reuters, so we add 52,869 additional terminations from the I/B/E/S database.¹⁰ This suggests that around one in seven terminations was exogenous over the 2000-2005 period ($16,253/(64,596+52,869)$).

B. Descriptive and Summary Statistics

Figure 1 plots smoothed histograms for the average number of days between sample terminations, estimated by broker or by stock using a kernel density estimator. Sample terminations clearly cluster at the broker level: For a given broker, the majority of drops occurs in a small time window (mean=12.3 trading days). This reflects the fact that brokerage houses terminated coverage of many stocks at the same time, e.g., when they quit equity research or terminated coverage of one or more sectors. At the stock level, on the other hand, the time between sample drops is spread quite evenly over a range from a few days to several years (mean=248.8 days). This lack of clustering at the stock level is consistent with sample terminations being exogenous: If a stock had been dropped by multiple brokers due to negative stock-specific information, we would have expected terminations to cluster. Though not shown, we find a similar lack of clustering at the sector level.

The figure also shows the density of the average number of days by stock between the 53,092 Reuters non-sector terminations, which we screened out as likely endogenous. These drops cluster visibly (mean=141.1 days), suggesting common factors (and so potential endogeneity) in brokers'

¹⁰ These are defined crudely as stocks covered by a broker in one year but not the next, as we lack an I/B/E/S coverage table with precise coverage dates. For the same reason, Kecskes and Womack (2007), who use I/B/E/S data, conduct their analysis at an annual frequency rather than in event time.

decisions to drop coverage on these stocks.

Table I compares summary statistics for the size, liquidity, and return volatility of the stocks in our sample to the CRSP universe of publicly traded U.S. stocks and to the universes of U.S. stocks covered by at least one brokerage firm according to Reuters Estimates and I/B/E/S, respectively. Sample stocks are on average larger and more liquid than the average stock in the CRSP universe. The comparison to the Reuters and I/B/E/S universes shows that this reflects a general tendency among analysts to cover larger and more liquid stocks, rather than a characteristic of companies experiencing coverage terminations. In addition, sample stocks – like Reuters and I/B/E/S stocks more generally – are less volatile than CRSP stocks. On average, 9.5 analysts covered a sample stock before a coverage drop, in line with the corresponding averages for the Reuters universe and the Reuters-I/B/E/S union. Overall, the average sample firm looks similar to the average firm in the Reuters and I/B/E/S databases, in terms of size, liquidity, volatility, and analyst coverage.

C. Extent of Prior Coverage

Table II reports a breakdown of the sample by the number of other brokers covering the stock in the six months preceding the coverage termination (based on pooling data from Reuters Estimates and I/B/E/S). Of the 16,253 sample terminations, 838 stocks were left orphaned following the termination, 4,031 were still covered by between 1 and 5 other analysts, 3,783 by between 6 and 10 others, 2,866 by between 11 and 15, and 4,735 by more than 15 other analysts.

Orphaned stocks are not necessarily smaller than other stocks that experience coverage terminations. The average orphaned stock had a market capitalization of \$1.4 billion one month before the termination – more than in the next bin (\$740.9 million). Only stocks covered by 11-15 or >15 other analysts are markedly larger on average (\$5.4 billion and \$22.1 billion, respectively). Median market capitalization, on the other hand, increases monotonically in coverage.

D. Matching Firms

We use a series of difference-in-difference (DiD) estimators to test for changes in sample firm characteristics around coverage terminations. Difference-in-difference tests are commonly used to remove biases due to omitted variables (Ashenfelter and Card (1985)). For instance, informational efficiency may change over time due to changes in market architecture or regulations rather than due to termination of coverage. The DiD tests compare the change in a variable of interest for sample terminations to the contemporaneous change in the same variable for a set of control firms matched to have similar characteristics. This removes biases due to secular trends affecting similar companies at the same time, such as omitted regulations.

The control groups consist of firms we expect to behave similarly to the event firms but that are themselves unaffected by the event. We follow Daniel et al. (1997), henceforth DGTW, and randomly choose as controls for firm i five unique firms in the same size, book-to-market, and momentum quintile in the month of June prior to a termination, subject to two conditions: a) Control stocks must not have experienced a coverage termination during the estimation window; and b) each control stock must have sufficient data to estimate the statistic of interest.¹¹

Table II shows the size, book-to-market, and momentum quintiles that sample terminations are assigned to. Terminated stocks are of average size (mean quintile=3.1), have below-average book-to-market ratios (mean quintile=2.5), and exhibit average momentum (mean quintile=2.9). When we condition on the extent of prior coverage, we again find that the extent of coverage is positively related to size. There is no obvious relation to book-to-market or momentum.

II. Exogeneity

What sets our analysis apart from contemporaneous studies of the effect of changes in coverage is our focus on stocks that are dropped for exogenous reasons. We have argued that only exogenous

¹¹ In our difference-in-difference tests, we lose 1,395 events involving stocks that do not satisfy DGTW's data conditions. We also drop observations with fewer than three viable controls for a given test. Depending on the variable of interest and estimation window, the number of observations ranges from around 10,000 to around 14,000.

terminations allow us to measure the value and function of analyst research consistently and have constructed our sample accordingly.

Perhaps the strongest evidence that our empirical results are free of endogeneity bias comes from the 20 brokerage closures. It is difficult to argue that a brokerage firm decided to quit equity research because the future prospects of the companies and sectors it covered were poor. Had prospects simply been poor, the brokerage firm could have changed its coverage universe, and brokerage firms do so routinely. The decision to quit research must thus reflect considerations that are uncorrelated with the future prospects of the affected companies and sectors. By comparing the results we obtain in the full sample of 16,253 terminations to those in the subsample of 1,792 closure-related terminations, we can gauge the success of our identification strategy. Our empirical results are very similar in either sample (with one exception, which we will point out later).

A large part of our sample comes from sector terminations. This could introduce a source of bias if brokers systematically chose to drop sectors whose characteristics correlate with variables of interest. We therefore investigate the empirical determinants of which sectors a broker chooses to drop using a McFadden (1974) choice model. The choice set includes all sectors the broker covers at the time of the decision. Brokers are assumed to select those sectors whose termination will maximize their expected utility (i.e., profit). Brokers might terminate sectors with poor past or poor anticipated performance, those covered by low-quality analysts, or those in which they add little value. The model is estimated using probit with or without random brokerage-firm effects. The dependent variable equals one if a sector is chosen for termination and zero otherwise.

Table III reports the results. We find no evidence that brokers chose which sectors to terminate based on prior sector performance, whether we focus on market-adjusted returns or earnings surprises during the previous four quarters. Nor do brokers appear to be clairvoyant: Their choices are unrelated to future sector performance measured over the following four quarters, both in terms

of average market-adjusted returns and earnings surprises. These coefficients are not only statistically insignificant, they are also economically small, as the marginal effects in the table indicate. Instead, brokers are significantly more likely to drop less volatile sectors and those covered by few other brokers (in the random-effects specification). Brokers tend to retain sectors covered by “all-star” or more experienced analysts. If higher-quality analysts tend to produce more valuable research, this implies that our estimates of the value of research will be conservative.

Finally, we test directly whether terminations signal private information about future prospects by analyzing subsequent earnings surprises. Suppose at time t , a broker drops coverage on a stock due to negative information about $t+1$ earnings that is not yet reflected in the consensus earnings forecast dated $t-1$. When earnings are eventually announced, they will fall short of consensus. For exogenous terminations, on the other hand, earnings will not disappoint systematically.

We implement this test in a panel tracking the universe of CRSP companies over the period 2000Q1 to 2006Q1 (filtering out REITs and companies with share codes greater than 12). The test includes the 53,092 non-sector Reuters terminations we filtered out as potentially endogenous alongside the 16,253 sample terminations.¹²

We first analyze earnings surprises (scaled by book value of assets) in a pooled OLS regression. Besides indicator variables for the sample and Reuters terminations, we control for lagged earnings surprises, share price returns and return volatility measured over the prior 12 months, the log number of brokers covering the stock, and year fixed effects. This yields the following estimates:

$$\begin{aligned}
 \text{earnings surprise} = & .364 \text{L(earnings surprise)} + .466 \text{return} - 4.814 \text{volatility} + .228 \text{no. brokers} \\
 & \quad \quad \quad .017 \quad \quad \quad .040 \quad \quad \quad .370 \quad \quad \quad .019 \\
 & - .126 \text{Reuters} + .016 \text{exog. termination} + \text{year effects} \\
 & \quad \quad \quad .043 \quad \quad \quad .040
 \end{aligned}$$

Adj. R^2 = 14.6% F = 84.5*** N = 65,629 firm - quarters (4,148 companies)

We report heteroskedasticity-consistent standard errors (clustered on CRSP permno) beneath the

¹² While there are 3.3 times as many Reuters terminations, they affect only 2.4 times as many firm-quarters as do the sample terminations (34.1% vs. 14.0%). As in Figure 1, this reflects the greater degree of clustering among the Reuters terminations.

coefficient estimates. Earnings surprises are serially correlated, more positive for high-return and low-volatility companies and among firms covered by many analysts. Controlling for these effects, earnings surprises are significantly more negative following Reuters terminations ($p = 0.004$). On average, a Reuters termination is followed by a quarterly earnings surprise that is 22.4% more negative than the sample average over this period. Sample terminations, by contrast, are unrelated to subsequent earnings surprises, consistent with our conjecture that they are exogenous.

The pooled OLS model ignores the panel structure of the data. Given serial correlation, we estimate a dynamic panel data model under the assumption that the errors follow an AR(1) process. To control for unobserved heterogeneity in firm characteristics, we include firm fixed effects, and we allow for unbalanced panels (see Baltagi and Wu (1999)). This yields the following estimates:

$$\begin{aligned} \text{earnings surprise} = & .449 \text{ return} - 2.238 \text{ volatility} - .011 \text{ no. brokers} \\ & \quad \quad \quad .024 \quad \quad \quad .348 \quad \quad \quad .036 \\ & - .107 \text{ Reuters} - .037 \text{ exog. termination} + \text{year effects} + \text{fixed effects} \\ & \quad \quad \quad .040 \quad \quad \quad .047 \\ \text{AR(1) term } (\rho) = & .304 \quad \quad \quad F = 42.9^{***} \quad \quad \quad N = 62,100 \text{ firm - quarters (3,995 companies)} \end{aligned}$$

We continue to find that Reuters terminations, but not sample terminations, correlate with subsequent earnings surprises and so appear to be endogenous.¹³

III. Price and Volume Patterns

A. The Value of Research

How best to quantify the value of sell-side equity research? One possibility is to compare the values of companies with and without coverage, though it is hard to hold other determinants of value constant in such a cross-sectional comparison. The alternative we propose is to focus on share price changes for the *same* company as the extent of its coverage changes. Over a sufficiently short time window, such as a few days, this ensures that other firm characteristics are held constant (as

¹³ As further evidence of the exogeneity of our sample, we replicate contemporaneous studies. In contrast to Kecskes and Womack (2007), sample firms do *not* underperform in the year before a termination (returning +6.4% above the market, compared to -2.7% in Kecskes and Womack). In contrast to Khorana, Mola, and Rau (2007), sample firms are *not* more likely to subsequently delist. And in contrast to Scherbina (2007), sample terminations do *not* predict negative abnormal returns around subsequent earnings announcements.

long as we focus on exogenous variation in coverage).

We compute abnormal returns using two benchmarks (the market model and the Fama-French three-factor model) from the close on the day before the termination announcement to the close on the announcement day [-1,0], one day later [-1,+1], or three days later [-1,+3]. Table IV reports the results. Regardless of benchmark or window, abnormal returns around coverage terminations are negative and statistically significant. They range from minus 47 to minus 67 basis points (not annualized) and increase somewhat the longer the event window.¹⁴ These findings suggest that exogenous reductions in coverage reduce firm value and thus that research coverage is valuable.¹⁵

To illustrate the economic magnitude of the capitalized value of research, we multiply a firm's equity market value in the month before a coverage termination by the Fama-French CAR for the [-1,+1] window. Following a termination, market value declines by around \$26.8 million on average, \$25.9 million if we winsorize 0.5% of observations in each tail, and \$1.95 million in the median event. Each of these point estimates is statistically significantly different from zero.

Though not tabulated, the corresponding abnormal returns in the subsample of 1,792 closure-related terminations average between minus 1.09% and minus 1.53%, depending on benchmark. The fact that market reactions are more negative when brokers quit equity research than when they selectively terminate sector coverage is consistent with our finding in Table III that brokers tend to drop sectors covered by lower-quality analysts. Below, we provide direct evidence that terminations involving lower-quality analysts involve smaller (less negative) market reactions.

Panel B of Table IV reports a breakdown of the CARs by the extent of prior coverage. Both economically and statistically, announcement returns are most negative among orphaned stocks.

¹⁴ Despite the widespread view that endogenous coverage terminations are implicit sell recommendations, we are not aware of studies estimating announcement returns for such implicit sells. For *explicit* sell recommendations (i.e., downgrades to sell), Womack (1996) finds average three-day excess returns of minus 3.87%.

¹⁵ The proviso is that investors realize the terminations are in fact exogenous. If investors mistrust a broker's stated termination reason, the share price falls shown in Table IV may instead capture their mistaken belief that the broker possessed negative private information about the affected stocks. However, contrasting terminations for which we have termination notices with brokerage closure-related terminations since 2003, we find no significant difference in market reaction.

These suffer abnormal returns of minus 90 basis points on the event day, minus 108-110 basis points by the end of day +1, and minus 131-138 basis points by day +3 ($p < 0.001$ in each case). Economically, this implies a fall in the market value of the average orphaned stock of \$8.4 million (\$9.2 million winsorized). CARs are negative and statistically significant in every bin and in nearly all specifications increase monotonically (become less negative) the more brokers continue to cover a stock. To illustrate, over the $[-1,+1]$ interval, orphaned stocks lose 108 basis points in the market model, while stocks covered by more than 15 other analysts lose only 40 basis points. This suggests that there are diminishing returns to coverage at the margin.

In Panel C, we report announcement returns by termination year. While there is some variation across years, average CARs are negative in every year and generally statistically significant. The largest market reactions occurred in 2002.

B. Cross-sectional Determinants of the Value of Research

In Table V, we report the results of OLS regressions with Fama-French CARs as the dependent variable. Besides firm fixed effects (using CRSP permnos) and announcement-year effects, we control for the extent of coverage by other brokers, analyst quality, and firm size. As in Table IV, we find that CARs are positively related to the extent of coverage (see column (1)). All else equal, a one-standard deviation decrease in the log number of other brokers covering the stock is associated with a 31 basis point lower (more negative) announcement effect. The market appears to react more negatively to terminations involving “all-stars” ($p=0.055$) or more senior analysts ($p=0.01$), confirming our conjecture that higher-rated and more experienced analysts add more value. The adjusted R -squared of 11.8% suggests this model has a reasonable fit.

In columns (2) through (4), we add proxies for the size of the broker/dealer dropping the stock. The larger a broker’s client base, the more investors are potentially directly affected by the reduction in information production, so we expect a larger (more negative) market impact. Our size

proxies are taken from the annual Factbook of the Securities Industry Association (SIA). While the SIA data are standardized, reporting is voluntary and availability of each data item varies across brokers. As the size proxies are highly correlated, we include them one at a time.

As predicted, we find a larger (more negative) announcement effect, the more brokerage accounts the broker manages ($p=0.021$), the greater the value of the brokerage accounts ($p=0.01$), and the larger the broker's sales force ($p=0.052$). Each of these effects is economically significant. To illustrate, one standard-deviation increases in these three proxies are associated with 10, 18, and 8 basis points lower (more negative) CARs. Thus, the adverse effect of a coverage termination appears to be proportional to the broker's reach.

Column (5) separately controls for the size of the retail and institutional sales forces. Interestingly, only the size of the retail sales force appears to matter ($p=0.003$). A one standard deviation increase in this variable (say, from Ladenburg, Thalman to Prudential Securities) is associated with a 13 basis point lower CAR. The fact that the reach of a broker's institutional-investor network appears to have no effect suggests that retail investors are more sensitive to changes in research coverage, perhaps because, in contrast to institutions, they typically receive research from only one brokerage house. In that case, a coverage termination may reduce their demand and hence their valuation for the stock. By contrast, many institutional investors have multiple broker relationships and possibly in-house research departments and so are less affected by reductions in sell-side research coverage.

C. Post-termination Return Behavior

Are these reductions in share prices transitory or persistent? The Fama-French CAR for the window $[-1,+30]$ is minus 43 basis points ($t = -2.75$), indicating there is no bounce back in the month after a coverage termination. To investigate share price behavior over the medium-term, we estimate standard calendar-time portfolio returns based on buying stocks in month 1 after a

coverage termination and selling them in month T , where $T = 6, 12, 18,$ or 24 months. We estimate abnormal performance as the intercept (alpha) in OLS regressions of these monthly portfolio returns in excess of the riskfree rate on the three Fama-French factors and the Carhart (1997) momentum factor. As an alternative to this factor model approach, we also compute style-adjusted portfolio returns by subtracting from each sample stock's monthly return the contemporaneous return on the relevant DGTW control portfolio, and weighting stocks in calendar time by the accumulated return since their entry into the portfolio. In this case, alpha is estimated as the time series average of the monthly time series of style-adjusted portfolio returns.

Table VI reports the results. For the sample as a whole, the monthly Fama-French-Carhart alphas over 6-, 12-, 18-, and 24-month holding periods are 0.50%, 0.29%, 0.27%, and 0.28%, respectively. These estimates are small economically, and none of them is statistically significant. The DGTW style-adjusted returns are smaller still and they too are statistically insignificant. When we split the sample by the extent of prior coverage, we continue to find that terminated stocks perform in line with risk and style benchmarks. These findings suggest that the announcement-day reductions in share prices shown in Table IV are permanent. Thus, the market neither over- nor underreacts to exogenous coverage terminations.

D. Trading Patterns Around Coverage Terminations

Coverage terminations are accompanied by abnormally high trading volumes as investors rebalance their portfolios. Using CRSP daily volume data, we define abnormal volume as the difference between log volume on each day over the window $[-5,+10]$ and average daily log volume over a 60-day estimation window ending 10 days before the termination announcement. Figure 2 plots the daily cross-sectional means of the abnormal changes in volume.

Trading volume is significantly higher on the announcement day and the day after, when a total of 5.9% more shares are traded than in the pre-event window. Volume remains higher than normal

for up to nine trading days, with cumulative abnormal volume of 13.0%. Curiously, trading volume is 2.8% higher on the day *before* the announcement ($p = 0.02$), though this is not accompanied by abnormal returns (+26 basis points on average in the market model).

E. Who Trades?

Upon announcement of a coverage drop, price falls and trading volume increases. Who trades? Absent high-frequency data identifying investors, we use CDA/Spectrum data to compute the change in the fraction of the dropped company's outstanding stock held by institutions required to file 13f reports.¹⁶ This is computed from the quarter-end before to the quarter-end after a termination announcement. Panel A of Table VII shows that 13f institutions as a group increase their holdings from 61.7% to 62.5% of shares outstanding following the average termination ($p < 0.001$).

This increase in institutional holdings could simply reflect broader portfolio trends unrelated to terminations, so we compute a difference-in-difference estimator using the control firms described in Section I.D. Institutional holdings in control firms are, on average, unchanged around sample terminations. The DiD estimate of 0.9% suggests that institutions increase their holdings in stocks suffering coverage drops by nearly one percentage points on average, a 1.5% average increase from their pre-termination holdings.

Is this significant? The time clustering of terminations evident in Figure 1 poses a problem for standard cross-sectional tests, so we compute p -values based on block bootstraps with 10,000 replications and Politis and White (2004) block lengths. The bootstrapped p -value indicates that the increase in institutional holdings is significant. Thus, 13f institutions are unusually large net buyers in the quarter of a termination announcement. By implication, the net sellers are retail investors and institutions without a duty to file 13f reports. This echoes our earlier finding that share price reductions around coverage terminations are related to the size of the broker's retail sales force.

¹⁶ Investment companies and professional money managers with over \$100 million under management are required to file quarterly 13f reports. Reports may omit holdings of fewer than 10,000 shares or \$200,000 in market value.

Panel B of Table VII shows a breakdown of the mean net changes in institutional holdings by the number of brokers covering the stock. Proportional institutional holdings increase the most when coverage drops to zero, by 3.3%, compared to 1.8% among stocks covered by between one and five other analysts.¹⁷ The relation between number of brokers and change in institutional holdings is broadly monotonic. The differences between orphans and the other four bins are all highly significant. These results are consistent with the view that retail investors (and small institutional investors) are particularly sensitive to changes in sell-side research coverage.

IV. Changes in Informational Efficiency

To test whether sell-side research contributes to informational efficiency, we use three popular proxies: The probability of informed trading (PIN), Amihud's (2002) illiquidity measure, and variance ratios. In addition, we examine changes in earnings surprises and return volatilities at earnings releases around coverage terminations. We use difference-in-difference estimators to help isolate structural breaks in informational efficiency due to terminations, as opposed to random, temporary fluctuations in efficiency. For each measure, we test the null that terminations have no effect on informational efficiency: $DiD = \text{mean} [(post_i - pre_i) - \overline{(post - pre)_{Control\ Group\ i}}] = 0$.

A. Probability of Informed Trading

Easley, Kiefer, and O'Hara (1997) propose the probability of informed trading (PIN) as a measure of informational efficiency. This is based on the idea that the presence of privately informed traders can be noisily inferred from order flow imbalances. One plausible result of reduced analyst coverage is that less information for the firm is publicly revealed, leading to an increased incidence of investors trading on private information and hence an increase in PIN.

Using quarterly PIN estimates, we calculate the difference $PIN_{t+1} - PIN_{t-1}$ for both sample and control firms, where t denotes the quarter of termination. Estimating PIN can be difficult and

¹⁷ Xu (2006) finds that institutions *reduce* their ownership of orphaned stocks in a sample that does not screen out endogenous terminations. Xu's finding is consistent with the view that endogenous terminations are implicit sells.

requires an iterative convergence procedure over a constrained interval. The PIN data we use flags potentially problematic estimates (see Brown, Hillegeist, and Lo (2004)). We report results both with and without flagged estimates, giving sample sizes of 11,575 and 9,663, respectively.

Table VIII, Panel A reports the results. The average PIN for sample firms is 0.141 before and 0.145 after a coverage drop. Subtracting the change in PIN for control firms, the average DiD estimate is 0.003, an increase of 2.1% (bootstrapped p -value = 0.011). If we exclude flagged PIN estimates, the DiD estimate increases to 0.005 (+3.3%) with a p -value < 0.001.

This finding suggests that coverage terminations reduce informational efficiency. To illustrate its economic magnitude, note that PIN captures transaction costs. Easley et al. (1996) argue that a stock's bid-ask spread should approximately equal $(0.0193 \cdot \text{price} \cdot \text{PIN})$. The average orphaned stock (in terms of announcement CAR) has a share price of \$27.49 and a bid-ask spread of \$0.44 in the month before the termination, and its PIN increases by 0.0275 post-termination, net of its control group's PIN change. All else equal, its spread is expected to widen by \$0.0146, or about 3.3%.

B. Liquidity

Amihud's (2002) illiquidity measure (AIM) attempts to measure the price impact of trades. In more liquid stocks, a trade of given size should have a smaller price impact. AIM is measured as the absolute return on a stock divided by the dollar trading volume over a given estimation window. To ensure the resulting distributions of levels, changes, and DiD are well-behaved, we follow Amihud and rescale AIM as follows: $\log \text{AIM} = \ln(10^6 \cdot \text{average daily AIM over estimation window})$.

In the case of AIM, it is particularly important to use a DiD estimator rather than rely on own-differences. The numerator of AIM is stationary (a return) whereas the denominator is non-stationary (price \cdot volume). If returns are constant over time but dollar volumes increase, AIM will appear to be decreasing, i.e., stocks will appear to become more liquid over time. The DiD estimator

allows us to test whether sample stocks become more liquid at a slower rate than do control firms, i.e., whether they become less liquid in the wake of a coverage termination.

Table VIII, Panel B reports the results for three- and six-month estimation windows ending 10 days prior to the termination announcement or starting 10 days after the announcement. Using a three-month window, average logAIM is -8.055 before and -8.086 after a coverage termination, implying that dropped stocks became more liquid (less illiquid). The DiD, on the other hand, averages 0.026, implying a decrease in liquidity following coverage terminations. The corresponding estimate for the six-month window is 0.040. Though relatively small economically, both estimates are statistically significant.

C. Variance Ratio Tests

The less informationally efficient a stock, the less rapidly will its price incorporate new information, so its return will behave less like a random walk resulting in serial correlation. Under the null of a random walk with uncorrelated increments, the variance of returns is a linear function of the time interval over which returns are measured (Lo and MacKinlay (1988)), so the ratio of the N -day variance to N times the one-day variance equals one. Variance ratios greater (less) than one indicate positive (negative) serial correlation. Either indicates a departure from a random walk, so we focus on the absolute value of the variance ratio less one, i.e., $|VR-1|$.

Table VIII, Panel C reports the results for 2-, 3-, and 4-day returns measured over three- and six-month estimation windows. The variance ratios invariably increase following terminations, by 0.9% to 2.6% on average. This is consistent with reduced informational efficiency.¹⁸

D. Earnings Announcements

A promising setting in which to observe the consequences of reduced analyst coverage is around earnings announcements. As informational efficiency declines, we expect more volatile returns around earnings announcements. We also expect greater absolute earnings surprises to the extent

¹⁸ This is the only test in the paper where the subsample of brokerage closure-related terminations generates different results. Specifically, the DiD test shows that variance ratios *decline* around closure-related terminations.

that *post*-termination consensus forecasts reflect reduced information sets.¹⁹

Return volatility in the three days around an earnings announcement increases, suggesting that coverage terminations are associated with more volatile trading around subsequent information events such as earnings releases. This is consistent with West's (1988) prediction of larger price reactions to information events when the information set shrinks. As Table VIII, Panel D shows, the DiD estimate of the change in volatility following a termination is +0.049% ($p=0.002$). This implies an average increase in annualized return volatility around earnings announcements from 75.1% pre-termination to 78.8% post-termination. As for earnings surprises, the DiD estimate of 0.13 indicates that coverage terminations are associated with a 16.9% increase in the average magnitude of earnings surprises ($p=0.003$).

E. Summary

Terminations appear to lead to greater informational asymmetry which increases the risk of trading against better-informed investors and hence reduces liquidity. This will have a relatively greater effect on a retail investor (who typically bases trading decisions on public information) than on an institutional investor (who may have access to other sources of research). This may help explain the different reactions of institutional and retail investors found in Section III.

V. Changes in Expected Returns and Idiosyncratic Volatility

A. Expected Returns

Do coverage terminations affect investors' expected returns? To estimate expected returns, we estimate two standard factor models. The market-and-industry model (Roll (1988)) regresses daily excess returns (r_{jit}) on the CRSP value-weighted index return (r_{mt}) and the firm's value-weighted Fama-French industry return (r_{it}): $r_{jit} = \alpha_{ji} + \beta_{mj}r_{mt} + \beta_{ij}r_{it} + e_{jit}$. The Fama-French three-factor model regresses daily excess returns on r_{mt} and the return difference between small and big stocks

¹⁹ There is no conflict between this hypothesis and the results for earnings surprises shown in Section II. Unlike the exogeneity test in Section II, here we use *post*-termination consensus forecasts and *absolute* earnings surprises.

(SMB_t) and high and low book-to-market stocks (HML_t): $r_{jt} = \alpha_j + \beta_{mj}r_{mt} +$

$\beta_{SMB,j}SMB_t + \beta_{HML,j}HML_t + e_{jt}$. Regressions are estimated over three- or six-month windows

ending 10 days before or starting 10 days after the termination announcement. We report the cross-sectional means of the regression coefficients and the average change relative to the style-matched controls. As before, p -values are calculated using a block bootstrap. Results are robust to standard thin-trading adjustments (Bartholdy and Riding (1994) and Cowan and Sergeant (1996)).

Table IX reports the results. Market beta decreases significantly in both factor models. For instance, in the three-month Fama-French model, average market beta declines from 1.19 to 1.14 around coverage terminations. The change in the control group is approximately zero, resulting in a difference-in-difference estimate of -0.05 ($p < 0.001$). Loadings on the SMB and HML factors, on the other hand, increase significantly (in the six-month window), by 0.04 for SMB ($p = 0.019$) and 0.04 for HML ($p = 0.060$). These findings suggest that dropped stocks become less sensitive to aggregate market movements and begin to behave more like small stocks (which tend to be less informationally efficient) and like value stocks.

Combining the six-month Fama-French difference-in-difference estimates with historical mean factor returns (calculated over the period from 1963 and 2005, available on Kenneth French's website) suggests that annual expected returns increase by seven basis points, to 14.99%, following exogenous coverage terminations. This is commensurate with our point estimates of the announcement-day share price falls shown in Table IV. In the Gordon growth model, the percentage change in share price equals $(r_{before} - g)/(r_{after} - g) - 1$. For $g = 2\%$, 3% , or 4% , the increase in expected returns to 14.99% implies share price falls of 54, 58, or 64 basis points.

B. Volatility Changes

We next turn to volatility changes. A priori, it is unclear whether coverage terminations increase

or decrease volatility. Roll (1988) discusses the relation between firm-specific price fluctuations and information about fundamentals. He favors the view that high volatility reflects active trading by informed arbitrageurs who ensure that stock prices closely reflect fundamental value. In support of this view, Durnev et al. (2003) find that firms and industries with greater firm-specific return variation exhibit higher association between current returns and future earnings, suggesting their share prices are more informative. On the other hand, reduced information production may increase uncertainty and so, possibly, volatility.

We compute two annualized volatility measures: Raw volatility (based on a stock's return in excess of the riskfree rate), and idiosyncratic volatility (using residuals from the factor models). Raw volatility decreases significantly in both estimation windows. Three-month raw volatility decreases by 2.3 percentage points. After subtracting the control-group change, the DiD estimate is minus 1.39 percentage points ($p < 0.001$), reflecting the fact that firm-specific volatilities have generally trended down over the sample period. Over the six-month window, raw volatility falls by 1.85 percentage points. Idiosyncratic volatilities fall by somewhat larger amounts.

C. Volatility Trading Strategy

To illustrate the economic magnitude of the change in volatility around coverage terminations, we compute the profits of a trading strategy that takes a short position in a delta-neutral straddle in stocks dropped from coverage. Short straddles have positive payoffs if volatility falls. Significant returns to this trading strategy would suggest that the option market does not immediately price the termination-related decrease in volatility.

Coval and Shumway (2001) document that short delta-neutral (i.e., zero-beta) straddles generate significant returns net of the leverage effect and transactions costs, which they attribute to a stochastic volatility factor. As a result, it is doubly important that we compare the straddle returns to a benchmark, which we take to be an offsetting (i.e., long) straddle position in options on the

control stocks. Choosing style-matched control companies, rather than, say, options on the S&P 500 index, also helps hold transactions costs constant. Appendix A details the construction of the trading strategy. We simulate trading at the bid and ask but ignore other transactions costs.

Short straddles in firms subject to coverage terminations yield average daily returns of 1.36% ($p < 0.001$), consistent with Coval and Shumway (2001). These returns exceed those for the style-matched controls, which average 1.08% ($p < 0.001$). The average difference of 0.27% per day (6% per month) is statistically significant with a bootstrapped p -value of 0.002.

Short straddles have unlimited downside risk, and the long position in a straddle on control stocks is at best an imperfect hedge. An alternative trading strategy is a butterfly, which combines the short straddle with an upside and downside hedge (details can be found in Appendix B). The short butterfly on terminations yields average daily returns of 2.30% ($p < 0.001$), compared to 2.07% for the butterfly on the control stocks. The average net return is 0.23% per day ($p = 0.018$).

VI. Conclusions

We quantify the value of sell-side equity research by measuring the change in share price and market value of a firm that loses research coverage for exogenous reasons. The exogenous variation we exploit is due to structural changes in the research industry following recent regulatory interventions and adverse market developments which have caused brokerage firms to downsize their research departments and thus drop coverage of thousands of stocks.

We estimate that the capitalized value to a stock-market listed company of receiving sell-side research coverage averages \$8.4 million, or around 110 basis points of market value. The fact that coverage is valuable helps explain why investment banks successfully used promises of coverage when competing for underwriting mandates (Ljungqvist, Marston, and Wilhelm (2007)). The fact that there are diminishing returns to additional research coverage suggests that these promises were particularly compelling when a company starts out without research coverage, such as in an IPO. By

linking coverage to investment banking work, however implicitly, brokers effectively charged companies for the benefit of coverage. The 2003 Global Settlement bans this practice, and not surprisingly banks have scaled down their research operations in its wake.

An alternative way for brokers to monetize the value of research coverage is to charge companies directly, and our estimates suggest that it may be in shareholders' interest to pay for coverage. This model resembles "corporate brokerships" in the U.K. In the U.S., some smaller brokerage firms are experimenting with similar business models, though concerns about the independence of credit rating agencies, which have long charged borrowing companies for their credit ratings, sound a cautionary note.

One reason why analyst research is valuable appears to be that it helps improve informational efficiency in the stock market. If corporate investment is more efficient when stock prices are more informative, as Durnev, Morck, and Yeung (2004) argue, then diminished incentives to produce research may have macroeconomic costs. We aim to explore such real consequences of shifts in the supply of information in follow-on work.

Finally, our evidence suggests that individual investors and small institutions may value research more highly than do large institutional investors. A plausible explanation is that large investors have access to other sources of research (including their own) and so can overcome, or even benefit from, declines in informational efficiency, whereas the cost to small investors of replicating broker-provided research is, presumably, prohibitive. Thus, the recent regulatory interventions aimed at protecting smaller investors may have perverse side-effects, such as increasing trading costs and so reducing smaller investors' stock demands. Whether these side-effects are offset by gains in investor confidence is another open question for future research.

Appendix A: Volatility Trading Strategy - Straddles

The volatility trading strategy is implemented as follows. To avoid trading options subject to event-induced volatility, we establish a short straddle position on trading day +5 after the announcement of a coverage termination. Results are robust to using different start dates. We use options that have positive open interest, remaining time to maturity of between five and 44 trading days, and absolute delta between 0.2 and 0.8 (to ensure we trade at-the-money options). When multiple options satisfy these conditions, we choose the one with $|\Delta|$ closest to 0.5 and maturity closest to one month from the trade date. The mean time to maturity is 21 trading days.

We then take a straddle position consisting of short positions in one put and $w_{call} = \frac{-\Delta_{put}}{\Delta_{call}}$ calls, which ensures that the straddle has zero delta. The initial proceeds from establishing the short straddle are $w_{call} \cdot Premium_{call} + Premium_{put}$. We hold the options to maturity (assuming no interim dividends for simplicity), at which time we incur a payout of $w_{call} \cdot Payoff_{call} + Payoff_{put}$. The continuously compounded daily return on the short straddle is $\frac{1}{TTM} \ln\left(\frac{proceeds}{payout}\right)$, where TTM is time to maturity. The excess return is constructed by subtracting the equivalent straddle return for DGTW (1997) style-matched control stocks, constructed as in Section I.D.

We use OptionMetrics data and evaluate the trading strategy at the bid and ask quotes; investors who can trade inside the bid-ask spread will realize greater profits. The number of useable observations is 8,264.

Appendix B: Volatility Trading Strategy - Butterflies

The butterfly combines a short straddle (constructed as in Appendix A) with a long out-of-the-money put (to hedge extreme downside moves, to which short straddles are exposed) and a long out-of-the-money call (to hedge upside moves in the underlying). We match the hedge component weights to the straddle component weights, i.e., we buy one put and $w_{call} = \frac{-\Delta_{put}}{\Delta_{call}}$ calls. All options are chosen to have the same maturity as the straddle. We choose calls and puts that are three strike prices away from the strike price of the straddle; if these are not available, we use options that are two or one strikes away. This gives variable moneyness, so when constructing the butterfly positions for style-matched stocks, we match the hedge option moneyness as closely as possible to that used for the termination sample stock.

We apply the same data requirements as in Appendix A. Due to the need for out-of-the-money options with open interest, there are fewer tradable butterflies than straddles. The number of useable observations is 2,855.

References

- Agrawal, Anup, and Mark A. Chen, 2004, Analyst conflicts and research quality, Unpublished working paper, University of Alabama.
- Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time series effects, *Journal of Financial Markets* 5, 31-56.
- Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223-249.
- Ashenfelter, Orley, and David Card, 1985, Using the longitudinal structure of earnings to estimate the effect of training programs, *Review of Economics and Statistics* 67, 648-660.
- Baltagi, Badi H., and Ping X. Wu, 1999, Unequally spaced panel data regressions with AR(1) disturbances, *Econometric Theory* 15, 814-823.
- Barber, Brad, Reuven Lehavy, Maureen McNichols, and Brett Trueman, 2001, Can investors profit from the prophets? Security analyst recommendations and stock returns, *Journal of Finance* 56, 531-564.
- Bartholdy, Jan, and Alan Riding, 1994, Thin trading and estimation of betas: The efficacy of alternative techniques, *Journal of Financial Research* 17, 241-255.
- Bhojraj, Sanjeev, Charles M.C. Lee, and Derek Oler, 2003, What's my line? A comparison of industry classification schemes for capital market research, *Journal of Accounting Research* 41, 745-771.
- Boehmer, Ekkehart, James Musumeci, and Annette B. Poulsen, 1991, Event-study methodology under conditions of event-induced variance, *Journal of Financial Economics* 30, 253-272.
- Brown, Stephen, Stephen A. Hillegeist, and Kin Lo, 2004, Conference calls and information asymmetry, *Journal of Accounting and Economics* 37, 343-366.
- Brown, Stephen J., and Jerold B. Warner, 1980, Measuring security price performance, *Journal of Financial Economics* 8, 205-258.
- Carhart, Mark, 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57-82.
- Coval, Joshua D., and Tyler Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983-1009.
- Cowan, Arnold R., 1992, Nonparametric event study tests, *Review of Quantitative Finance and Accounting* 2, 343-358.
- Cowan, Arnold R., and Anne M.A. Sergeant, 1996, Trading frequency and event study test specification, *Journal of Banking and Finance* 20, 1731-1757.

- Cowen, Amanda, Boris Groysberg, and Paul M. Healy, 2006, Which types of analyst firms make more optimistic forecasts? *Journal of Accounting and Economics* 41, 119-146.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristics-based benchmarks, *Journal of Finance* 52, 1035-1058.
- Das, Somnath, Carolyn B. Levine, and K. Sivaramakrishnan, 1998, Earnings predictability and bias in analysts' earnings forecasts, *The Accounting Review* 73, 277-294.
- Demiroglu, Cem, and Michael Ryngaert, 2005, Analyst coverage initiations on previously uncovered stocks, Unpublished working paper, University of Florida.
- Dugar, Amitabh, and Siva Nathan, 1995, The effects of investment banking relationships on financial analysts' earnings forecasts and investment recommendations, *Contemporary Accounting Research* 12, 131-160.
- Durnev, Artyom, Randall Morck, Bernard Yeung, and Paul Zarowin, 2003, Does greater firm-specific return variation mean more or less informed stock pricing?, *Journal of Accounting Research* 41, 797-836.
- Durnev, Artyom, Randall Morck, and Bernard Yeung, 2004, Value-enhancing capital budgeting and firm-specific stock return variation, *Journal of Finance* 59, 65-105.
- Easley, David, Nicholas M. Kiefer, and Maureen O'Hara, 1997, A day in the life of a very common stock, *Review of Financial Studies* 10, 805-835.
- Easley, David, Nicholas M. Kiefer, Maureen O'Hara, and Joseph B. Paperman, 1996, Liquidity, information, and infrequently traded stocks, *Journal of Finance* 51, 1405-1436.
- Francis, Jennifer, and Donna Philbrick, 1993, Analysts' decisions as products of a multi-task environment, *Journal of Accounting Research* 31, 216-230.
- Gleason, Cristi A., and Charles M.C. Lee, 2003, Analyst forecast revisions and market price discovery, *The Accounting Review* 78, 193-225.
- Grossman, Sanford J., and Joseph Stiglitz E., 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393-408.
- Hong, Harrison, and Marcin Kacperczyk, 2007, Competition and bias, Unpublished working paper, Princeton University.
- Irvine, Paul, 2004, Analysts' forecasts and brokerage-firm trading, *The Accounting Review* 79, 125-150.
- Jackson, Andrew, 2005, Trade generation, reputation, and sell-side analysts, *Journal of Finance* 60, 673-717.
- James, Christopher, and Jason Karceski, 2006, Strength of analyst coverage following IPOs, *Journal*

of Financial Economics 82, 1-34.

Jegadeesh, Narasimhan, Joonghyuk Kim, Susan D. Krische, and Charles Lee, 2004, Analyzing the analysts: When do recommendations add value?, *Journal of Finance* 59, 1083-1124.

Kecskes, Ambrus, and Kent Womack, 2007, Adds and drops of analyst coverage: Does the stock market overreact?, Unpublished working paper, Dartmouth College.

Khorana, Ajay, Simona Mola, and P. Raghavendra Rau, 2007, Is there life after loss of analyst coverage?, Unpublished working paper, Purdue University.

Lim, Terence, 2001, Rationality and analysts' forecast bias, *Journal of Finance* 56, 369-385.

Lin, Hsiou-wei, and Maureen F. McNichols, 1998, Underwriting relationships, analysts' earnings forecasts and investment recommendations, *Journal of Accounting and Economics* 25, 101-127.

Ljungqvist, Alexander, Felicia Marston, and William J. Wilhelm, 2007, Scaling the hierarchy: How and why investment banks compete for syndicate co-management appointments, Unpublished working paper, New York University.

Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41-66.

McFadden, Daniel L., 1974, Conditional Logit Analysis of Qualitative Choice Behavior, in *Frontiers in Econometrics*, edited by Paul Zarembka, New York: Academic Press.

McNichols, Maureen, and Patricia C. O'Brien, 1997, Self-selection and analyst coverage, *Journal of Accounting Research* 35, 167-199.

Michaely, Roni, and Kent L. Womack, 1999, Conflict of interest and the credibility of underwriter analyst recommendations, *Review of Financial Studies* 12, 653-686.

Politis, Dimitris N., and Halbert White, 2004, Automatic block-length selection for the dependent bootstrap, *Econometric Reviews* 23, 53-70.

Roll, Richard, 1988, R-squared, *Journal of Finance* 43, 541-566.

Scherbina, Anna, 2007, Suppressed negative information and future underperformance, Unpublished working paper, Harvard Business School.

West, Kenneth D., 1988, Dividend innovations and stock price volatility, *Econometrica* 56, 37-61.

Womack, Kent, 1996, Do brokerage analysts' recommendations have investment value?, *Journal of Finance* 51, 137-167.

Xu, Li, 2006, Institutional investing activities and firms' information environments before and after sell-side analyst coverage initiation and termination, Unpublished working paper, Duke University.