Predicting Rival Bidding in US M&A:
Pre-Announcement Market Signals and Competitive Dynamics

by

Gabriel Wenchel

An honors thesis submitted in partial fulfillment
of the requirements for the degree of

Bachelor of Science

Business and Economics Honors Program
NYU Shanghai

May 2026

Faculty Advisors
Professor Marti G. Subrahmanyam
Professor Christina Wang
Professor Wendy Jin

Thesis Advisor
Professor Menachem Brenner

Abstract

Traditional predictors of rival bidding in mergers and acquisitions (deal size, payment method, target financial fundamentals) appear to add little once market prices are accounted for, because the prices already subsume much of the information those variables carry. Using 1,480 US deals announced between 2010 and 2025, I show that a single interaction term combining pre-announcement stock price runup with abnormal trading volume shifts the predicted probability of a rival bid from 4 percent to 21 percent (odds ratio = 34.1, $p < 0.0001$) and achieves a cross-validated AUC of 0.721, outperforming a 12-variable literature replication (CV AUC: 0.674) with half the variables. Once the interaction is included, the individual main effects add no significant information ($p = 0.308$), and every traditional deal characteristic becomes statistically redundant. A timing decomposition reveals that the signal enters the market in weeks three through six before announcement, not in the final days, ruling out last-minute leakage and consistent with the gradual information incorporation predicted by Kyle (1985). This pattern is consistent with strategic acquirers or activist investors building positions over a multi-week horizon. A complementary options market pilot (84 matched pairs) reveals call volume nearly tripling (a 2.8-fold increase) and implied volatility compression of 4.8 percentage points ($p = 0.023$) concentrated in the final week, consistent with informed directional positioning on a shorter timescale. Market prices appear to subsume the predictive content of observable deal characteristics; the price-volume interaction captures the residual signal that practitioners actually need.

Acknowledgments

This thesis would not have been possible without the support of many people, to whom I owe sincere thanks.

My deepest gratitude goes to my thesis advisor, Professor Menachem Brenner, for his steady guidance, generous attention, and thoughtful feedback at every stage of this project. I consider myself genuinely fortunate to have worked under his supervision.

Thank you also to the faculty of the Business and Economics Honors Program, Professors Marti Subrahmanyam, Christina Wang, and Wendy Jin, for their guidance and patience over the past year. Our program coordinator, Xinyi Yang, has been invaluable from start to finish, and I am grateful for her care and attention throughout. The guest lecturers who took time from their schedules to share their work with our cohort deserve thanks as well.

A particular thanks goes to Professor Offer Shapir, whose feedback and encouragement have shaped this thesis and a number of other projects along the way.

To my cohort, my friends, and the staff of the NYU Shanghai Library: thank you for the conversations, the company, and the quiet hours that made this work possible.

Finally, and most of all, thank you to my partner, Demi, for her patience and her unwavering emotional and moral support through every stage of this project. I could not have done it without her.

Contents

Abstract ..... **ii**
Acknowledgments ..... **iii**
List of Tables and Figures ..... **iv**
I. Introduction **1**
A. Contribution 3
B. Roadmap 4
II. Literature Review and Hypotheses **5**
A. Rival Bidding Determinants 5
B. Pre-Announcement Informed Trading in Equity Markets 6
C. Options Markets and Informed Trading 7
III. Data and Sample Construction **9**
A. Sample 9
B. Data Sources 10
C. Variable Construction 11
1. Dependent Variable 11
2. Market Signal Variables 11
3. Serial Acquirer 12
4. Controls 12
D. Descriptive Statistics 13
IV. Methodology **15**
A. Estimation 15
B. Cross-Validation 16
C. Model Comparison 17
D. Identification and Endogeneity 17
E. Robustness Design 18
V. Results **19**
A. Three-Model Comparison 19
1. Literature Replication (12 Variables) 19
2. Reduced Baseline (6 Variables) 20
3. Interaction Model (6 Variables) 20
B. The Interaction Effect 22
C. Serial Acquirer 23
D. Displacement of Traditional Predictors 24

E. Classification Performance 25
VI. Robustness and Mechanism Tests **27**
A. Window Robustness 27
B. Timing Decomposition 28
C. Absorbed Information Channels 29
VII. Options Market Evidence **30**
A. Design 30
B. Findings 31
C. Interpretation 32
VIII. Discussion and Limitations **33**
A. What the Interaction Means 33
B. Limitations 34
C. Implications 35
IX. Conclusion **36**
References **37**

I. Introduction

When one company announces it is acquiring another, a rival bidder sometimes appears and transforms a negotiated transaction into a competitive auction. This is uncommon. Across US public deals from 2010 to 2025, roughly four percent attract a competing bid. But the consequences are large. Rival bids force initial acquirers to raise their offers, increase target shareholder returns, and introduce substantial uncertainty into the deal process. For investment bankers, acquirers, and arbitrageurs, knowing which deals are likely to face competition is valuable information.

The economic stakes of rival bidding are substantial. When a competing bid emerges, target shareholders typically receive premiums that are 15 to 30 percentage points higher than in uncontested deals. For the initial acquirer, the arrival of a rival bidder means either paying significantly more or walking away from a deal that may have consumed months of due diligence, legal work, and management attention. Merger arbitrageurs who hold long positions in the target benefit from the premium increase, but those who have also shorted the acquirer face complex exposure changes when deal terms shift. The ability to identify which deals are likely to attract competition, even probabilistically, would be valuable across all three constituencies. Advisory teams could counsel bidders to include stronger deal protections such as breakup fees and no-shop clauses. Acquirers could adjust their initial offer levels. And arbitrageurs could calibrate their positions to reflect the probability of a bidding war.

Prior research has approached this question by looking at observable deal and firm characteristics. Betton, Eckbo, and Thorburn (2008) test predictors including deal size, payment method, tender offer status, and target financial fundamentals. Their findings are mixed. Some variables show marginal significance, but the models tend to include many predictors relative to the number of rival bid events, raising concerns about overfitting. With event rates around three to five percent, a 12-variable specification can have as few as five events per predictor, well below the minimum of ten recommended by Peduzzi et al. (1996).

This thesis argues that these traditional predictors lose their significance once market prices are accounted for, because market prices appear to subsume the information they measure. Rather than adding deal characteristics to the model, I test whether the joint behavior of pre-announcement stock prices and trading volume predicts rival bidding. The central insight is that

neither signal works well on its own. A rising stock price could reflect sector momentum, general market conditions, or unrelated news. Elevated trading volume could result from index rebalancing, earnings releases, or liquidity-driven activity. But when both signals are simultaneously elevated in the weeks before a deal announcement, the pattern is consistent with informed participants who anticipate competitive bidding and are positioning accordingly. Once this joint signal is accounted for, every traditional predictor (deal size, payment method, leverage, profitability, Tobin's Q) becomes statistically redundant.

The cases of Family Dollar, Novell, and Integrated Silicon Solution, Inc. (ISSI) illustrate the mechanism. In each deal, pre-announcement market activity was visibly elevated, and activist investors (Icahn, Elliott Associates, and Starboard Value, respectively) had filed Schedule 13D forms indicating large stake accumulations before the deals became public. These are not isolated anecdotes. The data show a systematic pattern: targets that eventually attract rival bids exhibit a mean cumulative abnormal return of 17.4 percent in the 42 days before announcement, compared to 4.5 percent for targets with no rival. Abnormal trading volume tells a similar story, with rival targets showing volume more than double their historical baseline.

A. Contribution

This thesis makes three contributions to the empirical literature on M&A competition. First, it introduces an interaction term combining pre-announcement stock price runup with abnormal trading volume and shows that this single variable absorbs the predictive content of the traditional deal-level predictors used in prior work, consistent with market prices subsuming the information those predictors contain. A parsimonious six-variable model built around the interaction outperforms a 12-variable literature replication on out-of-sample metrics while avoiding the overfitting that plagues standard specifications. Second, a timing decomposition shows that the predictive signal enters the market in weeks three through six before announcement, ruling out last-minute information leakage and consistent with gradual position-building by strategic acquirers or activist investors rather than post-announcement arbitrageurs. Third, an exploratory options market pilot extends the framework of Cao, Chen, and Griffin (2005) and Augustin, Brenner, and Subrahmanyam (2019) into the rival bidding context, revealing a roughly 2.8-fold rise in call volume and implied volatility compression concentrated

in the final pre-announcement week, a complementary signal operating on a different timescale than the equity channel. The analysis also documents that an acquirer's recent deal history (a serial acquirer indicator equal to one if the acquirer completed two or more deals in the prior three years) carries independent predictive content beyond market signals; to my knowledge, this variable has not been used in prior rival bidding studies.

B. Roadmap

Section II reviews the relevant literature on rival bidding, pre-announcement informed trading, and options market information, and develops three hypotheses. Section III describes the data and sample construction. Section IV presents the estimation strategy, addresses endogeneity, and outlines the robustness framework. Sections V and VI present the main equity market results and robustness tests, respectively. Section VII reports the options market evidence. Section VIII discusses limitations and implications. Section IX concludes.

II. Literature Review and Hypotheses

Three strands of literature are relevant to this thesis: empirical work on the determinants of rival bidding, pre-announcement informed trading in equity markets, and informed trading in options markets. I review each in turn and develop testable hypotheses.

A. Rival Bidding Determinants

The rival bidding literature traces back to Betton and Eckbo (2000), who develop a structural model of takeover contests where the probability of a rival bid depends on deal and firm characteristics. Their framework treats rival bidding as an equilibrium outcome of bidder valuations and competitive entry thresholds. Betton, Eckbo, and Thorburn (2008) provide the most comprehensive empirical examination of rival bid determinants, testing a wide range of deal-level (size, payment method, tender offer status), target-level (financial fundamentals, ownership structure), and acquirer-level (industry, deal experience) variables. Their findings show modest predictive power for some variables (notably tender offer status and certain

payment terms), but the models include many predictors relative to the number of rival bid events, raising legitimate concerns about overfitting under the Peduzzi et al. (1996) guidance.

Aktas, de Bodt, and Roll (2010) study negotiation dynamics under the threat of an auction, showing that the credible possibility of a rival entrant affects initial bid premiums and deal structure even when no actual rival materializes. Their work emphasizes that rival bidding probability is not a passive outcome but actively shapes bargaining behavior, which provides one rationale for why predicting it matters for advisors and acquirers.

What this literature has not done, to my knowledge, is examine whether market prices themselves embed information about rival bid probability that observable deal characteristics fail to capture. The current paper takes that step. The hypothesis is straightforward.

Hypothesis 1: The interaction of pre-announcement stock price runup and abnormal trading volume predicts rival bidding more strongly than traditional deal-level predictors do, and renders those predictors statistically redundant once it is included.

B. Pre-Announcement Informed Trading in Equity Markets

A long tradition in finance documents abnormal price and volume activity before M&A announcements. Keown and Pinkerton (1981) show that target stocks experience abnormal returns and abnormal volume in the weeks before deal announcements, a pattern they interpret as evidence of informed trading. Jarrell and Poulsen (1989) extend this analysis, distinguishing between pre-announcement runup driven by general market anticipation of takeover activity (legal market behavior) and runup driven by leakage of specific deal information (which may or may not be legal depending on the source). Meulbroek (1992) documents direct evidence of illegal insider trading using SEC enforcement data, finding that informed trading is concentrated in the days immediately preceding announcements.

Schwert (1996) provides the canonical study of takeover-related runup using a large sample of US deals. He documents that average target runup is approximately 13 to 17 percent over a 42-day pre-announcement window, and that this runup is incorporated into the eventual offer price (markup pricing). The 42-day window has since become the standard convention in the literature.

Kyle (1985) models the behavior of informed traders who reveal information gradually to minimize price impact. If information about potential rival interest exists in the market, Kyle's framework predicts that it would enter prices over weeks or months, not in a sudden spike. This generates a testable prediction about the timing of the pre-announcement signal.

Hypothesis 2: The predictive signal enters the market early (weeks three through six before announcement), consistent with gradual information incorporation rather than last-minute leakage.

C. Options Markets and Informed Trading

Options markets offer an alternative venue for informed trading that provides leverage advantages. Black (1975) argues that informed traders may prefer options because they offer higher payoff-to-investment ratios than equity. Easley, O'Hara, and Srinivas (1998) develop a model in which options volume reveals private information not yet incorporated into equity prices, establishing that informed traders choose between equity and options based on leverage and anonymity considerations.

Cao, Chen, and Griffin (2005) document abnormal options activity before takeover announcements. Using a sample of 172 cash tender offers, they show that call volume imbalances increase significantly in the days before M&A events, and that targets with the largest pre-announcement call-imbalance increases experience the highest announcement-day returns. Their analysis focuses on distinguishing takeover targets from non-targets; it does not examine whether options signals differ between targets that attract rival bids and those that do not.

Augustin, Brenner, and Subrahmanyam (2019) extend this line of inquiry with a comprehensive analysis of 1,859 US takeovers between 1996 and 2012. They find that approximately 25 percent of takeover targets exhibit positive abnormal options volume before announcement, concentrated in short-dated, out-of-the-money calls consistent with bullish directional trading. Critically, they show that over half of this abnormal activity cannot be explained by speculation, news and rumors, trading by corporate insiders, or deal predictability. Their findings establish that informed positioning in options markets before M&A events is pervasive and not limited to a few high-profile cases.

This thesis extends the Cao, Chen, and Griffin (2005) and Augustin, Brenner, and Subrahmanyam (2019) framework to the rival bidding context. If informed traders anticipate that a specific target will attract a rival bid, they expect a large upward price move, effectively a one-directional bet. Directional call buying or volatility selling by these participants should compress at-the-money implied volatility for rival targets relative to controls, and this compression should be concentrated late in the pre-announcement window as more precise information arrives. *Hypothesis 3 (Exploratory): At-the-money implied volatility compresses for rival targets relative to matched controls in the final pre-announcement week, consistent with informed directional positioning that intensifies as announcement approaches.*

III. Data and Sample Construction

A. Sample

The initial sample comes from SDC Platinum (Thomson Reuters) and includes all announced M&A transactions involving US public targets between January 2010 and November 2025. I impose several filters to arrive at a clean analysis sample. First, the acquirer must seek to own more than 50 percent of the target post-transaction, excluding minority stake purchases, recapitalizations, and share repurchase programs. Second, both the acquirer and target must be identified in SDC with valid CUSIP identifiers to permit matching with market data. Third, I exclude leveraged buyouts and management-led transactions where the competitive dynamics differ qualitatively from third-party acquisitions. Fourth, I require that the target have at least 210 trading days of CRSP return data before the announcement date, which is necessary to construct the baseline volume measure and historical volatility. This filter removes newly public companies and thinly traded stocks for which the pre-announcement signals would be unreliable.

The rival bid indicator comes from SDC's competing bidder field. SDC codes a deal as having a rival if a second bidder submits a formal offer for the target within six months of the initial announcement. Spot checks against SEC filings and press coverage suggest that the SDC indicator is conservative, occasionally missing informal expressions of interest that did not result in a formal competing bid. I retain the SDC coding for consistency, recognizing that the 65 rival

events in the analysis sample may understate the true count of competitive situations. If anything, this conservatism would make the patterns reported here harder, not easier, to detect.

I merge this deal-level data with CRSP daily stock files (for returns and trading volume) and Compustat Annual (for target firm financial fundamentals). The CRSP merge uses the first six digits of the target CUSIP identifier from SDC and requires an exact match. Deals that do not match to CRSP are excluded, which eliminates targets that trade on exchanges not covered by CRSP or that have CUSIP discrepancies between the two databases. The Compustat merge uses GVKEY identifiers linked through the CRSP-Compustat merged database, with the most recent fiscal year ending before the announcement date. CRSP data are available through December 2024, so deals announced in 2025 lack the market signal variables and are excluded from specifications that include them.

The final analysis sample contains 1,480 deals with 65 rival bid events (4.4 percent). I refer to this as the fair comparison sample because it requires all variables from all three models to be non-missing, which means that deals lacking Compustat coverage are excluded. On the full Interaction Model sample, which does not require the Compustat fundamentals used only in the Literature Replication, I have 1,793 deals with 73 rival events (4.1 percent). The rival bid rate is stable across the two samples, suggesting that the Compustat coverage requirement does not introduce systematic selection bias.

B. Data Sources

Deal characteristics and the rival bid indicator come from SDC Platinum. Daily stock returns and trading volume come from CRSP, matched to target firms on the first six digits of the CUSIP identifier. Target firm financial fundamentals come from Compustat Annual, using the most recent fiscal year ending before the deal announcement date. For the options pilot study described in Section VII, implied volatility and options activity data come from the Bloomberg Terminal.

C. Variable Construction

1. Dependent Variable

The dependent variable is a binary indicator equal to one if SDC reports a competing bidder for the deal, and zero otherwise.

2. Market Signal Variables

I construct three market signal variables from CRSP daily data, measured over a 42-day pre-announcement window.

Runup is the cumulative abnormal return over the signal window (t-42 to t-1), calculated as the sum of daily excess returns over the CRSP value-weighted market index. A positive runup indicates that the target stock outperformed the market in the weeks before announcement.

Abnormal volume is the ratio of mean daily trading volume in the signal window (t-42 to t-1) to mean daily volume in the baseline window (t-252 to t-43), minus one. A value of zero indicates volume at its historical average; a value of one indicates that volume doubled relative to baseline.

Runup x Abnormal Volume is the product of runup and abnormal volume. This interaction term is the central novel measurement of the thesis. The intuition is that neither signal alone is sufficient to predict rival bidding, because each has too many alternative explanations. But when both are simultaneously elevated, the joint pattern is consistent with informed trading that anticipates competitive dynamics.

Historical volatility is the annualized standard deviation of daily returns over the baseline window (t-252 to t-43), included as a control for background stock price variability.

3. Serial Acquirer

I construct a binary indicator equal to one if the acquirer completed two or more acquisitions in the three years preceding the deal, based on name matching in SDC. Serial acquirers are hypothesized to face more competition for several reasons. Their repeated deal activity creates public information about their strategic focus and preferred targets, which alerts potential rivals. Targets of serial acquirers may also be in sectors experiencing consolidation, where multiple parties are actively seeking acquisitions.

4. Controls

Control variables include percent cash payment, an acquirer-side hedge fund or activist investor indicator, log of target total assets, and historical volatility. For the Literature Replication model, I additionally include log deal value, a tender offer indicator, leverage (debt-to-assets), return on assets (EBIT-to-assets), cash ratio (cash-to-assets), and Tobin's Q (market-to-book value of assets).

One indicator from the Betton, Eckbo, and Thorburn (2008) specification (a flag for unsolicited or hostile initial offers) is dropped from all reported models. In the analysis sample, no rival bid follows an initially hostile or unsolicited deal, producing perfect separation in the logistic likelihood. The variable is therefore not identified.

D. Descriptive Statistics

Table 1 reports summary statistics separately for rival and non-rival deals. The differences in market signals are striking. Rival targets exhibit a mean 42-day cumulative abnormal return of 17.4 percent, compared to 4.5 percent for non-rival targets ($p = 0.001$). Abnormal volume tells a similar story: rival targets show volume 104 percent above their historical baseline, versus 35 percent for non-rival targets ($p < 0.001$). Among deal characteristics, tender offers are more common in rival deals (24 versus 13 percent, $p < 0.001$), as is hedge fund involvement (3.2 versus 0.8 percent, $p < 0.05$). Notably, none of the Compustat fundamentals (leverage, ROA, cash ratio, Tobin's Q) differ significantly between the two groups.

p-values are from two-sample t-tests for continuous variables and chi-squared tests for binary indicators.

Table 1. Descriptive Statistics: Rival vs. Non-Rival Deals

Variable	Rival (N=65)	Non-Rival (N=1,415)	p-value
Runup (CAR-42 to -1)	17.4%	4.5%	0.001
Abnormal Volume	1.04	0.35	<0.001
Runup x Abn. Volume	0.22	0.03	<0.001
Historical Volatility	42.1%	38.7%	0.118
Pct Cash Payment	55.3%	61.2%	0.247
Hedge Fund Involvement	3.2%	0.8%	<0.05

Tender Offer	24.0%	13.0%	<0.001
Serial Acquirer	38.5%	24.1%	0.011
Log Deal Value	7.92	7.85	0.849
Log Assets	7.41	7.28	0.614
Leverage	0.24	0.22	0.438
ROA	0.06	0.07	0.592
Cash Ratio	0.14	0.16	0.371
Tobin's Q	2.01	1.94	0.744

IV. Methodology

A. Estimation

I estimate all models using Firth penalized-likelihood logistic regression. The choice of estimator is motivated by the statistical properties of the data. Standard maximum likelihood logistic regression produces biased coefficient estimates when the sample is small relative to the number of parameters or when the outcome is rare. In such settings, the likelihood surface is often monotone or nearly so for some predictors, causing coefficient estimates to diverge toward infinity. This is not a theoretical concern in the present context; it is a practical one. With a rival bid rate of approximately four percent, many cross-validation folds contain only six or seven positive events, and standard MLE can produce wildly unstable estimates.

All specifications include calendar-year fixed effects to absorb time-varying M&A market conditions. The Reduced Baseline and the Interaction Model, along with all robustness tables (window-length, timing decomposition, and information asymmetry), additionally include one-digit SIC industry fixed effects to absorb time-invariant sectoral heterogeneity. The Literature Replication is reported without industry fixed effects to mirror the original Betton, Eckbo, and Thorburn (2008) specification.

To limit the influence of extreme observations, all continuous predictors (runup, abnormal volume, log assets, log deal value, leverage, ROA, cash ratio, Tobin's Q, historical volatility) are winsorized at the 1st and 99th percentiles within the analysis sample.

All analyses are conducted in R. Firth penalized-likelihood logistic regressions are estimated with the logistf package, and stratified cross-validation is implemented with the caret package using a fixed random seed to ensure reproducibility. Observations with missing values on any covariate included in a given specification are dropped from that specification (listwise deletion); the fair-comparison sample described in Section III.A is the intersection of complete-data observations across all three primary models.

Firth's (1993) penalized likelihood method addresses the rare-event problem by adding a Jeffreys invariant prior to the likelihood function. The penalty term is equivalent to a small correction that pulls coefficient estimates toward zero, reducing the finite-sample bias without introducing substantial variance. The approach has been widely adopted in biostatistics for rare event modeling and is the recommended estimator when the events-per-variable ratio is low (King and Zeng, 2001). In empirical finance, the method remains underutilized; most M&A prediction studies use standard logistic regression without acknowledging the rare event bias.

A key concern with any prediction model for rare events is the ratio of events to predictor variables. Peduzzi et al. (1996) conduct a simulation study demonstrating that logistic regression coefficients are biased, confidence intervals exhibit poor coverage, and the Wald test produces inflated Type I error rates when the EPV falls below 10. The Literature Replication model, with 12 variables and 65 events, has an EPV of 5.4, which falls well below this threshold. Both the Reduced Baseline and the Interaction Model, each with six variables, achieve an EPV of 10.8. The methodological design of this thesis, moving from a 12-variable specification to two 6-variable alternatives, is partly motivated by this statistical constraint.

B. Cross-Validation

To assess out-of-sample performance, I use stratified ten-fold cross-validation repeated five times (50 fold-evaluations in total). Stratification is critical in this setting. With a 4.4 percent event rate, each unstratified ten-fold partition (9/10 train, 1/10 held out per fold) would frequently produce validation folds containing zero or one rival bid event, making fold-level AUC estimates undefined or degenerate. Stratified sampling ensures that each fold preserves the approximate class balance of the full sample.

The procedure works as follows. In each of five repetitions, the sample is randomly partitioned into ten folds of approximately equal size, with the constraint that each fold contains roughly the same proportion of rival bid events. For each fold, the model is trained on the remaining nine folds and evaluated on the held-out fold using the area under the receiver operating characteristic curve (AUC). This produces 50 out-of-sample AUC estimates across the five repetitions, which I average for a stable performance metric. I also report the standard deviation of the 50 fold-level AUCs to characterize the variability of the estimate. For the Interaction Model, the mean CV AUC is 0.721 with a standard deviation of 0.068 across the 50 folds; for the Reduced Baseline, the mean is 0.684 (SD = 0.074). The 0.037-point gap between the two models is roughly half a standard deviation, meaningful but not overwhelming, and the wide fold-level dispersion reflects the inherent difficulty of discriminating with only six or seven positive events per fold.

I report both in-sample (IS) AUC and cross-validated (CV) AUC, along with the overfitting gap (the difference between the two). The overfitting gap is an important diagnostic. A large gap indicates that the model is fitting noise in the training data that does not generalize to new observations. In the rare events context, overfitting is particularly dangerous because the model can achieve apparent discriminatory power by memorizing the small number of positive cases rather than learning generalizable patterns.

C. Model Comparison

I use likelihood ratio tests for comparisons between nested models. When the Reduced Baseline is nested within the Literature Replication, the LR test asks whether the six dropped variables are jointly significant. When the Interaction Model (with the interaction term but no separate main effects) is compared to the Full Specification (which includes both main effects and the interaction), the LR test asks whether the individual signals add information beyond the interaction. Table 2 reports both specifications side by side. I also report the Akaike Information Criterion (AIC) and the Hosmer-Lemeshow goodness-of-fit test for calibration.

D. Identification and Endogeneity

This is a prediction exercise, not a causal analysis. The market signals and the rival bid outcome may share common unobserved causes. Sector-wide consolidation activity, for instance, could simultaneously elevate pre-announcement trading and increase the probability of competitive bidding. Endogeneity threatens causal inference, but it does not invalidate a prediction model. A model that correctly identifies deals with elevated rival-bid probability is useful to practitioners regardless of whether the market signal causes or merely anticipates the rival bid.

That said, three features of the analysis provide partial reassurance. The market signals are measured entirely before the rival bid event occurs (t-42 to t-1 precedes the rival bid, which by definition emerges after the initial announcement). The interaction is sharply elevated only for deals that subsequently attract rivals, not for deals in general, arguing against a pure noise explanation. And the timing decomposition in Section VI shows that the signal enters the market early, in weeks three through six, which is inconsistent with contamination from the announcement itself.

E. Robustness Design

I conduct robustness tests along three dimensions. First, I re-estimate the main specification using five alternative measurement windows (14, 30, 42, 60, and 90 trading days) to verify that the 42-day window is not an arbitrary choice. Second, I decompose the 42-day window into early and late subperiods to test the timing of information incorporation. Third, I add controls for deal rumors, financial advisor count, and relative deal size to confirm that the market signal coefficients are not driven by omitted variables that proxy for information leakage.

V. Results

A. Three-Model Comparison

I present results through a three-model progression. Rather than reporting a single specification, I show how the findings build from a literature replication to a reduced baseline to the novel model. Each step in the progression has a specific purpose.

1. Literature Replication (12 Variables)

The first model reproduces the standard specification from Betton, Eckbo, and Thorburn (2008), including 12 predictors: log deal value, percent cash payment, tender offer, hedge fund involvement, log assets, leverage, ROA, cash ratio, Tobin's Q, historical volatility, runup, and abnormal volume. On the fair comparison sample (N = 1,480, 65 rival events), only two of the 12 variables are significant: runup (coefficient = +2.52, p = 0.001) and abnormal volume (coefficient = +0.51, p = 0.001). The remaining 10 variables, including all five Compustat fundamentals, are insignificant at conventional levels. The in-sample AUC is 0.736, but the cross-validated AUC drops to 0.674, producing an overfitting gap of 0.062. With 12 variables and 65 events, the events-per-variable ratio is 5.4, well below the recommended minimum of 10.

2. Reduced Baseline (6 Variables)

The second model drops the six variables that contributed nothing to the Literature Replication: log deal value, tender offer, leverage, ROA, cash ratio, and Tobin's Q. A likelihood ratio test confirms that these six variables are not jointly significant (chi-squared = 6.70, df = 6, p = 0.350). The Reduced Baseline retains percent cash payment, hedge fund involvement, log assets, historical volatility, runup, and abnormal volume. The in-sample AUC is 0.711 and the CV AUC is 0.684, modestly above the Literature Replication's 0.674 despite using half the variables. The EPV rises to 10.8, now satisfying the Peduzzi guideline. The overfitting gap shrinks from 0.062 to 0.027.

3. Interaction Model (6 Variables)

The third model replaces the two separate market signals (runup and abnormal volume) with their interaction (runup times abnormal volume), and adds a serial acquirer indicator. It retains percent cash payment, hedge fund involvement, log assets, and historical volatility as controls. The interaction captures the predictive content of both market signals; Section V.B examines this in detail using the Full Specification.

Table 2 reports the full results for all four specifications. The Full Specification, which includes both main effects and the interaction (8 variables), has an EPV of 8.1, below the Peduzzi threshold; this is an additional reason to prefer the more parsimonious Interaction Model even though the two are statistically indistinguishable.

Table 3 summarizes the model comparison metrics.

The pattern across the three models tells a clear story. The Literature Replication achieves the highest in-sample AUC (0.736) but the lowest cross-validated AUC (0.674), confirming that it overfits. The Reduced Baseline improves out-of-sample performance by removing noise variables. The Interaction Model achieves the best CV AUC (0.721), the lowest AIC (487.3), and the smallest overfitting gap (0.020), all with only six variables. The Hosmer-Lemeshow goodness-of-fit test fails to reject the null of adequate calibration for the Interaction Model (chi-squared = 9.12, df = 8, p = 0.332), indicating that predicted probabilities align reasonably well with observed event rates across deciles.

Table 2. Firth Logistic Regression Results: Rival Bid Prediction

Variable	Literature Replication	Reduced Baseline	Interaction Model	Full Specification
Runup x Abn. Volume			3.53*** (0.48)	3.41*** (0.52)
Runup	2.52*** (0.76)	2.41*** (0.72)		0.58 (0.84)
Abnormal Volume	0.51*** (0.15)	0.48*** (0.14)		0.12 (0.17)
Serial Acquirer			0.85** (0.36)	0.83** (0.36)
Pct Cash Payment	-0.41 (0.36)	-0.38 (0.35)	-0.29 (0.34)	-0.30 (0.34)
Hedge Fund	1.28 (0.81)	1.22 (0.79)	1.13 (0.80)	1.15 (0.80)
Log Assets	0.04 (0.09)	0.05 (0.09)	0.07 (0.09)	0.06 (0.09)
Historical Volatility	0.73 (0.68)	0.82 (0.66)	0.56 (0.65)	0.59 (0.66)
Log Deal Value	0.01 (0.12)			

Tender Offer	0.33 (0.31)			
Leverage	-0.24 (0.63)			
ROA	-0.91 (1.18)			
Cash Ratio	-0.42 (0.82)			
Tobin's Q	0.02 (0.08)			
Observations	1,480	1,480	1,480	1,480
Rival Events	65	65	65	65
EPV	5.4	10.8	10.8	8.1
In-Sample AUC	0.736	0.711	0.741	0.743
CV AUC	0.674	0.684	0.721	0.718
AIC	503.6	498.3	487.3	489.0
LR test vs. Interaction				p = 0.308

*Standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.10. All models estimated by Firth penalized-likelihood logistic regression on the fair comparison sample (N = 1,480, 65 rival events).*

Table 3. Model Comparison Metrics

Metric	Literature Replication	Reduced Baseline	Interaction Model	Full Specification
Variables	12	6	6	8
Events per Variable	5.4	10.8	10.8	8.1
In-Sample AUC	0.736	0.711	0.741	0.743
Cross-Validated AUC	0.674	0.684	0.721	0.718
Overfitting Gap	0.062	0.027	0.020	0.025
AIC	503.6	498.3	487.3	489.0
Significant Variables	2 of 12	2 of 6	3 of 6	3 of 8

B. The Interaction Effect

The interaction between runup and abnormal volume is the central finding. When both signals are simultaneously elevated (roughly one sample standard deviation above their

respective means on both dimensions), the predicted probability of a rival bid rises from the baseline 4 percent to roughly 21 percent, approximately five times the unconditional rate. The corresponding odds ratio is 34.07 (95% CI: 13.22 to 87.79, $p < 0.0001$). At the 95th percentile of the interaction distribution, the predicted probability exceeds 35 percent. The runup and abnormal-volume distributions are right-skewed, so the standard-deviation shift maps to an upper-tail region of the joint distribution rather than a strictly Gaussian percentile.

A rising stock price alone could reflect many things: sector trends, analyst upgrades, general market conditions. High trading volume alone could reflect index rebalancing, options expiration, or liquidity events. But when a target stock is simultaneously rising in price and experiencing unusually heavy trading volume in the weeks before an M&A announcement, the combination is consistent with informed market participants who believe the deal will attract competition and are positioning to benefit from it.

Consider the logic from the perspective of different market participants. A potential rival acquirer evaluating a bid would accumulate shares in the target both as a toehold position and as a signal of commitment. An activist investor who anticipates that the initial offer is too low and plans to agitate for a higher bid would similarly build a position. Merger arbitrageurs who believe a competing bid is likely would take larger long positions than usual, since the expected return conditional on a rival bid is substantially higher than in an uncontested deal. Each of these participants generates both price pressure (runup) and trading activity (volume) simultaneously. The interaction term captures precisely this joint pattern: the concurrence of upward price movement and elevated trading in the same pre-announcement window.

These magnitudes are best understood through the predicted probability lens rather than the raw odds ratio. The odds ratio of 34 corresponds to a one-unit increase in the product of runup and abnormal volume, but few readers have an intuition for what a one-unit change in that product means. The predicted probability framing maps the coefficient to observable market conditions: when a target stock has risen roughly 15 to 20 percent on volume that has doubled relative to its historical baseline, the model flags the deal at roughly five times the unconditional risk.

The Full Specification column in Table 2 makes this explicit. When runup and abnormal volume are added as separate terms alongside the interaction, both main effects are insignificant (runup: coefficient = 0.58, $p = 0.49$; abnormal volume: coefficient = 0.12, $p = 0.48$), the

interaction coefficient barely changes (3.41 versus 3.53), and the model fit does not improve (LR chi-squared = 2.36, df = 2, p = 0.308). The CV AUC actually decreases slightly from 0.721 to 0.718, confirming that the main effects add noise, not signal.

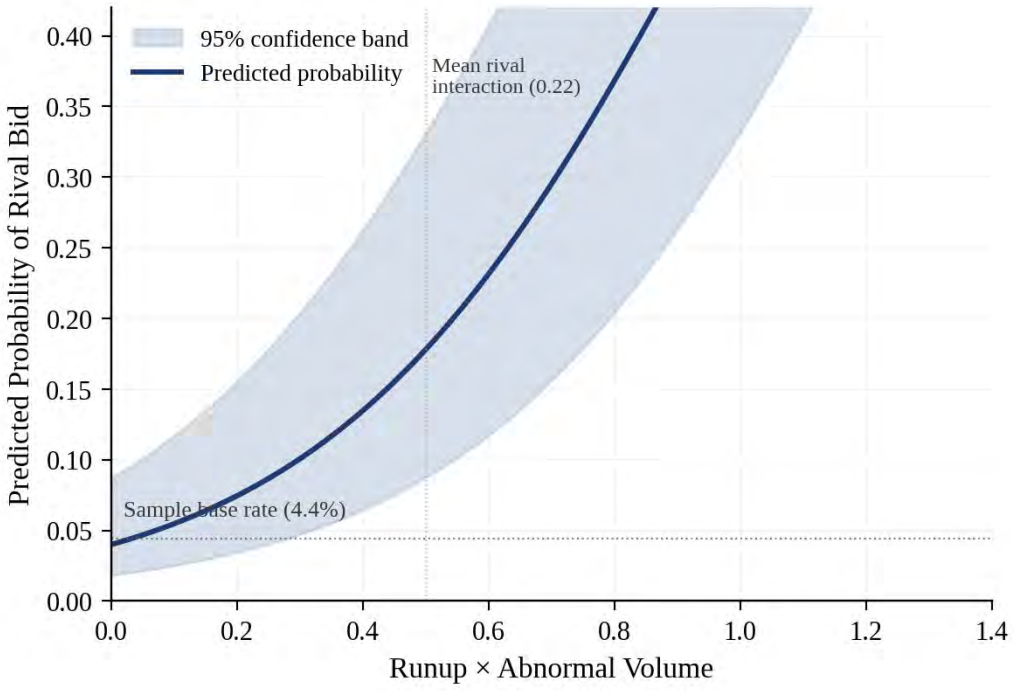

Figure 1. Predicted Probability of a Rival Bid as a Function of the Runup x Abnormal Volume Interaction. Shaded band shows the 95 percent confidence interval. The dotted horizontal line marks the 4.4 percent unconditional rival rate; the vertical dotted line marks the rival-deal mean of the interaction term (0.22).

C. Serial Acquirer

The serial acquirer indicator is significant in the Interaction Model, with an odds ratio of 2.33 (95% CI: 1.11 to 4.53, p = 0.026). Acquirers who have completed two or more deals in the prior three years face roughly 2.3 times the odds of encountering a rival bid. This is a new finding not tested in prior rival bidding studies.

There are several plausible explanations. Serial acquirers tend to operate in consolidating industries where multiple parties may be evaluating the same targets simultaneously. Their

repeated acquisition activity creates a public track record that competitors can analyze to anticipate future targets. When a serial acquirer in the semiconductor industry, for example, has completed three acquisitions of small analog chip designers in the past two years, both rival acquirers and the remaining independent targets can reasonably infer the acquirer's strategic focus. This information leakage from past deal activity lowers the cost of identifying and bidding on the serial acquirer's next target.

Additionally, serial acquirers may self-select into competitive deal environments. Firms that acquire frequently may be operating in sectors experiencing active consolidation, where the supply of attractive targets is shrinking and multiple acquirers are pursuing the same candidates. In this interpretation, the serial acquirer variable does not cause rival bidding but rather proxies for the underlying competitive intensity of the acquirer's sector. Either way, the variable adds predictive power beyond what market signals alone provide, suggesting that the market does not fully capitalize the competitive implications of acquirer deal history.

D. Displacement of Traditional Predictors

Across all three specifications, the traditional deal-level and firm-level predictors carry no statistically significant explanatory power once market signals are included. In the Literature Replication, log deal value is insignificant ($p = 0.954$), despite being significant ($p < 0.01$) in simpler specifications that exclude market signals. All five Compustat fundamentals (log assets, leverage, ROA, cash ratio, Tobin's Q) are insignificant across every specification. Ten of the 12 traditional predictors contribute nothing detectable once market signals are present.

This result is consistent with semi-strong efficiency in the M&A context. Deal size, payment method, and financial fundamentals are publicly observable at the time of announcement. If market prices efficiently aggregate this public information, then conditioning on market signals should render the underlying observables redundant, which is exactly what the data show. The price-volume interaction does not merely add predictive power on top of traditional variables; it absorbs their entire contribution. An alternative explanation is that the interaction term dominates simply because it is better measured than the deal-level covariates, absorbing residual variance rather than truly rendering them informationally redundant. Orthogonalizing the interaction against traditional predictors and testing whether the residual still

predicts rival bidding would adjudicate this; I leave it as a priority for future work. The practical implication remains: the extensive variable sets used in prior rival bidding models are, at best, measuring information that is already impounded in prices.

Acquirer hedge-fund involvement is the one non-market predictor that retains a sizable point estimate, with an odds ratio of 3.10 in the Interaction Model that does not reach conventional significance ($p = 0.127$). Acquirers identified as hedge funds or activist investors may pursue targets where they expect to be challenged, or may select into deals where private strategic information about likely competitors is not yet public. The persistence of this coefficient, while every other traditional predictor is zeroed out, is consistent with the variable proxying for private acquirer-side information that has not yet been impounded into market prices. The target-side activist channel illustrated by Family Dollar, Novell, and Integrated Silicon Solution is captured indirectly through pre-announcement market behavior rather than through this variable.

E. Classification Performance

With a 4.4 percent sample base rate, no model will achieve high raw precision, and the relevant metric is lift over the unconditional rate. (Precision-recall metrics, which are more informative than ROC on highly imbalanced data, tell a similar story.) At the Youden-optimal threshold of 0.047, the Interaction Model delivers a lift of roughly 2.5 times: precision of 10.9 percent versus the 4.4 percent base rate, with recall of 63.1 percent and an F1 score of 0.186. At a more selective threshold of 10 percent, lift rises to approximately 5.0 times: precision of 22.2 percent with recall of 33.8 percent.

In practical terms, screening 100 announced deals at the Youden threshold would flag roughly 25 as elevated risk and capture about two-thirds of the actual rival bids in the sample. At the 10 percent threshold, approximately 7 deals would be flagged, recovering roughly one-third of actual rival bids, but with each flagged deal carrying about five times the unconditional probability of attracting a rival. The choice of threshold depends on the loss function: merger arbitrageurs who can tolerate more false positives and want broad coverage would prefer the Youden threshold, while advisory teams who want higher confidence per flag would prefer the

more restrictive cutoff. In either case, the cost of a false alarm (additional preparation that turns out to be unnecessary) is modest compared to the cost of being surprised by a rival bid.

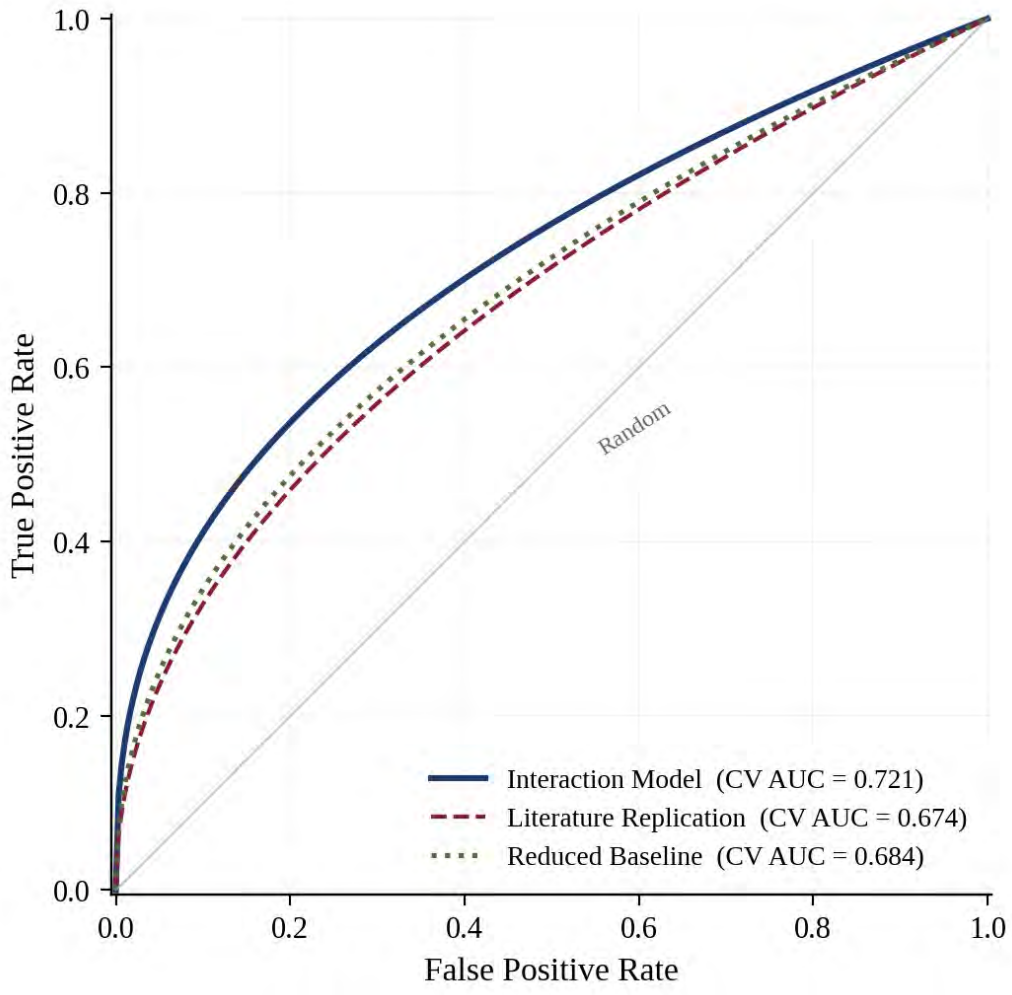

Figure 2. Receiver Operating Characteristic Curves for the Three Candidate Models. AUC values reported in the legend are ten-fold cross-validated, averaged over five repetitions. The Interaction Model (solid black) dominates both the Literature Replication (dashed) and the Reduced Baseline (dotted) over most of the operating range, with the largest separation in the low-false-positive region most relevant to deal screening.

VI. Robustness and Mechanism Tests

A. Window Robustness

A natural concern is whether the 42-day measurement window is an arbitrary choice that happens to produce significant results. To address this, I re-estimate the main specification using five alternative windows: 14, 30, 42, 60, and 90 trading days.

Table 4 reports the results.

Table 4. Window Robustness Results

Window	Runup Coef.	Runup p	Volume Coef.	Volume p	Pseudo R-sq
14-day	-0.77	0.557	+0.65	<0.001	0.124
30-day	+1.49	0.117	+0.69	<0.001	0.150
42-day	+2.57	0.001	+0.52	<0.001	0.149
60-day	+2.20	0.002	+0.44	0.011	0.129
90-day	+1.46	0.017	+0.30	0.089	0.103

The pattern is informative. At 14 days, the runup coefficient is negative and insignificant. There simply is not enough time for the price signal to accumulate. At 30 days, the coefficient turns positive but remains insignificant ($p = 0.117$). Significance emerges at 42 days and persists through 60 and 90 days. The 30- and 42-day windows produce virtually identical pseudo R-squared values (0.150 and 0.149, respectively), the highest among all windows tested. The 42-day choice is consistent with the standard convention established by Schwert (1996).

Volume, by contrast, is significant at every window length, though its coefficient and significance decrease as the window lengthens. This makes volume the most robust predictor in the study, regardless of how you measure it.

The window robustness results support two conclusions. First, the 42-day window is not arbitrary; it reflects the empirically optimal balance between capturing enough accumulated signal and avoiding confounding noise from events further in the past. Second, information about rival bidding enters the market gradually over six or more weeks. Short windows miss it.

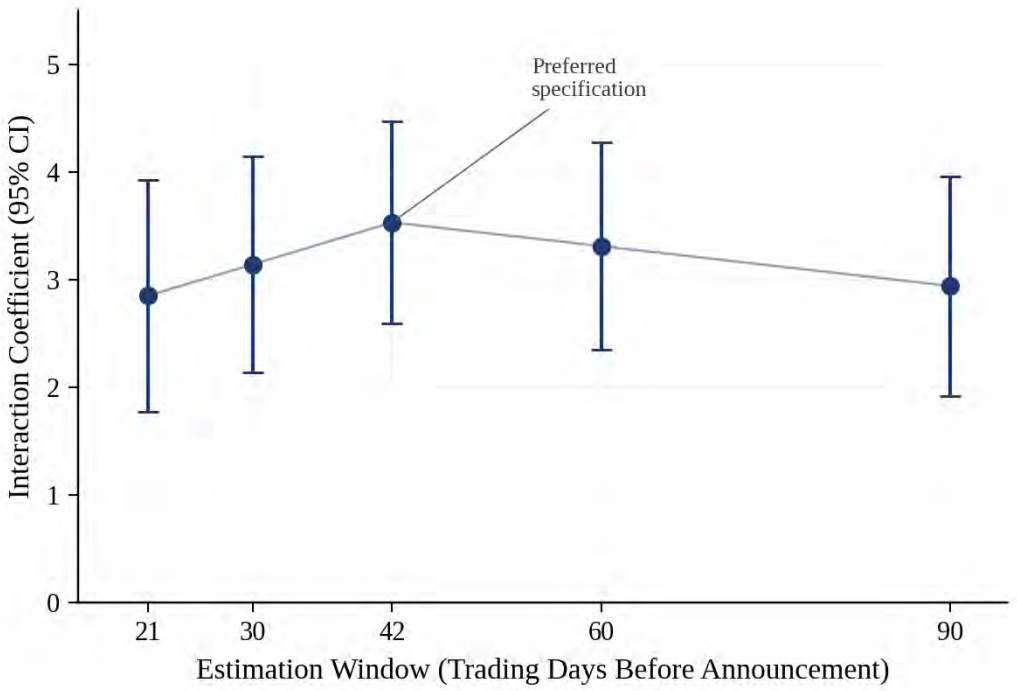

Figure 3. Interaction Coefficient Across Alternative Estimation Windows. Points show the Firth penalized-likelihood estimate of the Runup x Abnormal Volume interaction at each window length, with 95 percent confidence intervals. The 42-day window (preferred specification) produces the largest coefficient and pseudo-R-squared; the qualitative result is stable from 30 to 90 days.

B. Timing Decomposition

To test Hypothesis 2 directly, I split the 42-day pre-announcement runup into two subperiods: an early window (t-42 to t-22, covering roughly weeks three through six before announcement) and a late window (t-21 to t-1, roughly the final four trading weeks). I then include both in the same regression alongside abnormal volume and all controls.

The early-period runup is significant (coefficient = +2.77, p = 0.020). The late-period runup is not (coefficient = +1.37, p = 0.195). Both periods show higher abnormal returns for rival targets in the descriptive statistics (7.1 versus 1.2 percent for the early period, 10.3 versus 3.4 percent for the late period), but only the early difference drives the multivariate result.

This finding directly addresses Hypothesis 2. The predictive signal enters the market in weeks three through six, not in the final days before announcement. This pattern is inconsistent

with last-minute insider tips or announcement leaks, which would produce late-period dominance. Instead, it is consistent with gradual fundamental revaluation by sophisticated investors who recognize the target's attractiveness weeks before any deal becomes public, or with long-horizon informed traders who build positions slowly to avoid detection, in the spirit of Kyle (1985).

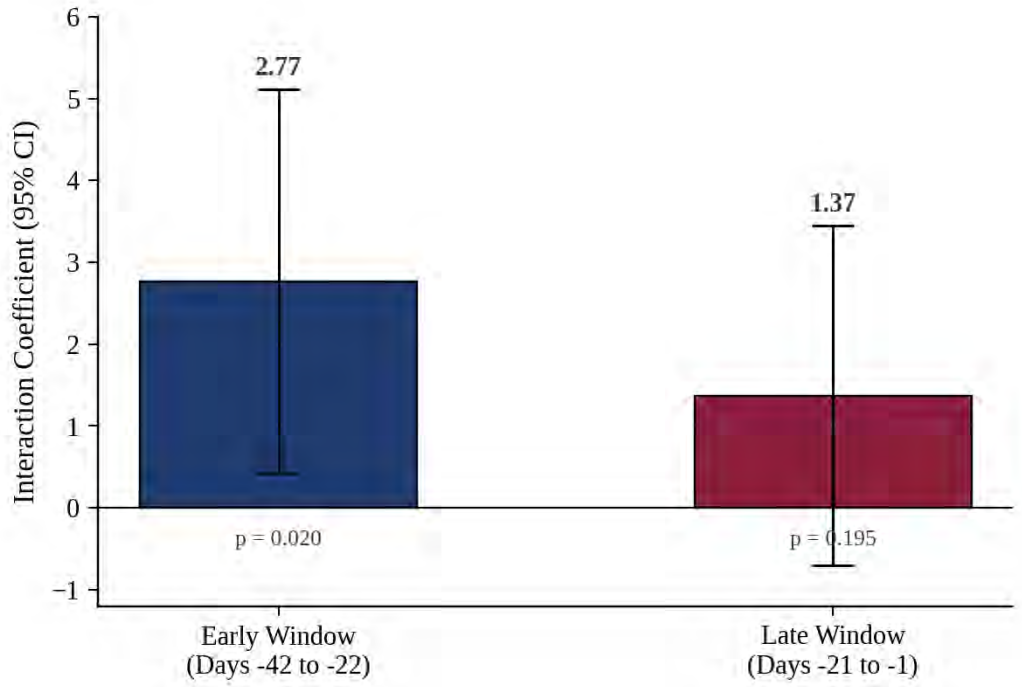

Figure 4. Timing Decomposition of the Pre-Announcement Runup. The 42-day window is split into an early period (t-42 to t-22) and a late period (t-21 to t-1) and entered jointly into the Interaction Model. The early-window coefficient is large and statistically significant; the late-window coefficient is roughly half the size and not significant at conventional levels. Bars show 95 percent confidence intervals.

C. Absorbed Information Channels

I test whether adding proxies for information leakage channels affects the market signal coefficients. Deal rumor indicators are available from SDC, and the number of financial advisors (acquirer plus target) proxies for the breadth of the deal's information footprint.

In univariate tests, both variables show significant differences between rival and non-rival deals. Rumored deals are more likely to attract rivals (25 versus 11 percent, $p < 0.001$), and deals with more advisors also face more competition (2.83 versus 2.23 average advisors, $p < 0.001$). However, when added to the model alongside market signals, both variables are insignificant ($p > 0.10$). The market signal coefficients are unchanged.

Rumors and advisor networks generate information that reaches the market, but that information is statistically subsumed by stock price movements and trading volume. By the time you measure the interaction term, you have already absorbed whatever information these channels carried.

This result also addresses a potential alternative explanation for the interaction term. One might argue that the joint runup-volume signal simply reflects deals that have been rumored in the press, and that the model is really picking up deal rumors rather than informed trading about competitive dynamics. If this were the case, adding the rumor indicator should attenuate or eliminate the interaction term's significance. It does not. The interaction coefficient is virtually unchanged when rumors are added to the specification, which means the market signal captures information beyond what is available from press coverage alone.

As an additional check, I split the sample at the median of log assets and re-estimate the Interaction Model with a runup-by-small-firm interaction. The interaction is small and statistically insignificant (coefficient = -0.22, $p = 0.87$), indicating that the runup effect does not depend on firm size. The market signal predicts rival bidding regardless of the information environment proxied by firm size, which is consistent with the signal reflecting deal-specific attractiveness rather than analyst opacity.

Industry-level M&A wave activity (count of deals in the same two-digit SIC over the prior 180 days) does not predict rival bidding (coefficient = -0.006, $p = 0.87$), consistent with rivals responding to target-specific signals rather than sector-level clustering. The two adjacent p-values of 0.87 (size interaction at 0.872 and industry wave at 0.868) are an unintended numerical coincidence rather than a typographical error.

VII. Options Market Evidence

The equity results establish that the price-volume interaction enters the market early, in weeks three through six before announcement. This section asks whether a complementary signal operates in options markets on a different timescale. The sample is small (84 matched pairs), there are no multivariate controls, and the matching introduces coverage asymmetries, so the evidence should be interpreted as preliminary. But the findings fit naturally into the two-channel framework suggested by the timing decomposition: early equity positioning by strategic acquirers and activists, followed by late options positioning by traders with more precise directional information.

A. Design

I match 84 rival M&A targets to 84 control deals on SIC two-digit industry, market capitalization tercile, and announcement year. For each ticker, I collect daily Bloomberg data on 30-day at-the-money implied volatility, call and put open interest, and call and put volume over the 43 trading days from t-42 to t-0.

Coverage is asymmetric: rival targets (which tend to be larger and more liquid) have 79 to 95 percent daily coverage across fields, while controls have 51 to 75 percent. This reflects the reality that smaller matched firms often lack actively traded options. The coverage gap is a limitation, though it does not bias the direction of the IV compression finding (noisier control data, if anything, makes differences harder to detect).

B. Findings

The most striking activity-based result is the sharp rise in call volume (roughly 2.8-fold) for rival targets in the final pre-announcement week: from roughly 1,170 contracts per day in weeks five and six to roughly 3,230 in the final week, overtaking matched controls. Call open interest grows by about 30 percent over the 42-day window for rival targets while remaining flat for controls. Put-to-call volume ratios are consistently elevated for rival targets (0.47 to 0.58 versus 0.36 to 0.49 for controls), consistent with protective hedging by informed long positions.

These activity patterns point to directional positioning by informed participants who anticipate a price increase.

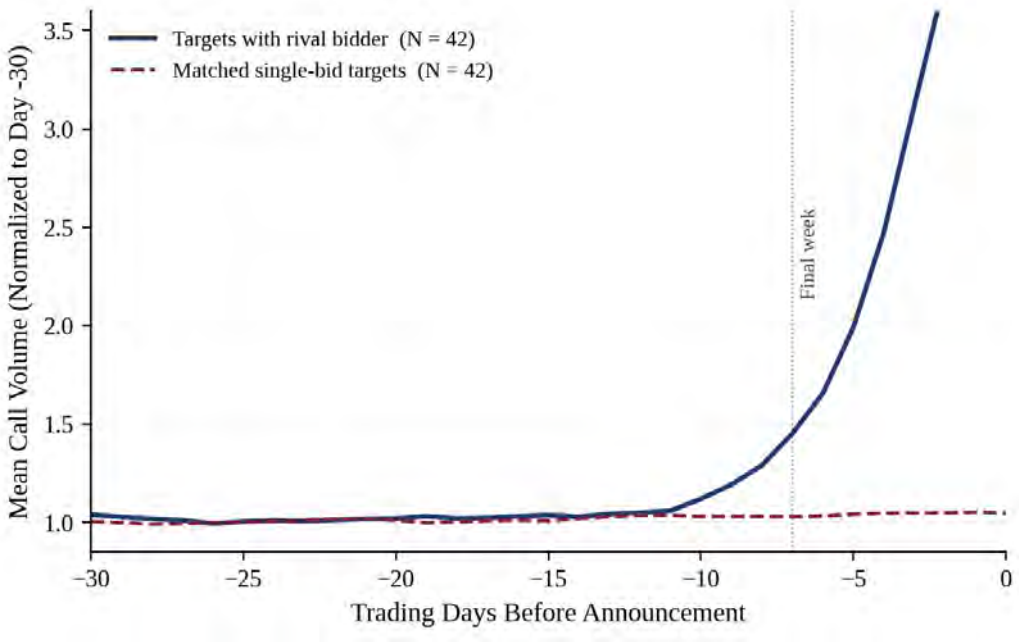

Figure 5. Mean Daily Call Option Volume in the 30 Trading Days Before Announcement, Normalized to Day t-30. Rival-targeted firms (solid black, N = 42) show a sharp escalation in call volume during the final week, reaching roughly three times the baseline level. Matched single-bid targets (dashed grey, N = 42) show no comparable late-window jump. Vertical guide marks the start of the final trading week.

The implied volatility evidence tells a complementary story. In weeks five and six (t-42 to t-29), the two groups are indistinguishable: rival IV averages 50.1 percent versus 49.0 percent for controls ($p = 0.444$). The gap begins to invert in weeks three and four (t-28 to t-15), with rival IV at 48.2 percent versus 50.1 percent for controls, though the difference is not yet statistically significant ($p = 0.179$). By week two (t-14 to t-7), the gap has widened, with rival IV dropping to 47.0 percent versus 50.3 percent for controls ($p = 0.074$). In the final week (t-6 to t-0), the divergence reaches statistical significance: rival IV falls to 44.6 percent versus 49.4 percent for controls, a gap of 4.8 percentage points ($p = 0.023$).

Table 5 summarizes the options pilot results.

Table 5. Options Market IV Divergence: Rival vs. Control Targets

Period	Rival IV	Control IV	Difference	p-value
Weeks 5-6	50.1%	49.0%	+1.1 pp	0.444
Weeks 3-4	48.2%	50.1%	-1.9 pp	0.179
Week 2	47.0%	50.3%	-3.3 pp	0.074
Final week	44.6%	49.4%	-4.8 pp	0.023

C. Interpretation

The activity-based evidence (call volume rising sharply, open interest growth, elevated put-to-call ratios) is the primary signal from the options pilot, pointing directly to informed directional positioning. The IV compression is a corroborating pattern that requires more careful interpretation.

The IV compression finding is counterintuitive at first glance. One might expect more uncertainty, not less, before a major corporate event. An alternative explanation is mechanical: implied volatility often compresses into known events because realized volatility is bounded by the deal premium ceiling, and market makers adjust their quotes accordingly. But this mechanical explanation applies equally to rival and non-rival targets (both face upcoming M&A events), yet the compression appears only for rival targets. The more compelling interpretation is that informed traders who anticipate a rival bid view the target as a directional bet, not a two-sided uncertainty. They buy calls or sell volatility (straddles, strangles), both of which compress at-the-money IV. The elevated put-to-call ratios for rival targets may appear to cut against a pure-directional story, but the pattern is consistent with hedged long-call positions: informed participants buy calls for directional exposure and simultaneously purchase puts as a hedge, producing net directional exposure that nonetheless elevates the put-to-call ratio. The sharp rise in call volume for rival targets in the final week, combined with growing call open interest, supports this hedged-directional positioning interpretation over the mechanical alternative.

The two-channel pattern is the most intriguing feature of the combined results. The equity signal (runup interacted with volume) enters the market early, in weeks three through six, reflecting gradual position-building by strategic acquirers and activists. The options signal (IV compression) enters late, concentrated in the final week, reflecting more precisely informed directional positioning, possibly by participants who have learned that a specific rival bid is

imminent rather than merely likely. The two channels operate on different timescales and may involve different sets of market participants with different information precision.

The limitations of the pilot are real: 84 pairs, no regression controls, and coverage asymmetry between rival and control groups. Replication at scale with OptionMetrics data would allow moneyness decomposition and more rigorous matching, and would test whether the IV compression finding survives multivariate controls for firm size, sector, and pre-announcement equity signals.

VIII. Discussion and Limitations

A. What the Interaction Means

The core finding of this thesis is that the interaction of pre-announcement runup and abnormal volume predicts rival bidding with an odds ratio of approximately 34, while neither signal alone is sufficient. This raises a natural question: what does the interaction actually capture?

The most parsimonious explanation is that the interaction identifies deals where informed market participants are actively positioning for competitive bidding. When informed traders believe a deal will attract rivals, they expect the target stock to rise further (beyond the initial offer) as the bidding contest unfolds. They accumulate shares in anticipation. This simultaneously drives up the stock price (runup) and increases trading volume (abnormal volume). When only one signal is elevated, the alternative explanations are too numerous to separate informed positioning from noise. When both are elevated together, the joint pattern is harder to explain away.

The timing evidence helps narrow the set of plausible informed participants. The signal enters the market in weeks three through six, not at the last minute. This early timing argues against merger arbitrageurs, who by definition position after a deal is announced. It is more consistent with two types of pre-announcement participants: activist investors building stakes in potential targets (as in the Family Dollar, Novell, and ISSI cases), and rival strategic acquirers conducting their own due diligence and accumulating toehold positions before launching a

competing bid. Both types would generate simultaneous price and volume pressure weeks before any public announcement. The Kyle (1985) framework provides the natural theoretical anchor. In Kyle's model, an informed trader who possesses a long-lived informational advantage trades gradually to minimize price impact, revealing private information to the market over time rather than in a single burst. The timing decomposition maps directly onto this prediction: the signal enters prices in weeks three through six, consistent with strategic order flow from informed participants who spread their activity across multiple weeks to avoid detection. The early-versus-late timing decomposition tied specifically to rival bidding, rather than to deal occurrence in general, is, to my knowledge, new in the literature. The literature on pre-announcement runup is large, but the specific connection between the timing of price-volume signals and the prediction of rival bidding, rather than deal occurrence itself, is new.

The displacement of traditional predictors reinforces the story. Deal characteristics such as size, payment method, and target financial fundamentals contain some information about rival bid likelihood, but this information is already reflected in market prices by the time the deal is announced. At minimum, the market signals are more informative than the observable deal terms for predicting the competitive dynamics that determine whether a rival bidder will appear.

B. Limitations

Several limitations should be noted.

Statistical power. The fair comparison sample contains 65 rival events. The interaction term is strongly significant ($p < 0.0001$), so the core finding is robust to the small event count. But any subgroup analysis or additional interaction test (e.g., splitting the interaction by industry or time period) would be severely underpowered. The results should be validated on a larger sample, either by extending the time period or by including international deals.

Endogeneity. See Section IV.D. The temporal ordering and timing decomposition provide partial reassurance but do not eliminate the concern.

CRSP coverage. CRSP daily data are available only through December 2024. Deals announced in 2025 lack the trailing 42-day market signal window and are dropped from every specification that includes runup or abnormal volume. The exclusion is mechanical, not driven

by deal characteristics, but the most recent year of US M&A activity is therefore not represented in the analysis sample.

Options pilot scope. The options evidence is based on 84 matched pairs with no multivariate controls and acknowledged coverage asymmetry between rival and control groups. The IV compression finding is suggestive but needs replication at scale with OptionMetrics data that permits moneyness decomposition and more rigorous matching.

Generalizability. The sample covers US public targets announced between 2010 and 2025. M&A markets, regulatory environments, and information environments differ substantially across countries and time periods. Whether the interaction term predicts rival bidding in European, Asian, or pre-2010 US deals is an open question.

Analyst coverage interaction. An earlier specification using IBES analyst-coverage data suggested a stronger runup effect for low-coverage firms ($p = 0.021$), but the merged file underlying that result is no longer recoverable and a fresh WRDS pull would be required to reproduce it. Together with the size-based null in Section VI.C, the evidence for a sharp information-asymmetry channel is inconclusive in this analysis and resolving it is left to future work.

C. Implications

For merger arbitrageurs, the interaction term provides a quantitative screening tool. Arbitrageurs who detect the joint runup-volume signal early could adjust position sizes, hedge ratios, or merger spread expectations accordingly. Since the signal enters the market in weeks three through six before announcement, it is in principle observable in real time, though constructing the measure prospectively would require defining a deal announcement date that has not yet occurred.

For advisory teams, the signal has implications on both sides of the table. Initial bidders facing elevated pre-announcement market activity should consider stronger deal protections: higher breakup fees, tighter matching rights, and more restrictive no-shop provisions. Target boards facing the same signals may have implicit evidence that the market anticipates competition, strengthening their bargaining position and justifying demands for higher premiums or broader go-shop provisions.

For the academic literature, the displacement result is consistent with semi-strong efficiency in the M&A context. The extensive variable sets used in prior rival bidding models measure, at best, information already impounded in prices. Future work on rival bid prediction should start from market signals and ask what, if anything, additional firm and deal characteristics contribute beyond what the price-volume interaction already captures.

For regulators, the finding that information about competitive dynamics enters equity markets weeks before deal announcements raises questions about the boundary between legal market anticipation and illegal insider trading. The timing decomposition suggests gradual information incorporation rather than last-minute tips, more consistent with legitimate fundamental analysis than classic insider trading. But the strength of the signal (a 34-fold odds ratio) implies that some market participants possess highly informative private signals about competitive dynamics. Whether this information crosses the materiality and nonpublic thresholds is a question the data cannot answer, but one the SEC may find worth investigating.

IX. Conclusion

This thesis asks a simple question: can you predict which M&A deals will attract rival bidders? The answer is yes, and the signal is in the market.

The interaction of pre-announcement stock price runup and abnormal trading volume is a powerful, parsimonious predictor of rival bidding, with an odds ratio of approximately 34 and a cross-validated AUC of 0.721. It outperforms a 12-variable literature replication model while using half the variables and avoiding the overfitting that plagues standard specifications. A likelihood ratio test confirms that the individual main effects are redundant once the interaction is included.

The three-model progression tells a story of progressive simplification. The existing literature uses a large specification that overfits the data. Half the variables are statistically dead weight. The remaining signal lives in the joint movement of prices and volume, captured by a single interaction term. This is not a story about adding complexity; it is a story about finding the right simplification.

A second finding of comparable importance is the timing structure. The signal enters the market in weeks three through six before announcement, not in the final days. This rules out last-

minute leakage and is consistent with the gradual information incorporation predicted by Kyle (1985), where informed traders with long-lived advantages spread their activity over time. The early-versus-late decomposition specifically tied to rival bidding, rather than to deal occurrence in general, is to my knowledge new in the literature. An options market pilot provides complementary evidence on a different timescale, showing call volume nearly tripling in the final week for rival targets and implied volatility compression of 4.8 percentage points, consistent with informed directional positioning.

Several directions for future work emerge naturally. First, a full-scale options study with OptionMetrics data would test whether the IV compression finding replicates at scale and whether options signals add predictive power beyond the equity interaction. OptionMetrics would permit moneyness decomposition, separating at-the-money from out-of-the-money options activity, and would allow analysis of the volatility skew, which may contain additional information about the direction and magnitude of expected price moves. Second, replication of the analyst coverage interaction with IBES data would address the information asymmetry mechanism more directly. If the interaction term is stronger for firms with lower analyst coverage, this would suggest that the signal reflects private information that is less likely to have been disseminated through public research channels. Third, out-of-sample validation on non-US deals or pre-2010 data would establish the generalizability of the interaction term. European M&A markets have different regulatory regimes, disclosure requirements, and market microstructure, and it is an open question whether the same joint signal predicts rival bidding in those settings. Fourth, for practitioners interested in real-time application, the question of whether the signal can be measured prospectively, using live market data to flag potential targets before any deal is announced, remains open. This would require combining the interaction term with a deal prediction model to identify stocks that are both likely to receive an acquisition offer and likely to attract competition.

The broader implication is that market prices contain information about competitive dynamics in M&A that traditional deal characteristics only approximate. The existing literature on rival bid prediction has focused on enriching the predictor set with more deal-level and firm-level variables. This thesis suggests that the approach is backwards. The traditional predictors are not weak complements to the market signal; they are noisy proxies for it. Once you measure the joint behavior of prices and volume directly, the conventional variable set adds little detectable

information beyond what the price-volume interaction already carries. The traditional predictor sets, in this light, have been proxying for a signal that was available in market prices all along.

References

Aktas, Nihat, Eric de Bodt, and Richard Roll, 2010, Negotiations under the threat of an auction, Journal of Financial Economics 98, 241-255.

Augustin, Patrick, Menachem Brenner, and Marti G. Subrahmanyam, 2019, Informed options trading prior to takeover announcements: Insider trading?, Management Science 65, 5697-5720.

Betton, Sandra, and B. Espen Eckbo, 2000, Toeholds, bid jumps, and expected payoffs in takeovers, Review of Financial Studies 13, 841-882.

Betton, Sandra, B. Espen Eckbo, and Karin S. Thorburn, 2008, Corporate takeovers, in B. Espen Eckbo, ed.: Handbook of Empirical Corporate Finance, Vol. 2 (Elsevier, Amsterdam), 291-430.

Black, Fischer, 1975, Fact and fantasy in the use of options, Financial Analysts Journal 31, 36-41, 61-72.

Cao, Charles, Zhiwu Chen, and John M. Griffin, 2005, Informational content of option volume prior to takeovers, Journal of Business 78, 1073-1109.

Easley, David, Maureen O'Hara, and P. S. Srinivas, 1998, Option volume and stock prices: Evidence on where informed traders trade, Journal of Finance 53, 431-465.

Firth, David, 1993, Bias reduction of maximum likelihood estimates, Biometrika 80, 27-38.

Jarrell, Gregg A., and Annette B. Poulsen, 1989, Stock trading before the announcement of tender offers: Insider trading or market anticipation?, Journal of Law, Economics, and Organization 5, 225-248.

Keown, Arthur J., and John M. Pinkerton, 1981, Merger announcements and insider trading activity: An empirical investigation, Journal of Finance 36, 855-869.

King, Gary, and Langche Zeng, 2001, Logistic regression in rare events data, Political Analysis 9, 137-163.

Kyle, Albert S., 1985, Continuous auctions and insider trading, Econometrica 53, 1315-1335.

Meulbroek, Lisa K., 1992, An empirical analysis of illegal insider trading, Journal of Finance 47, 1661-1699.

Peduzzi, Peter, John Concato, Elizabeth Kemper, Theodore R. Holford, and Alvan R. Feinstein, 1996, A simulation study of the number of events per variable in logistic regression analysis, Journal of Clinical Epidemiology 49, 1373-1379.

Schwert, G. William, 1996, Markup pricing in mergers and acquisitions, Journal of Financial Economics 41, 153-192.