Are acquirer stock price reactions to M&A announcements in any way predictable? A machine-learning analysis

João Quariguasi Frota Neto
University of Manchester, Manchester, UK

Konstantinos Bozos
Leeds University Business School, University of Leeds, Leeds, UK

Marie Dutordoir
University of Manchester, Manchester, UK

Konstantinos Nikolopoulos
University of Durham, Durham, UK

ABSTRACT

We examine whether acquirer stock price reactions to M&A deal announcements can be forecasted based on ex ante acquirer, target, deal, and macroeconomic characteristics. We employ machine learning methodologies with out-of-sample testing and standard cross-validation procedures to assess the forecasting accuracy of various parametric and nonparametric models. While overall predictability is low, nonparametric models exhibit some ability to forecast acquirer stock price reactions to M&A announcements, whereas parametric models do not. Feature importance analyses reveal that a handful of predictors, including acquirer size and (relative) deal size, contribute most to the predictions. Our findings have practical implications for corporate managers and various corporate stakeholders.

Are acquirer stock price reactions to M&A announcements in any way predictable? A Machine-Learning analysis.

Anonymous

September 19, 2025

Abstract

We examine whether acquirer stock price reactions to M&A deal announcements can be forecasted based on ex ante acquirer, target, deal, and macroeconomic characteristics. We employ machine learning methodologies with out-of-sample testing and standard cross-validation procedures to assess the forecasting accuracy of various parametric and nonparametric models. While overall predictability is low, nonparametric models exhibit some ability to forecast acquirer stock price reactions to M&A announcements, whereas parametric models do not. Feature importance analyses reveal that a handful of predictors, including acquirer size and (relative) deal size, contribute most to the predictions. Our findings have practical implications for corporate managers and various corporate stakeholders.

Keywords: Mergers and Acquisitions, Forecasting, Shareholder value, Investor perceptions, Machine Learning

1 Introduction

Mergers and acquisitions (M&A) are among the most critical and consequential strategic decisions companies can make (Betton et al., 2008; Mulherin et al., 2017). Over the past two decades, the global M&A market has recorded more than 880,000 transactions, with a total deal value exceeding \$63 trillion.[1] For an acquiring firm, an M&A announcement can have dramatic positive or negative effects on shareholder value (Moeller et al., 2004; Martynova and Renneboog, 2008; Tunyi, 2021). For example, when Charter Communications (NASDAQ: CHTR) announced its acquisition of Cox Communications for \$34.5 billion on May 16, 2025, Charter's stock price promptly rose by 2%, resulting in a one-day market value gain of approximately \$2 billion. Conversely, when Global Payments (NYSE: GPN) announced its acquisition of Worldpay for \$24.5 billion on April 17, 2025, Global Payment's stock price declined by 17%, leading to a one-day market value loss of roughly \$3.9 billion.[2]

Given the potentially substantial stock price effects of M&A deals for acquiring firms, an important question is whether these stock price reactions are, to some extent, forecastable using information available prior to the announcement.[3] This question is the key focus of our paper. Our analysis is grounded in the belief that forecasting investor perceptions of a given M&A announcement, as captured by the announcement-period stock price reaction, can benefit acquiring firms in several ways.[4] More particularly, if a negative investor reaction is anticipated, the acquirer may reconsider the deal, adjust its terms (e.g., the proportion of stock financing), or refine its communication strategy to better highlight the transaction's strategic value. Conversely, if a positive investor reaction is forecasted, this may strengthen the acquirer's case for proceeding with the deal. A broader set of corporate stakeholders, including board members, suppliers, customers, and policymakers, should also be interested in these forecasts, given the economic magnitude of the wealth creation or destruction associated with M&A activity. Forecasted announcement returns may help these parties make more informed decisions about the desirability of a proposed (or rumored) acquirer–target combination, even before a deal is formally announced.

To examine the forecastability of acquirer stock price reactions, we use a sample of 9,517 M&A announcements by US-domiciled public acquirers between 1992 and 2022, constructed using standard screening criteria (Netter et al., 2011; Jaffe et al., 2013). As independent variables, we use publicly-available acquirer, target, deal, and macroeconomic characteristics considered by previous event studies on M&A shareholder value effects (Moeller et al., 2004; Harford and Li, 2007; Jaffe et al., 2013; Ishii and Xuan, 2014; Elnahas and Kim, 2017; Adra et al., 2020; Tunyi, 2021).

Consistent with most previous studies on acquiring-firm stock price reactions (Andrade et al., 2001; King et al., 2004; Renneboog and Vansteenkiste, 2019; Hu et al., 2020), we measure acquirer abnormal stock returns net of contemporaneous "normal" stock returns calculated with a standard market model approach, thus capturing the incremental stock price effect of M&A announcements (Eden et al., 2022).

[1] Source: authors' calculations based on data from Securities Data Company (SDC) Platinum.
[2] Sources: `https://tinyurl.com/myw55hmn` and `https://tinyurl.com/2f4w4rcw`. For comparison, the US stock market recorded returns of 0.12% and 0.70% on the respective announcement dates.
[3] We use the terms forecastable and predictable interchangeably throughout the paper.
[4] Although technically the term investor may include (convertible) bondholders, in this paper we use it as a synonym for shareholder.

Following a standard event study methodology (Brown and Warner, 1985; Kothari and Warner, 2007; Eden et al., 2022), we consider the acquirer's cumulative abnormal stock return (CAR) in the three trading days centered around the deal's announcement date.

To assess the predictability of acquirer CAR, we use Machine Learning (ML) methodology. Two features of this methodology are particularly attractive for our research design. First, ML tends to rely on out-of-sample forecastability and concomitant procedures of cross-validation (Makridakis et al., 2023; Valizade et al., 2024). This approach is in accordance with the widely-accepted agreement, within the forecasting community, that forecasting methods ought to be compared based on their accuracy using out-of-sample testing to avoid overfitting and mitigate data-mining concerns (Makridakis, 1990; Tashman, 2000; Campbell and Thompson, 2008; Ferson et al., 2013).[5] Second, while canonical statistical methods rely on strict distributional assumptions concerning an unknown data generation process (Valizade et al., 2024), ML is underpinned by a model-free, algorithmic approach (Ranta et al., 2023). This model-agnostic way of analyzing data is suitable for our research purpose, because the process by which investors embed ex ante features in their reaction to corporate announcements is unknown (Liberti and Petersen, 2019).

Consistent with the recent literature on empirical forecasting competitions (Makridakis et al., 2020), we evaluate a range of well-known methods for cross-sectional forecasting, including three parametric models (Ordinary Least Squares (OLS), Ridge, and Lasso regressions) and three nonparametric models (Random Forest (RF), k-Nearest Neighbors (KNN), and Light Gradient-Boosting Machine (LGBM)). Our motivation for considering different models is that it is not possible to know which algorithm will outperform the others without testing their accuracy on the data, a notion popularised in the "No free lunch theorem" (Wolpert, 1996).

Also following common practice in the ML literature, we conduct the estimations using a rolling-window approach with a five-fold cross-validation to select hyperparameters over the training dataset. Our approach avoids a look-ahead bias and accounts for the fact that acquirer stock price reactions and their determinants might evolve over the research period (Alexandridis et al., 2017; Cao and You, 2024).

We caution that it would be unrealistic to expect a high degree of forecastability for acquirer stock price reactions. Daily abnormal returns around major corporate announcements are known for having a very high noise-to-signal ratio (Chacko et al., 2008), due to market microstructure effects, pre-announcement rumors, information asymmetries, and divergent investor interpretations of publicly-available information (Liberti and Petersen, 2019). In the context of M&A, additional noise may arise from managers' strategic disclosure of stock-price-relevant information prior to the deal's formal announcement (Ahern and Sosyura, 2014).

Our main empirical results are the following. We obtain the highest out-of-sample R^2 (henceforth referred to as R^2_{OS}) for nonparametric methods, indicating the benefits from adopting a model-agnostic approach when forecasting acquirer stock price reactions. However, even the best-performing approach (LGBM) has only a modest power to forecast the magnitude of acquirer stock price reactions, with R^2_{OS} not exceeding 2.5% for any of the models analysed. Compared to KNN and Random Forest, the superior

[5] Data mining is still possible in out-of-sample forecasts, although substantially less likely than for in-sample forecasting (Timmermann, 2018).

performance of LGBM likely reflects its ability to flexibly capture nonlinearities and complex interactions in M&A data through gradient-based boosting. None of the parametric methods is able to beat using an average CAR as a forecast. In follow-up robustness tests, we obtain a similar conclusion if we conduct the analysis separately for subsamples of deals based on payment type, and for private targets. We also obtain similar results if we include an additional explanatory variable capturing deals withdrawn after their announcement.

A subsequent feature importance analysis using SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017; Ranta et al., 2023) reveals that a small set of predictors—most notably acquirer size and (relative) deal size—contribute disproportionately to the predictions of the nonparametric models. Interestingly, many of these variables are standard "bread-and-butter" controls in prior event studies of acquirer announcement returns, rather than focal independent variables. Our analysis suggests that they are, in fact, key drivers of acquirer stock price reactions.

To the best of our knowledge, ours is the first paper to address the question of whether investor perceptions of M&A announcements can at all be forecasted with a set of straightforward-to-obtain, publicly-available features. As such, our paper departs from previous studies on M&A shareholder value effects, which mostly rely on in-sample OLS regressions with a focus on description rather than forecasting. While most studies find positive target-firm stock price reactions, empirical evidence on the sign and magnitude of acquiring-firm stock price reactions is mixed (Martynova and Renneboog, 2008; Netter et al., 2011; Alexandridis et al., 2017). Findings on the determinants of acquiring-firm announcement returns are equally inconclusive (Travlos, 1987; Fuller et al., 2002; King et al., 2004; Malmendier and Tate, 2008; Deng et al., 2013; Jaffe et al., 2013; Eckbo et al., 2018; Renneboog and Vansteenkiste, 2019; King et al., 2021). The overall explanatory power of regression models of acquirer returns tends to be very low, with in-sample R^2s hovering around 5% (Travlos, 1987; Fuller et al., 2002; Malmendier and Tate, 2008; Deng et al., 2013; Jaffe et al., 2013; Eckbo et al., 2018), and Campbell et al. (2016) calling the lack of systematic cross-sectional evidence on acquirer announcement return determinants "disheartening".

Whilst we use parametric and nonparametric models, we wish to emphasize that our contribution goes beyond the mere assessment of whether more complex methods outperform the OLS model, a staple of previous studies on the drivers of acquirer stock price reactions. Prior to our paper, even the OLS model had not yet been formally evaluated for its out-of-sample forecasting properties of acquirer stock price reactions. Our key research question is, therefore, whether investor perceptions of M&A announcements, as captured by acquirer announcement returns, can be forecasted in the first place using *any* method available to decision makers. We are also the first to examine the contribution of each feature to the final prediction using Explainable AI (XAI), through a SHAP analysis. Finally, our paper contributes to previous studies using ML methods to forecast stock returns (Gu et al., 2020; Fieberg et al., 2023; Zhou et al., 2023). While these papers focus on general, unconditional stock return predictions, our paper forecasts incremental stock price reactions to a major corporate announcement, i.e., the news that the firm will acquire a particular other firm.

The remainder of this article is organised as follows. The next section provides a brief overview of the three main strands of literature relevant to our paper. Section 3 outlines the data collection, measurement of acquirer announcement returns, and independent variables used for forecasting these returns. Section 4 describes the forecasting, cross-validation and hyperparameter selection methods used in our analysis. Section 5 provides and discusses the methods' performance in predicting acquirer CAR. Section 6 discusses the results of the variable importance analysis. Section 7 concludes with a summary of our key findings, their implications for practitioners and academics, and avenues for future academic research.

2 Position in the literature

In this section, we position our work relative to previous studies on acquiring-firm stock price reactions, financial forecasting, and machine learning.

2.1 Acquiring-firm stock price reactions to deal announcements

A central theoretical pillar underpinning the vast empirical literature on stock price reactions to M&A announcements is the efficient market hypothesis (Fama, 1970). This hypothesis posits that stock prices immediately reflect the discounted incremental cash flow effects of any publicly-available corporate news, without delays or biases. Even if some investors behave irrationally, rational, well-informed arbitrageurs will quickly restore prices to their "efficient" levels (Brealey et al., 2018). In the context of M&A, stock market efficiency would imply that any change in the acquirer's stock price following a deal's announcement fully and accurately reflects shareholders' assessment of the (positive or negative) discounted net cash flow implications of the announced transaction for the firm.

Building on this key premise of efficient market theory, a large body of literature has used acquirer stock price reactions to M&A deals as a measure of the discounted net cash flows generated by the deal for the acquirer—or, in other words, the deal's net present value (NPV) (Andrade et al., 2001; King et al., 2004; Jaffe et al., 2013; Junni et al., 2015; Renneboog and Vansteenkiste, 2019; Tunyi and Machokoto, 2021). Corporate finance theory yields conflicting predictions regarding the NPV of M&A deals for the acquiring firm (Bruner, 2002). Value-increasing theories posit that M&A generate positive cash flow effects, driven by synergies from cost savings, revenue enhancement, and the elimination of inefficient management (Manne, 1965; Bradley et al., 1988; Houston et al., 2001). These theories predict a positive acquirer CAR around the deal announcement. Value-decreasing theories, in contrast, suggest that M&A are motivated by factors that reduce cash flows for the acquirer, such as managerial hubris (Roll, 1986), overconfidence (Malmendier and Tate, 2008), empire building (Jensen, 1986), and entrenchment (Shleifer and Vishny, 1989). These theories predict a negative acquirer announcement return. Perhaps not surprising in light of these opposing theoretical predictions, empirical evidence on the sign and magnitude of shareholder wealth effects for acquiring-firm shareholders is inconclusive (Martynova and Renneboog, 2008). Some studies find positive average effects (Netter et al., 2011; Alexandridis et al., 2017), other studies negative average effects (Andrade et al., 2001; Moeller et al., 2004), yet others find no significant stock price reactions at all (Datta et al., 1992; Bruner, 2002).

Although acquirer stock price reactions are a ubiquitous measure of M&A value creation, we caution against interpreting these returns as a market-based assessment of deal performance. The reason is that, in addition to investors' assessment of the deal's NPV, stock price reactions to deal announcements may also reflect a reassessment of the acquiring firm's standalone value, inferred from the deal's timing and payment terms (Malmendier et al., 2018; Ben-David et al., 2025). More specifically, in a context of information asymmetry between managers and investors, stock-financed M&A announcements may signal to the market that the acquiring firm is overvalued, leading to a negative acquirer announcement return (Myers and Majluf, 1984; Travlos, 1987). Acquirer stock price reactions in stock-financed deals may also reflect downward price pressure from shorting transactions by merger arbitrageurs (Mitchell et al., 2004; Dutordoir et al., 2022), who want to hedge their equity exposure to the firm. Further complicating matters, some studies document evidence consistent with bounded investor rationality in the interpretation of M&A deal announcements, and advocate for the adoption of a behavioral perspective on investor reactions to M&A news (Schijven and Hitt, 2012; Campbell et al., 2016).

Given the challenges in interpreting acquirer CAR as a reliable measure of deal value creation, we instead view these stock price reactions as broader reflections of investor perceptions of M&A announcements. This framing allows acquirer announcement returns to reflect not only investor assessments of the deal's NPV but also of the standalone value of the acquirer. It also accommodates potential departures from market efficiency and acknowledges bounded investor rationality. Accordingly, our main aim is to assess the extent to which investor perceptions of M&A announcements can be predicted using information available at the time of the deal announcement.

Forecasting investor reactions to a planned deal may be useful for corporate managers for at least three reasons. Firstly, assessing predicted investor reactions may help managers optimize deal structure and communications. For example, if a negative stock price reaction is forecasted, managers may choose to rely less on stock financing for the deal. They may also want to time the announcement strategically, e.g. by aligning it with positive earnings surprises, and make a greater effort to proactively highlight synergies to investors. Secondly, forecasting investor reactions may provide the prospective bidding firm with a competitive edge. More particularly, if rival firms are considering the same target, a favorable predicted stock market reaction could strengthen the bidding firm's negotiating position. Thirdly, data-backed insights into potential stock price reactions could help strengthen the acquiring firm's manager's case for the deal towards the board of directors and other stakeholders, such as employees, supply chain partners, and regulators.[6]

Of course, it is uncertain whether managers would effectively take these investor reaction predictions into account in their decision-making processes. Finance theory suggests there may effectively be a feedback effect from stock returns following corporate events to management decisions (Dye and Sridhar, 2002). More particularly, while the manager undoubtedly has first-hand information about his firm, the

[6]To gauge Generative Artificial Intelligence's answer on why acquiring-firm managers may want to predict stock price reactions to M&A announcements, we prompted ChatGPT with the following question, in March 2025: "Imagine I am a corporate CEO of a publicly-listed, US company considering the takeover of company X. Imagine that someone has developed a tool that can predict my stock price reaction to the announcement of the takeover of company X. Why would I want to use this tool?" ChatGPT's answer mentioned: "This tool would provide a strategic advantage in risk management, deal structuring, and investor relations—all critical factors in executing a successful acquisition."

power of the stock market lies in the aggregation of information from many different, often sophisticated traders. These market participants may have relevant information on different dimensions of the deal, for example related to its anticipated synergies (Goldstein, 2023). As long as there is some information that managers do not have, they should rationally update their beliefs on the deal's quality based on stock price reactions to the deal announcement. However, managers may refuse to listen to the market due to hubris (Roll, 1986) or agency conflicts (Kau et al., 2008), or because they know that the investor reactions are inflated by the management's own, overly favorable news releases prior to the M&A announcement (Ahern and Sosyura, 2014). Empirical evidence on whether managers withdraw their deals following a negative stock market reaction to their M&A announcement is mixed: Jennings and Mazzeo (1991) find no feedback effect from the stock market to managers, while Luo (2005), Kau et al. (2008), and Abed and Abdallah (2017) find that bidding firms do respond to the stock market. While these studies focus on the link between past stock price reactions and managerial decisions, we argue that it could be even more useful for managers and other stakeholders to have an (albeit imperfect) forecast of the stock price reaction to a future, as yet unannounced deal they are considering. Whether acquirer announcement returns are at all forecastable is, therefore, the key research question addressed in our paper. To the best of our knowledge, no previous study has addressed this question.

2.2 Financial forecasting

Our study aims to predict acquirer stock price reactions, and is therefore also part of the financial forecasting literature. Timmermann (2018) outlines the challenges associated with obtaining good forecasting results in finance as "(...) the difficulty of establishing predictability in an environment with a low signal-to-noise ratio, persistent predictors, and instability in predictive relations arising from competitive pressures and investors' learning." Most of this literature relates to the prediction of asset prices, with an increasing weight being given to out-of-sample predictability (Brown et al., 1987; Çakmaklı and van Dijk, 2016; Timmermann, 2018; Grønborg et al., 2021). Overall, predictability of asset returns is very low even when compared to other challenging forecasting problems, such as predicting microeconomic indicators (Timmermann and Granger, 2004; Timmermann, 2018).

A few asset pricing papers, however, have obtained moderate success. Kanas (2003), for instance, examines the US stock market annual returns spanning the period 1872–1999 and observes these to have some degree of predictability. Çakmaklı and van Dijk (2016) show that factor-augmented predictive regressions have a superior performance to existing benchmarks in predicting monthly US excess stock returns. More recently, Haase and Neuenkirch (2023) attempt to forecast S&P500 stock returns and obtain some success in terms of risk-adjusted performance, although the proposed models do not outperform commonly-used benchmarks.

Next to studies on the predictability of asset returns, a smaller group of papers focuses on predicting corporate finance decisions or outcomes, including corporate financial distress and bankruptcy (Altman, 1968; Shumway, 2001; Jones and Hensher, 2004), security choices (Bayless and Chaplinsky, 1991; Lewis et al., 1999), and the likelihood of corporate restructuring activity (Palepu, 1986; Shumway, 2001) and of

M&A deals in particular (Rodrigues and Stevenson, 2013; Song and Walkling, 2000; Futagami et al., 2021; Tunyi, 2021). We contribute to this literature by assessing the predictability of acquirer announcement returns for given acquirer–target combinations, rather than the probability of a deal happening.

2.3 Machine Learning applications in business research

Our paper contributes to a fast-growing literature that uses ML techniques to address technically complex management questions. ML is a multidisciplinary field that merges insights from computer science and statistical learning to build algorithms capable of learning patterns and associations from data without human supervision (Makridakis et al., 2023; Valizade et al., 2024). In finance, ML is most commonly applied in asset pricing, including algorithmic trading, risk analysis and assessment, fraud detection, and portfolio optimization.[7] In contrast, applications of ML in corporate finance remain relatively limited. For example, Li et al. (2021) use ML techniques on earnings call transcripts to construct a corporate culture metric, which they link to major corporate events such as M&A. Beyond finance, ML has also been adopted in other management areas, including human resources management (Pei et al., 2024) and marketing (Ngai and Wu, 2022).

Consistent with our study, several papers use ML approaches to make financial forecasts. As such, the ML and the financial forecasting literature partially overlap. For example, Obaid and Pukthuanthong (2022) find that an investor sentiment index obtained using ML methods on newspaper pictures can predict market return reversals and trading volume. Erel et al. (2021) apply ML techniques for predicting company director performance, and Gu et al. (2020), Fieberg et al. (2023) and Zhou et al. (2023) use ML to predict stock prices. While these papers study unconditional return forecasts, our paper forecasts abnormal stock returns conditional on a major corporate announcement, i.e., the incremental stock price reaction to the news that the firm will acquire a particular other firm. More similar to our study, Bozos and Nikolopoulos (2011) use a variety of ML approaches to forecast stock price reactions to seasoned equity offering announcements and find these to be partly predictable.

We contribute to this literature strand by using ML approaches to assess whether acquirer announcement returns can be forecasted. We emphasize that our contribution lies in addressing the M&A announcement return forecastability question per se, rather than in merely comparing ML with OLS models. Before our study, even OLS models had not yet been formally assessed for their forecasting ability in the context of acquirer stock price reactions.

3 Sample and variables

In this section, we describe the construction of the dataset of M&A deals, the measurement of acquirer announcement returns, and the selection of independent variables used in the forecasting analysis.

[7] A comprehensive review of ML applications in asset pricing is beyond the scope of this article. For detailed surveys, see Buchanan (2019), Ozbayoglu et al. (2020), Hoang and Wiegratz (2023), and Ranta et al. (2023).

3.1 M&A screening process

In a first step, we collect a sample of 363,706 M&A deals between the dates of 01/01/1992 and 31/12/2022 from the Securities Data Company Platinum (henceforth SDC), the reference database for empirical research on M&A (Bollaert and Delanghe, 2015), which is currently embedded in the London Stock Exchange Group (LSEG) Workspace platform. We start in 1992 because some evidence suggests that SDC's deal coverage is incomplete before that year (Netter et al., 2011). We then impose a number of data screens, which are standard in the empirical literature on M&A (Fuller et al., 2002; Netter et al., 2011; Jaffe et al., 2013; Eckbo et al., 2018). More particularly, we exclude observations for which we have missing data on the deal value, or where the final target share owned by the acquirer is lower than 50% (Netter et al., 2011). We limit our study to acquisitions made by firms headquartered in the US. Finally, we remove financial institutions and utility firms from our sample by excluding deals whose acquirer primary SIC code is between 6000 and 6999 or between 4900 and 4999. The exclusion of financial institutions and utility firms is common in corporate finance research and M&A studies in particular (Andrade et al., 2001; Eckbo et al., 2018), because these firms tend to be more regulated and have a different balance sheet structure than industrial firms (Li et al., 2016). We remove deal announcements whereby the same acquirer makes multiple announcements on the same date, as it is impossible to associate the stock price reaction with a given set of deal characteristics for these announcements.

From the remaining sample, we discard M&A for which acquirer balance sheet information or daily stock returns are not available from Compustat Fundamentals Annual and the Center for Research in Security Prices (henceforth CRSP), respectively, for the fiscal year-end preceding the deal's announcement date, which we obtain from SDC. Subsequently, we exclude observations for which the deal size is smaller than $10 million, as in Eckbo et al. (2018). The final dataset contains 9,517 observations. Since we closely follow recommendations from seminal work in M&A throughout our data collection and screening process, we are confident that observations within our sample are representative of a "typical" M&A, and therefore suitable for the purposes of the study.

3.2 Acquiring-firm stock price reactions

To calculate acquirer stock price reactions to M&A announcements, we follow a similar event study methodology as in prior research on M&A shareholder wealth effects (Louis, 2005; Cai and Sevilir, 2012; Ishii and Xuan, 2014; Croci and Petmezas, 2015). In a first step, we calculate normal acquirer stock returns using a standard market model approach (Kothari and Warner, 2007). We regress the acquirer's stock return R_{it} on the market return R_{mt}, which we measure as the return over the CRSP equally-weighted stock market index, i.e.:

$$R_{it} = a_i + \beta_i R_{mt} + e_{it},$$

with e_{it} a mean zero, independent disturbance term at time t.

The event study literature does not have standard rules for the exact length and end date of the estimation window in market model regressions. Practices vary substantially across event studies and likely reflect a tradeoff between including more observations to increase statistical accuracy and avoiding going too far back from the announcement date, in case the parameters of the return-generating mechanism have shifted (Strong, 1992). We take the pragmatic approach of using the default estimation period settings in Eventus, the software we use for the event study estimation.[8] These consist of an estimation period length of maximum 255 trading days, ending on day 46 before the event day 0 (the announcement date of the deal), to reduce the risk of event-related contamination of the market model results. In unreported robustness tests, we find that the abnormal return estimations are largely insensitive to the use of alternative estimation period lengths and ending days.[9]

For each deal, we calculate the acquirer's normal return over the event date using the intercept α_i and slope coefficient β_i obtained from the market model regression. We obtain the acquirer's abnormal stock return on the event date by deducting this normal return from the acquirer's raw return over the event date. We follow a similar procedure for trading days -1 and +1 around the event date, and aggregate the resulting abnormal stock returns in a cumulative abnormal stock return or *CAR*. The use of a three-day event window is standard practice in the M&A literature (Louis, 2005; Harford and Li, 2007; Ishii and Xuan, 2014; Becht et al., 2016). It has the advantage of accounting for potential pre-announcement date information leakage (through the inclusion of day -1) and for announcements made on a non-trading day or after stock market closure on a trading day (through the inclusion of day 1). In an unreported robustness test, we obtain a weaker forecasting accuracy when we measure *CAR* over larger event windows, e.g. a window ranging from trading days -2 to +2. This reduced accuracy can be explained by the additional noise resulting from the inclusion of additional days in the event window (Kothari and Warner, 2007). We therefore focus on the *CAR[-1,1]*, or shortly *CAR*, throughout the paper.

Table 1: Summary statistics for the acquirer *CAR[-1,+1]*

Mean	1.10%***
Median	0.60%***
Standard Deviation	7.65%
Percentage Negative	43.62%
t-test	14.07%
	(0.00%)
Wilcoxon test	1.85E+7
	(0.00%)
N	9,517

Notes: This table reports summary statistics for acquirer cumulative abnormal returns (*CAR*) over the event window [-1, +1] around M&A announcements. $^{***}p < 0.01$.

Table 1 describes the acquirer *CAR*. Consistent with a number of previous event studies (Cai et al., 2011; Netter et al., 2011; Alexandridis et al., 2017), we find a positive mean and median *CAR* for the full sample (1.10% and 0.60% respectively). A t-test and Wilcoxon test indicate that the *CAR* is significantly different from zero. We also find substantial variation in the *CAR*, with 43.62% of the deal announcements provoking a negative stock price reaction.

[8]Eventus performs event studies using stock price data obtained directly from the CRSP database.

[9]Detailed results of all unreported robustness tests mentioned in the paper are available from the corresponding author.

3.3 Independent variables

A rich empirical literature examines the determinants of acquirer announcement returns, albeit with few consistent findings across the different studies. To identify suitable determinants of acquirer announcement returns, we conduct a review of relevant studies published in top finance and general management journals from the early 1990s until mid-2024. We focus on variables that are commonly used in the academic literature and can be obtained from standard databases.

Since our primary goal is forecasting, we do not distinguish between control and main variables, as all predictors are used jointly to forecast the CAR. We are also not overly concerned about multicollinearity, for two reasons. First, our focus is on prediction accuracy rather than on interpreting the individual effects of variables. Second, although multicollinearity can pose problems for in-sample analyses using traditional parametric methods such as OLS, the built-in variable selection and dimension reduction capabilities of ML methods make them well suited to handle correlated predictors (Gu et al., 2020).

Broadly speaking, the variables included in previous studies on acquirer stock price reactions fall into four main categories: acquirer, target, deal, and macroeconomic characteristics. The inclusion of these variables is typically justified by referring to key M&A theories (Manne, 1965; Jensen, 1986; Roll, 1986; Travlos, 1987; Bradley et al., 1988; Shleifer and Vishny, 1989; Dong et al., 2006; Malmendier and Tate, 2008; King et al., 2021; Schneider and Spalt, 2021). We also include proxies for investor sentiment in our set of variables, to allow for behavioral elements in the interpretation of M&A announcements. We now briefly describe the four categories of variables used in our analysis.

The first and largest category consists of a set of standard acquirer characteristics, as in Moeller et al. (2004), Harford and Li (2007), Ishii and Xuan (2014), Becht et al. (2016), Elnahas and Kim (2017), and Adra et al. (2020), among others. We obtain these variables from Compustat Fundamentals Annual. We measure the acquirer characteristics as of the fiscal year-end before the deal's announcement date, since we want to replicate the information set that is available to decision makers as of the time of the deal selection. In particular, we consider the acquirer's Size, Return on Assets, Cash, Free Cash Flow, Market to Book, Leverage, and Research and Development (R&D) Intensity. We also construct two variables capturing the acquirer's financial distress costs (Altman Z and Bankruptcy Risk). We furthermore include a High Tech Industry dummy variable identifying acquirers in technology-intense industries. Finally, to capture acquirer learning effects and potential overconfidence coming with more experienced acquirers (Jaffe et al., 2013; Renneboog and Vansteenkiste, 2019), we use an Experience variable measuring the number of deals the same acquirer has announced over the previous years.

M&A rationales do not yield clear predictions regarding the direction of the impact of these acquiring-firm characteristics on acquirer announcement returns. For example, the Market to Book ratio could capture acquirer growth opportunities, and therefore be associated with a synergistic rationale for M&A and more positive acquirer announcement returns (Bradley et al., 1988). But it could also capture acquirer overvaluation, and therefore be associated with an opportunistic rationale for M&A and more negative acquirer announcement returns (Shleifer and Vishny, 2003; Dong et al., 2006; Eckbo et al., 2018). Perhaps unsurprisingly given this association of given proxies with conflicting theories, empirical studies do not obtain conclusive evidence on their impact. For example, while Moeller et al. (2004) and Field and Mkrtchyan (2017) find a negative impact of acquirer size proxies on acquirer announcement returns, Ishii and Xuan (2014) do not find a significant impact.

The second category of independent variables consists of a set of standard target characteristics, as in Asquith et al. (1983), Travlos (1987), Cai and Sevilir (2012), and Becht et al. (2016), among others. In particular, we obtain the following four target characteristics from SDC: Deal Size (capturing the size of the target), Relative Deal Size (capturing the size of the target relative to the acquirer), a Public Target dummy, and an International Target dummy.[10] We do not have clear predictions for Deal Size and Relative Deal Size, because these proxies can be linked with opposing theoretical rationales. To give an example, Asquith et al. (1983) find a positive impact of Relative Deal Size on acquirer announcement returns, which they attribute to the fact that the synergies of the deal are amplified for larger deals (Schneider and Spalt, 2021; Kališ et al., 2023), while Alexandridis et al. (2013) find a negative impact, which they attribute to the higher complexities associated with the post-merger integration of larger targets. Relative Deal Size can also capture price pressure effects of merger arbitrageurs in stock-financed deals (Mitchell et al., 2004). Some empirical studies find a negative impact of a Public Target dummy variable, which could be attributable to the fact that acquirers may receive a better price for private targets (Fuller et al., 2002; Becht et al., 2016). Moeller and Schlingemann (2005) find more negative acquirer announcement returns for acquirers involved in deals with international (non-US) targets.

The third category of independent variables consists of a set of standard deal-related characteristics, as in Jaffe et al. (2013), Alexandridis et al. (2017), and Eckbo et al. (2018). We obtain these variables from SDC, except for Dormant Period which we construct based on the dataset's observations. We identify All Stock and All Cash deals with corresponding dummy variables, the remaining deals being financed with a combination of stock and cash. We predict a more negative acquirer stock price reaction for deals with a higher percentage of acquirer stock financing, due to the adverse signal that the decision to use stock may send about acquiring-firm overvaluation (Travlos, 1987; Eckbo et al., 2018). We also identify Hostile deals, which are predicted to have more negative stock price effects since acquirers are more likely to overpay in hostile deal scenarios in order to secure the target (Servaes, 1991; Jaffe et al., 2013). We account for the similarity between acquirer and target business activities through an Industry Relatedness dummy variable. We predict a positive impact for this variable, because similar target and acquirer activities could signal more valuable synergies (Morck et al., 1990; Louis, 2005). We furthermore include a Merger of Equals dummy capturing deals between similar-sized firms. Stock price reactions to these deals may be more negative, since investors may anticipate difficulties associated with the post-merger integration of two similar-sized firms (Zaheer et al., 2003). We include three dummy variables capturing the effects of the timing of a deal. A Friday deal dummy identifies deals announced on a Friday, as some evidence suggests these deals might be met by weaker stock price reactions due to investor inattention (Reyes, 2018). A Winter deal dummy captures the stock price effects of negative weather-induced moods,

[10]Similar to other empirical studies, we are constrained from obtaining additional target-specific variables by the fact that Compustat only provides data for public firms. Most importantly, as only 17% of the targets are listed on a stock market, we cannot include the takeover premium as an explanatory variable as this measure is unavailable for private targets (Officer, 2007).

pessimism and risk aversion amongst investors (Tunyi and Machokoto, 2021). A Dormant Period deal captures the number of days since the last M&A announcement in the same industry (Cai et al., 2011). Finally, Percentage Sought captures the percentage of the target the acquiring firm seeks to acquire - by construction, this is higher than 50%.

The fourth category consists of a set of potentially relevant macroeconomic characteristics, obtained from the Federal Reserve Economic Data (FRED). We measure these as closely as possible, but prior to the M&A announcement. We include the monthly Federal Funds Rate for the US to capture acquirers' costs of debt financing for their deals. Higher financing costs are predicted to adversely affect M&A outcomes (Martynova and Renneboog, 2009; Adra et al., 2020). We furthermore add the VIX index, a Risk Sentiment index, and two US policy uncertainty indices (EPU 3C and EPU News) as proxies for overall uncertainty (Baker et al., 2016). Previous studies show an adverse effect of uncertainty on M&A decisions and their outcomes (Bhagwat et al., 2016; Nguyen and Phan, 2017; Bonaime et al., 2018). We also add dummies for the COVID Pandemic and US Election years, and include changes in the US Gross Domestic Product (GDP). Finally, we include a Merger Wave variable capturing periods with a high volume of deals. Table 2 summarizes the definitions of the independent variables.

Table 2: Measurement of independent variables

Variable & Measurement	
Size	Acquirer log(Total Assets)
Return on Assets	Acquirer Net Income / Total Assets
Cash	Acquirer Cash and Short-term Investments / Total Assets
Free Cash Flow	Acquirer (Operating Income before Depreciation and Amortization - Interest & Expenses - Income Taxes - Capital Expenditures)/Total Assets
Market to Book	Acquirer Market Value of Equity / Total Assets
Leverage	Acquirer Total Liabilities/Total Assets
R&D Intensity	Acquirer R&D Expenditure/ Total Assets
Altman Z	3.3 (EBIT/Total Assets) + 0.99 (Sales/Total Assets) + 0.60 (Market Capitalization / Liabilities) + 1.2 (Working Capital/Total Assets) + 1.4 (Retained Earnings/ Total Assets) (all variables are the acquirer's)
Bankruptcy Risk	Boolean variable, equal to 1 if Altman Z < 1.81 and equal to 0 otherwise
High Tech Industry	Boolean variable, equal to 1 if the acquirer SIC $\in$ {3571, 3572, 3575, 3577, 3578, 3661, 3663, 3669, 3674, 3812, 3823, 3825, 3826, 3827, 3829, 3841, 3845, 4812, 4813, 4899, 7370, 7371, 7372, 7373, 7374, 7375, 7378, 7379} and equal to 0 otherwise
Experience	The number of acquisitions by the same acquirer over the past two years represented in our sample
Deal Size	Size of the deal in millions (USD)
Relative Deal Size	Size of the deal in millions (USD)/(Acquirer) Total Assets
Public Target	Boolean variable, equal to 1 if the target is publicly quoted and equal to 0 otherwise
International Target	Boolean variable, equal to 1 if the target is not domiciled in the US and equal to 0 otherwise
All Stock	Boolean variable, equal to 1 if deal is funded 100% with stocks and equal to 0 otherwise
All Cash	Boolean variable, equal to 1 if the deal is funded 100% with cash and equal to 0 otherwise
Hostile	Boolean variable, equal to 1 if the takeover is hostile and equal to 0 otherwise
Industry Relatedness	Boolean variable, equal to 1 if acquirer and target have the same primary SIC code and equal to 0 otherwise
Merger of Equals	Boolean variable, equal to 1 if acquirer and target have approximately the same market capitalization and equal to zero otherwise
Friday Deal	Boolean variable, equal to 1 if the M&A announcement was on a Friday and equal to 0 otherwise

Continuation ...	
Winter Deal	Boolean variable, equal to 1 if the M&A announcement was made during the winter and equal to 0 otherwise
Dormant Period	Number of days prior to the the last M&A in the same industry (same primary SIC code)
Percentage Sought	Percentage of shares the acquirer is seeking to purchase in the transaction
Federal Funds Rate	Interest rate charged to US commercial banks and other depository institutions on loans they receive from their regional Federal Reserve Bank's lending facility
VIX	Chicago Board Options Exchange volatility index
Risk Sentiment	Equity Market Volatility Tracker capturing investor sentiment about overall risk levels, based on articles in 3,000 US newspapers that contain terms including "economy", "equity", "volatility", "pandemic" and more. The higher the count, the higher the value of the tracker
EPU 3C	US Monthly Economic Policy Uncertainty Three Component Index
EPU News	US Monthly Economic Policy Uncertainty News Based Policy Uncertainty Index
Pandemic	Boolean variable, equal to 1 if the announcement was made in the years 2019, 2020 or 2021 and equal to 0 otherwise
Election	Boolean variable, equal to 1 if the announcement was made in a US presidential election year and equal to 0 otherwise
GDP Change	Percentage change in US GDP
Merger Wave	Total number of M&As in the preceding year represented in our sample

Notes: This table provides detailed descriptions of all independent variables used in the analysis. The variables capture acquirer characteristics, target characteristics, deal-specific features, and contextual or macroeconomic factors. Boolean variables equal 1 if the specified condition is met, and 0 otherwise.

Table 3 presents descriptive statistics for the independent variables, which are similar to those reported by other studies on M&A.[11]

To further gauge the similarity between our sample and the M&A deals used in prior studies, we run an in-sample OLS regression of the *CAR* on the independent variables. Table 4 provides the results. We obtain an in-sample R^2 of 0.041 and an adjusted R^2 of 0.038, in line with the magnitudes of in-sample R^2 reported in previous event studies on acquirer announcement returns.

Although in-sample description is not our key focus, we briefly describe the regression findings. For the acquirer characteristics (for which we did not have clear prior expectations), we find a negative impact of Size, Cash, Free Cash Flow, R&D Intensity, and High-Tech Industry, and a positive impact of Leverage. Investors may be concerned that acquirers with high cash and free cash flow levels are more

[11] For exposition purposes, we report full descriptive statistics for the dummy variables, although these variables are by construction either 0 or 1.

likely to engage in empire-building behavior, whereas a high leverage ratio may act as a disciplinary mechanism that curbs such tendencies (Jensen, 1986). Regarding the target characteristics, we find a positive impact of Relative Deal Size, consistent with Asquith et al. (1983), and a negative impact of the Public Target and International Target dummies, in line with previous findings (Fuller et al., 2002; Moeller et al., 2004; Becht et al., 2016). Our results for the deal-specific characteristics are largely consistent with expectations, except for the significantly negative impact of Industry Relatedness. Finally, for the macroeconomic variables, we observe a negative impact of the Federal Funds Rate and Risk Sentiment, consistent with predictions.

We caution that we obtain high Variance Inflation Factor (VIF) values for two of the features included in our study, namely 12.35 for the policy uncertainty index EPU 3C and 11.59 for the policy uncertainty index EPU News. These values indicate strong multicollinearity. Maximum VIF values for all other features are below 3. While multicollinearity may complicate the interpretation of OLS coefficients and statistical significance, we reiterate that in-sample inference is not our primary focus, and that ML methods are very capable of dealing with highly correlated explanatory variables. In unreported tests, we find that dropping these two policy uncertainty indices does not affect our results. Given these considerations, we retain them in our forecasting analysis, prioritizing a comprehensive set of predictors over multicollinearity concerns, consistent with other forecasting studies.

Table 3: Summary statistics of the independent variables used in the forecasting analysis

	Count	Mean	St.Dev.	Pct(25)	Median	Pct(75)
Size	9,517	7.374	1.988	5.994	7.296	8.639
Return on Assets	9,517	0.033	0.185	0.016	0.053	0.089
Cash	9,517	0.183	0.195	0.037	0.108	0.260
Free Cash Flow	9,517	0.037	0.128	0.019	0.057	0.092
Market to Book	9,517	4.665	28.382	1.861	2.892	4.753
Leverage	9,517	0.502	0.231	0.349	0.500	0.636
R&D Intensity	9,517	0.040	0.063	0.000	0.014	0.058
Altman Z	9,517	6.918	25.469	2.731	4.000	6.157
Bankruptcy Risk	9,517	0.121	0.326	0.000	0.000	0.000
High Tech Industry	9,517	0.344	0.475	0.000	0.000	1.000
Experience	9,517	2.408	2.380	1.000	2.000	3.000
Deal Size	9,517	715.787	4,252.920	32.000	92.000	310.000
Relative Deal Size	9,517	0.291	1.673	0.024	0.080	0.236
Public Target	9,517	0.170	0.376	0.000	0.000	0.000
International Target	9,517	0.210	0.408	0.000	0.000	0.000
All Stock	9,517	0.090	0.286	0.000	0.000	0.000
All Cash	9,517	0.352	0.478	0.000	0.000	1.000
Hostile	9,517	0.005	0.069	0.000	0.000	0.000
Industry Relatedness	9,517	0.656	0.475	0.000	1.000	1.000
Merger of Equals	9,517	0.002	0.040	0.000	0.000	0.000
Friday Deal	9,517	0.116	0.320	0.000	0.000	0.000
Winter Deal	9,517	0.245	0.430	0.000	0.000	0.000
Dormant Period	9,517	156.928	419.981	7.000	30.000	123.000
Percentage Sought	9,517	98.832	6.256	100.000	100.000	100.000
Federal Funds Rate	9,517	2.640	2.019	0.750	2.250	4.860
VIX	9,517	19.371	7.414	13.800	17.880	22.960
Risk Sentiment	9,517	0.402	0.391	0.172	0.306	0.500
EPU 3C	9,517	113.541	41.300	82.384	102.161	137.470
EPU News	9,517	125.081	57.720	84.450	108.758	156.504
Pandemic	9,517	0.107	0.309	0.000	0.000	0.000
Election	9,517	0.216	0.411	0.000	0.000	0.000
GDP Change	9,517	0.011	0.006	0.009	0.012	0.015
Merger Wave	9,517	2,102.769	527.238	1,763	1,985	2,415

Notes: This table presents summary statistics for all independent variables used in the forecasting analysis. "Count" refers to the number of non-missing observations. "Mean" and "St.Dev." are the sample mean and standard deviation. "Pct(25)", "Median", and "Pct(75)" indicate the 25th percentile, median, and 75th percentile, respectively. Continuous variables are expressed in raw or ratio terms, while dummy variables (e.g., Public Target, Pandemic, High Tech Industry) take the value of 1 if the condition is met and 0 otherwise. The dataset covers 9,517 M&A deals.

Table 4: In-sample analysis of determinants of CAR[-1,+1]

	Dependent variable: CAR[1-,1]
	coefficient (st. error)
Size	-0.005^{***} (0.0005)
Return on Assets	0.003 (0.005)
Cash	-0.018^{***} (0.005)
Free Cash Flow	-0.026^{***} (0.008)
Market to Book	-0.00001 (0.00003)
Leverage	0.011^{***} (0.004)
R&D Intensity	-0.022 (0.015)
Altman Z	0.00002 (0.00003)
Bankruptcy Risk	0.003 (0.003)
High Tech Industry	-0.003^{*} (0.002)
Experience	0.0001 (0.0004)
Deal Size	-0.00000 (0.00000)
Relative Deal Size	0.002^{***} (0.0005)
Public Target	-0.013^{***} (0.002)
International Target	-0.006^{***} (0.002)
All stock	-0.008^{***} (0.003)
All cash	0.007^{***} (0.002)
Hostile	-0.008 (0.011)
Industry Relatedness	-0.004^{***} (0.002)
Merger of Equals	-0.004 (0.020)
Friday Deal	0.002 (0.002)
Winter Deal	0.001 (0.002)
Dormant Period	0.00000^{**} (0.00000)
Percentage Sought	-0.0002 (0.0001)
Federal Funds Rate	-0.002^{***}(0.001)
VIX	-0.00000 (0.0001)
Risk Sentiment	-0.008^{***} (0.002)
EPU 3C	-0.0001 (0.0001)
EPU News	0.00004 (0.00005)
Pandemic	0.002 (0.003)
Election	0.002 (0.002)
GDP Change	0.189 (0.154)
Merger Wave	0.00000 (0.00000)
Constant	0.071^{***}
	(0.014)
Observations	9,517
R^2	0.041
Adjusted R^2	0.038
Residual Std. Error	0.075 (df = 9483)
F Statistic	12.392^{***} (df = 33; 9483)

Notes: This table reports the results from an in-sample OLS regression examining the determinants of acquirer cumulative abnormal returns (CAR) over the [-1, +1] event window surrounding M&A announcements. Reported values are coefficients with standard errors in parentheses. The model includes firm-level financials, deal characteristics, and macroeconomic controls. Statistical significance is denoted by * p<0.10, ** p<0.05, and *** p<0.01.

4 Forecasting, cross-validation and hyperparameter selection methods

This section outlines the forecasting methods considered in our study. We also describe our cross-validation and hyperparameter selection methods.

4.1 Forecasting methods

We consider six commonly-used forecasting approaches, which we will describe in further detail in the remainder of this section. The first three methods are parametric, which means that they assume a specific functional form or model for the relationship between variables, typically based on a set of parameters. These models rely on underlying assumptions about the data-generating process, such as linearity or normality, and involve estimating a fixed number of parameters from historical data (Hyndman and Athanasopoulos, 2018). The final three models are nonparametric, which means that they adopt an agnostic, model-free approach to forecasting, without set assumptions about the data-generating process or number of parameters.

4.1.1 OLS regression

Within the literature on acquirer announcement effects, the OLS regression paradigm became particularly popular. Likely reasons are its simplicity, parsimony, and ability to directly test hypotheses pertaining to the cross-sectional variation of M&A announcement effects over a centered event window (Travlos, 1987; Fuller et al., 2002; Malmendier and Tate, 2008; Jaffe et al., 2013; Eckbo et al., 2018), albeit with few systematic findings. In forecasting studies, OLS regressions are widely-used alongside ML models to compare the accuracy of both classes of models (Lessmann and Voß, 2017; Cui et al., 2020). The OLS regression is therefore an obvious initial choice of model for our forecasting analysis. We emphasize that, despite its prevalence in event studies on M&A shareholder wealth effects, the ability of the OLS model to forecast acquirer announcement returns has not yet been formally examined. As such, we do not treat the OLS method as a benchmark against which other models are evaluated, but instead assess its as-yet-unestablished forecasting ability in its own right.

4.1.2 Ridge regression

The Ridge regression (also known as Tikhonov regularization) is a regularized version of the OLS regression. In the Ridge regression, the cost function equals $J(B) = MSE(B) + \alpha\frac{1}{2}\sum_{i=1}^{n}\beta_i^2$. The second term on the right-hand side introduces a penalty for overfitting that is not present in OLS regressions. Ridge regressions have been used in diverse applications, such as economic activity estimation (Exterkate et al., 2016) and microeconomic forecasting (Panagiotelis et al., 2019).

4.1.3 Lasso regression

We also use a least absolute shrinkage and selection operator (Lasso) approach. Like Ridge, Lasso is a regularized version of the OLS regression. First introduced in Tibshirani (1996), it minimizes the residual sum of squares subject to the sum of the absolute value of the coefficients being less than a constant. The Lasso regression cost function is as follows: $J(B) = MSE(B) + \alpha \sum_{i=1}^{n} |\beta_i|$, the second term on the right-hand side of this equation also imposing a penalty for overfitting, in contrast with OLS. Lasso is widely-used for forecasting and has been successful for predicting product returns (Cui et al., 2020), medical outcomes (Lantzy and Anderson, 2020), and microeconomic forecasting (Smeekes and Wijler, 2018).

4.1.4 Random Forest (RF) Regression

Random Forests are a comparably recent addition to regression and classification models and gained popularity since Breiman's seminal paper (Breiman, 2001). RF, a nonparametric approach, performs particularly well in the presence of nonlinearity and complex interactions, and is not prone to overfitting. A drawback of RF is that it is computationally taxing. RF has shown strong out-of-sample accuracy in finance applications such as credit risk approximation (Mercadier and Lardy, 2019). It has also been proven an effective forecasting method in other fields. In a review study, Couronné et al. (2018) compare RF with linear regressions in 243 real high-quality datasets, and observe that, in general, random forests outperform linear regressions.

4.1.5 k-Nearest Neighbor (KNN) Regression

KNN is a simple, computationally inexpensive and intuitively appealing method that can be deployed in regression and classification tasks. KNN has one clear advantage, namely its ability to deal with complex nonlinear behavior (Yankov et al., 2006). KNN is very popular in the forecasting literature and has been applied in diverse fields, e.g. to estimate cancer survival (Bjarnadottir et al., 2018), predict mortgage delinquency (Chen et al., 2021) and forecast wind power (Mangalova and Agafonov, 2014).

4.1.6 Light Gradient-Boosting Machine (LGBM)

Gradient-boosting starts with a simple model that makes predictions. Each subsequent model is trained to correct the errors of the combined previous models. The final model is a weighted sum of all individual models. This ensemble approach helps reduce overfitting while improving predictive performance. Developed in 2017, LGBM is reported to have a similar performance as other more computationally expensive gradient-boosting methods, but with a speed improvement of one order of magnitude (Ke et al., 2017). LGBM has been applied to credit scoring (Liu et al., 2022) and for predicting short-term wind power (Li et al., 2023) and carbon market volatility (Zhu et al., 2023).

4.2 Cross-validation and hyperparameter selection

We deploy state-of-the-art ML methodology similar to the approaches used by recent papers applying ML in an accounting or finance context (Mercadier and Lardy, 2019; Cao and You, 2024; van Binsbergen et al., 2020; Wainer and Cawley, 2021; Chen et al., 2022). We carry out all analyses in Python's scikit-learn package (Pedregosa et al., 2011).

We start by ordering the dataset chronologically, from older to more recent M&A. We then create the traditional training/validation/testing datasets using a fixed-size rolling window (Tashman, 2000). For a given year *n+1*, we use M&A deals in the previous n years as the training/validation dataset for the purpose of selecting optimal hyperparameters, with n set equal to five. The deals in year *n+1* are then used for testing. For instance, for forecasting the acquirer CAR for the year 1997, data for the following years are used for training/validation: 1992, 1993, 1994, 1995, 1996.

The training/validation dataset is subjected to a hyperparameter optimization procedure. A description of the candidate hyperparameters for each method (except OLS, for which there are no hyperparameters) can be found in Table 5. We use a Random Search procedure for this purpose (Bergstra and Bengio, 2012).

Table 5: Candidate hyperparameters for model selection

Model	**Candidate hyperparameters**
Ridge Regression	$\lambda \in \{0.05, 0.06, \ldots, 2.05\}$
Lasso Regression	$\lambda \in \{0.05, 0.06, \ldots, 2.05\}$
	Number of trees: $\{100, 500, 1000\}$
Random Forest (RF)	Max leaf samples: $\{1, 3, 5, 7, 11, 15, 20, 21\}$
	Max depth: $\{1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25\}$
K-Nearest Neighbors (KNN)	$k \in \{1, 2, \ldots, 100\}$
	Learning rate: $\{0.001, 0.005, 0.01, 0.1\}$
LightGBM (LGBM)	Max tree leaves: $\{2, 5, 7, 10, 20\}$
	Max depth: $\{2, 5, 7, 10, 20\}$

Notes: Summary of the candidate hyperparameters used for each model during random search optimization. The procedure uses 50 iterations and five-fold cross-validation within a rolling training/validation set. Hyperparameters not listed were kept at `scikit-learn`'s default values. Models are trained on past data only, using a chronologically-ordered rolling window to preserve the temporal structure of M&A events and avoid information leakage.

All models are allowed 50 iterations when choosing the hyperparameters. We set all other hyperparameters required for the various models at scikit-learn package's default values, consistent with common practice in the ML literature (Fitzpatrick and Mues, 2016). For the evaluation of each set of hyperparameters, we carry out a five-fold cross-validation within the training/validation dataset. We then test the respective models with their optimal hyperparameters in the testing dataset.

The next step is to move according to a fixed-size rolling window - the data from periods 2 to $n+1$ become the next training/validation dataset, and the testing dataset is updated to period $n+2$. The aforementioned steps are repeated until the end of the rolling window, which coincides with the end of our sample period.

Our approach is similar to a traditional nested cross-validation, but has two main advantages over the random partition and rotation of the training/validation samples associated with the latter procedure (Cao and You, 2024). Firstly, by working from a chronologically-ordered set of M&A deals, the training/validation/testing procedure recognizes the intertemporal nature of the events in our dataset, thereby preventing future events from being used to model stock price reactions to past events. Second, by gradually shifting the training/validation set and thereby updating the data, our approach recognizes that the determinants of acquirer stock returns may change over time (Alexandridis et al., 2012, 2017), for example due to changes in investor sensitivities, macroeconomic characteristics and technology.

In line with common practice in ML (Alonso-Robisco and Carbó, 2022), we conclude by scaling all independent variables using scikit-learn's standard scaler. To prevent information leakage, for each observation in any given testing set, normalisation uses the averages and standard deviations calculated from the training/validation datasets.

5 Forecasting acquirer announcement returns

In line with Campbell and Thompson (2008), we rely on the out-of-sample R^2, labeled R^2_{OS}, to evaluate the goodness of fit of the forecasting methods. The R^2_{OS} differs from the in-sample R^2 in two ways. First, while the in-sample R^2 relies on the same data for model testing and prediction purposes, the R^2_{OS} measures the predictive power of a given forecasting approach on unseen test data (Campbell and Thompson, 2008). Second, while both the in-sample and out-of-sample R^2 have a maximum value of one, the in-sample R^2 is by construction always higher than zero (because in-sample models are optimized to minimize errors), but the R^2_{OS} can be negative. A positive R^2_{OS} indicates that the model has at least some predictive power beyond simply using the average value of the dependent variable as a forecast. For example, an R^2_{OS} value of 2% implies that the forecasting model reduces the mean squared forecast error by 2% compared to just using the average CAR value from the training set as a forecast. Conversely, a negative R^2_{OS} suggests that the historical average would have provided a more reliable forecast than the model being tested.

Table 6 outlines the R^2_{OS} for the different forecasting approaches. As expected given the high noise-to-signal ratio in daily excess stock returns, the magnitudes of the R^2_{OS} are very modest, never exceeding 2.5%. Among the models, the three nonparametric approaches have a (limited) predictive ability, as indicated by their positive R^2_{OS}, with LGBM winning the horse race with an R^2_{OS} of 2.13%.

Conversely, the three parametric approaches (Lasso, Ridge, and OLS) have negative R^2_{OS}, indicating they perform worse than just taking the average CAR as a prediction.

A potential explanation for the dominance of nonparametric methods is that the performance disparity between parametric versus nonparametric model families tends to increase as the number of impertinent features rises (Athey and Imbens, 2019). Considering the high noise-to-signal ratio in financial forecasting, it is reasonable to anticipate a high prevalence of irrelevant independent variables in our research context. Models such as OLS are typically less successful in such a scenario, as they have no effective way to ignore potentially less relevant information. However, this is likely to be only a partial explanation, because Lasso (a parametric model) has redundant feature reduction built in, yet performs poorly. A second,

Table 6: Out-of-sample R^2 for the different forecasting methods tested

	R^2_{OS}
LGBM	2.13%
RF	1.23%
KNN	1.11%
Lasso	-0.32%
Ridge	-12.77%
OLS	-20.96%

Notes: Out-of-sample R^2, from larger to smaller. OLS is a linear regression using Ordinary Least Squares as estimator. Ridge and Lasso are regularized linear regressions using Ridge and Lasso approaches, respectively. RF is a Random Forest model. KNN is a K-Nearest Neighbor model, while LGBM refers to Light Gradient-Boosting Machine.

non-mutually exclusive explanation could be that the true relation between acquirer CAR and many of the features considered in our analysis is nonlinear, thereby giving the three nonparametric approaches (none of which presumes a linear relation between dependent and independent variables) an edge.

A plausible interpretation for LGBM's superiority over the other two nonparametric methods, in turn, is its ability to sequentially build trees, with each tree learning from the forecasting errors in the previous trees. The other two approaches either have no in-built learning (KNN)[12] or a more rudimentary tree-building procedure that is perhaps less suitable to capturing complex interactions (RF).

We next perform a series of robustness tests on the baseline findings reported in Table 6. As a first step, we examine the consistency of these findings across the sample period. A priori, we do not have strong expectations regarding the consistency of best- and worst-performing models across years, as both the magnitude and the drivers of acquirer stock price reactions may vary over time, for example, due to shifts in investor sentiment toward corporate deal-making. Figure 1 and Table 7 show the forecasting results on a year by year basis. We note the negative maximum R^2_{OS} values in two sample years (2006 and 2019), indicating that even nonparametric methods have no forecasting ability in these years. In all other years, we find that the best-performing method has a positive R^2_{OS}, suggesting some forecasting ability.

The outperformance of nonparametric over parametric approaches is fairly robust throughout the testing period. More particularly, in 19 out of the 26 sample years, a nonparametric method wins the forecasting horse race. Interestingly, we observe that KNN performs relatively well during turbulent periods such as the start of the COVID pandemic (2020) and the Global Financial Crisis (2008 and 2009). We tentatively argue that this may reflect KNN's ability to adapt quickly to regime shifts and structural breaks, as it makes no assumptions about the global structure of the data and instead relies on local similarity. In contrast, more complex models like LGBM and RF may rely too strongly on pre-crisis patterns and fail to recognize when underlying return drivers change dramatically. The forecasting process of KNN in these periods may also mirror investor behavior, as market participants perhaps rely more on analogical (nearest-neighbor) reasoning and less on fundamental valuation during times of uncertainty.

[12] KNN considers all independent variables when forecasting and does not learn which ones are more relevant for predicting the dependent variable.

Table 7: Out-of-sample R^2 for the different forecasting methods tested by year

Years	Model		R^2_{OS}	
	Maximum	Minimum	Maximum	Minimum
1997	LGBM	OLS	2.02%	-1.28%
1998	KNN	OLS	2.72%	-6.50%
1999	KNN	Lasso	1.21%	-0.10%
2000	KNN	OLS	2.43%	-3.55%
2001	LGBM	OLS	6.48%	-7.92%
2002	KNN	OLS	1.17%	-38.81%
2003	LGBM	OLS	1.21%	-7.90%
2004	KNN	Ridge	1.99%	-1.02%
2005	RF	OLS	11.84%	-10.13%
2006	LGBM	KNN	-0.30%	-2.98%
2007	Lasso	OLS	0.02%	-233.55%
2008	KNN	OLS	0.16%	-2.40%
2009	KNN	Ridge	2.29%	-1.11%
2010	KNN	OLS	0.52%	-3.65%
2011	Ridge	Lasso	4.05%	-0.07%
2012	KNN	OLS	1.55%	-2.88%
2013	LGBM	OLS	2.33%	-8.34%
2014	OLS	Lasso	4.92%	-0.69%
2015	OLS	Lasso	9.24%	-0.13%
2016	OLS	RF	2.94%	-1.61%
2017	Lasso	OLS	0.01%	-5.87%
2018	KNN	RF	1.70%	-1.56%
2019	Lasso	OLS	-1.80%	-14.2%
2020	KNN	OLS	0.94%	-309.46%
2021	LGBM	OLS	4.56%	-0.45%
2022	LGBM	Ridge	1.75%	-4.74%

Notes: Maximum and minimum R^2_{OS} and associated models for each of the 26 years forecasted, from 1997 to 2022. OLS is a linear regression using Ordinary Least Squares as estimator. Ridge and Lasso are regularized linear regressions using Ridge and Lasso approaches, respectively. RF is a Random Forest model. KNN is a K-Nearest Neighbor model, while LGBM refers to Light Gradient-Boosting Machine.

Figure 1: R^2_{OS} across the years

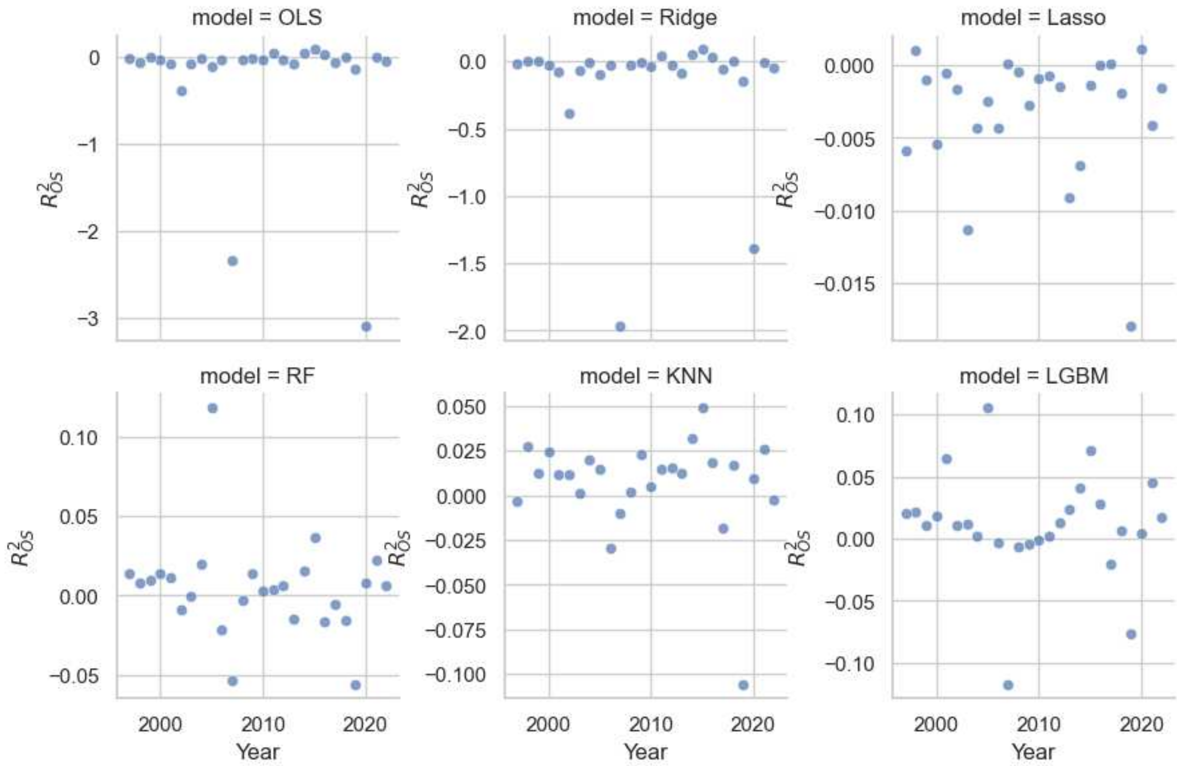

Notes: This figure presents the out-of-sample R^2_{os} values by year for six predictive models: OLS, Ridge, Lasso, Random Forest (RF), K-Nearest Neighbors (KNN), and LightGBM (LGBM). The metric R^2_{os} captures each model's predictive accuracy relative to a naïve benchmark. A positive value indicates better performance than the benchmark, while a negative value indicates worse performance.

In a next series of robustness tests, we examine whether better forecasting performance can be achieved by focusing on relevant subsets of M&A deals. We report the detailed results of these robustness tests in the Supplementary Materials for this paper. We first split the sample into deals paid fully in cash (All Cash dummy variable equal to one) and deals involving a portion of stock payment (All Cash dummy variable equal to zero). The forecasting results for these two subsamples resemble the pattern of the baseline results: nonparametric methods exhibit limited forecasting ability, while parametric methods show none. Next, we exclude public targets from the sample and repeat the analysis for private targets only. The results of this analysis again align closely with the baseline findings.[13] The lack of improvement in the forecastability of announcement returns for subsets of deals is not surprising. The reason is that the ML methods we use are already very apt at implicitly considering relevant subsets of deals, by creating splits (or tree branches) based on combinations of relevant features.

Subsequently, we repeat the baseline analysis with the addition of a “Completed” dummy variable, indicating whether a deal was completed (i.e., not withdrawn). Although investors are unlikely to know at the time of the deal announcement which deals will eventually be completed—hence our decision to exclude this variable from the baseline analysis—the Completed dummy variable may nonetheless proxy for investor perceptions of deal quality or anticipated regulatory opposition. As shown in the Supplementary Materials, we find that including this additional variable does not materially improve the

[13]The sample of public targets, constituting only 17% of the total sample, is too small for a separate analysis.

forecasting performance of any of the approaches. In an unreported in-sample analysis, we also find no significant effect of the Completed dummy on the acquirer's CAR.

While announcement-period returns are by far the most popular M&A performance metric, with over 90% of relevant studies relying on this measure (Ben-David et al., 2025), some studies use long-term stock returns following M&A announcements to capture M&A performance (Loughran and Vijh, 1997; Rau and Vermaelen, 1998; Campbell et al., 2024). However, as several authors have outlined, long-term stock returns are even more problematic than announcement-period returns, due to the potential confounding effects of other corporate events over longer windows, as well as the technical difficulties involved in measuring long-term abnormal stock price performance (Andrade et al., 2001; Billett et al., 2011; Malmendier et al., 2018).[14]For completeness, in a final additional test, we replicate the baseline analysis using acquirer CAR measured over the extended window from trading day +2 to trading day +60 following the announcement. We again present the results in the Supplementary Materials. We find that only LGBM and KNN exhibit some forecasting ability, albeit extremely limited, while the other approaches fail to outperform historical averages.

[14]Previous papers have also used accounting measures, post-merger diversification decisions, and subjective evaluations by corporate managers as alternative measures for M&A outcomes (Schoenberg, 2006; Ambrosini et al., 2011; Meschi and Métais, 2015; Avetisyan et al., 2020). As outlined in Bruner (2002), each of these measures has its own disadvantages.

6 Variable importance analysis for acquirer announcement returns

Having assessed the overall forecasting ability of the different methods, we next examine the features that are most important in driving the predicted stock price reactions. This represents an additional contribution of our paper, as, to the best of our knowledge, no prior studies have conducted a variable importance analysis in the context of out-of-sample forecasting of acquirer announcement returns.

We use SHAP values to better understand the importance of each independent variable for the three models with a positive R^2_{OS} - LGBM, KNN, and RF. SHAP values are derived from the concept of Shapley values in cooperative game theory (Shapley, 1997). In broad terms, the SHAP approach assesses the marginal contribution of each feature by calculating the prediction with and without the variable, and computing the difference between these two results (Lundberg and Lee, 2017).

SHAP values are widely recognised in the field of explanatory ML as one of the best methods for assessing variable importance. SHAP has been employed in previous research to enhance interpretability across diverse domains, including the forecasting of gold prices (Jabeur et al., 2021), sales volumes of non-fungible tokens (NFT) (Teplova et al., 2023), and credit scores (Liu et al., 2024).

Table 8 shows the average individual and cumulative SHAP values for the three nonparametric methods. For LGBM, the best-performing method, three variables, i.e. Relative Deal Size, (acquirer) Size, and R&D Intensity, have a cumulative average SHAP value that is higher than that of all remaining variables collectively. Whilst there is considerable overlap between LGBM and RF regarding the most important variables (both have an identical top three), KNN's results are somewhat different, although the acquirer Size also occupies a second place for that method, similar to the other two approaches.

Interestingly, many of the top-ranked variables are commonly employed as control variables in previous research on acquirer stock price reactions, rather than serving as focal variables of interest. Other commonly-used features such as acquirer Experience, Hostile Deal, and Merger of Equals, have consistently low SHAP values across the models.

Table 9 outlines the most important variables across all testing periods, from 1997 to 2022. We conclude that the important features identified by SHAP analysis are remarkably consistent over time.

Table 8: Summary of average SHAP values

	RF				KNN				LGBM		
Feature	SHAP	$\sum$ SHAP	$\sum$ SHAP (%)	Feature	SHAP	$\sum$ SHAP	$\sum$ SHAP (%)	Feature	SHAP	$\sum$ SHAP	$\sum$ SHAP (%)
Relative Deal Size	0.002	0.002	11.52%	Federal Funds Rate	0.002	0.002	8.71%	Relative Deal Size	0.003	0.003	22.83%
Size	0.002	0.004	22.33%	Size	0.002	0.004	15.88%	Size	0.002	0.005	43.20%
R&D Intensity	0.001	0.005	29.15%	Public Target	0.001	0.005	21.71%	R&D Intensity	0.001	0.006	54.27%
Federal Funds Rate	0.001	0.005	34.16%	High Tech Industry	0.001	0.006	26.56%	Public Target	0.001	0.007	61.76%
Cash	0.001	0.006	38.88%	GDP Change	0.001	0.007	30.83%	Cash	0.001	0.008	69.12%
Deal Size	0.001	0.007	42.98%	Risk Sentiment	0.001	0.008	35.00%	Deal Size	0.001	0.009	73.55%
Merger Wave	0.001	0.007	47.06%	Industry Relatedness	0.001	0.009	39.12%	Altman Z	0.000	0.009	76.43%
Public Target	0.001	0.008	50.78%	Cash	0.001	0.010	43.11%	All Stock	0.000	0.009	79.01%
Dormant Period	0.001	0.009	54.37%	R&D Intensity	0.001	0.011	47.10%	Merger Wave	0.000	0.010	81.43%
Risk Sentiment	0.001	0.009	57.59%	Merger Wave	0.001	0.011	50.94%	Return on Assets	0.000	0.010	83.80%
Altman Z	0.000	0.010	60.74%	International Target	0.001	0.012	54.70%	Risk Sentiment	0.000	0.010	86.11%
Free Cash Flow	0.000	0.010	63.87%	VIX	0.001	0.013	58.31%	Leverage	0.000	0.011	88.25%
EPU 3C	0.000	0.011	66.90%	Winter Deal	0.001	0.014	61.70%	Free Cash Flow	0.000	0.011	90.34%
VIX	0.000	0.011	69.92%	All Cash	0.001	0.015	65.06%	Market to Book	0.000	0.011	92.36%
Return on Assets	0.000	0.012	72.90%	Election	0.001	0.015	68.39%	Dormant Period	0.000	0.011	94.37%
Leverage	0.000	0.012	75.88%	All Stock	0.001	0.016	71.66%	Federal Funds Rate	0.000	0.012	96.21%
Experience	0.000	0.012	78.82%	Relative Deal Size	0.001	0.017	74.81%	Industry Relatedness	0.000	0.012	96.93%
Market to Book	0.000	0.013	81.49%	Leverage	0.001	0.018	77.88%	VIX	0.000	0.012	97.62%
EPU News	0.000	0.013	84.09%	Bankruptcy Risk	0.001	0.018	80.86%	EPU 3C	0.000	0.012	98.19%
All Stock	0.000	0.014	86.39%	EPU 3C	0.001	0.019	83.73%	High Tech Industry	0.000	0.012	98.70%
High Tech Industry	0.000	0.014	88.63%	EPU News	0.001	0.019	86.46%	GDP Change	0.000	0.012	99.20%
Industry Relatedness	0.000	0.014	90.72%	Experience	0.001	0.020	89.14%	EPU News	0.000	0.012	99.68%
GDP Change	0.000	0.015	92.56%	Dormant Period	0.000	0.020	91.08%	All Cash	0.000	0.012	99.88%
International Target	0.000	0.015	94.17%	Return on Assets	0.000	0.021	92.73%	International Target	0.000	0.012	99.98%
All Cash	0.000	0.015	95.48%	Free Cash Flow	0.000	0.021	94.38%	Experience	0.000	0.012	99.99%
Bankruptcy Risk	0.000	0.015	96.60%	Friday Deal	0.000	0.022	95.93%	Winter Deal	0.000	0.012	100.00%
Winter Deal	0.000	0.015	97.60%	Percentage Sought	0.000	0.022	96.93%	Election	0.000	0.012	100.00%
Election	0.000	0.016	98.45%	Deal Size	0.000	0.022	97.93%	Friday Deal	0.000	0.012	100.00%
Percentage Sought	0.000	0.016	99.01%	Altman Z	0.000	0.022	98.91%	Bankruptcy Risk	0.000	0.012	100.00%
Friday Deal	0.000	0.016	99.52%	Market to Book	0.000	0.022	99.45%	Hostile Deal	0.000	0.012	100.00%
Pandemic	0.000	0.016	99.86%	Pandemic	0.000	0.022	99.83%	Merger of Equals	0.000	0.012	100.00%
Hostile Deal	0.000	0.016	99.96%	Hostile Deal	0.000	0.022	99.96%	Pandemic	0.000	0.012	100.00%
Merger of Equals	0.000	0.016	100.00%	Merger of Equals	0.000	0.022	100.00%	Percentage Sought	0.000	0.012	100.00%

Notes: This table presents variable importance for three nonparametric methods, as measured by SHAP values. RF is a Random Forest model. KNN is a K-Nearest Neighbors model, while LGBM refers to Light Gradient-Boosting Machine. $\sum$ SHAP presents the cumulative sum of the SHAP values. $\sum$ SHAP (%) is the cumulative sum expressed as percentage of Max ($\sum$ SHAP).

Table 9: Most important acquirer *CAR[-1,+1]* predictor for three non-parametric models

Year			feature
	RF	KNN	LGBM
1997	Size	Winter Deal	Size
1998	Size	Risk Sentiment	Size
1999	Size	Cash	Size
2000	Size	GDP Change	Size
2001	Size	Risk Sentiment	Size
2002	Size	Interest Rate	Size
2003	Size	Interest Rate	Size
2004	Experience	Size	Size
2005	Size	Public Target	Size
2006	Interest Rate	Interest Rate	Relative Deal Size
2007	R&D Intensity	GDP Change	R&D Intensity
2008	R&D Intensity	R&D Intensity	Relative Deal Size
2009	R&D Intensity	Merger Wave	R&D Intensity
2010	R&D Intensity	Interest Rate	R&D Intensity
2011	R&D Intensity	International Target	Relative Deal Size
2012	R&D Intensity	International Target	R&D Intensity
2013	R&D Intensity	Industry Relatedness	R&D Intensity
2014	EPU 3C	GDP Change	Relative Deal Size
2015	Relative Deal Size	Relative Deal Size	Relative Deal Size
2016	Relative Deal Size	Interest Rate	Relative Deal Size
2017	Relative Deal Size	Interest Rate	Relative Deal Size
2018	Relative Deal Size	Interest Rate	Relative Deal Size
2019	Size	Risk Sentiment	Size
2020	Size	VIX	Relative Deal Size
2021	Size	International Target	Relative Deal Size
2022	Size	International Target	Size

Notes: This table presents the most influential variables in predicting *CAR[-1,+1]* for three nonparametric models, by year, as evaluated by SHAP values. RF is a Random Forest model. KNN is a K-Nearest Neighbors model, while LGBM refers to Light Gradient-Boosting Machine.

7 Summary, practical implications and avenues for future research

7.1 Summary of findings

M&A transactions can create or destroy value, sometimes to the tune of billions of US dollars. Acquirer stock price reactions to M&A announcements have been the focus of many academic studies and are often mentioned in the business press alongside other key deal characteristics (Andrade et al., 2001; King et al., 2004; Jaffe et al., 2013; Renneboog and Vansteenkiste, 2019; Hu et al., 2020). However, so far, the literature's focus has been on in-sample analysis of the magnitude and cross-sectional determinants of these returns. We, instead, address the novel question of whether investor perceptions of M&A deals, as captured by acquirer stock price reactions to M&A announcements, are in any way forecastable using pre-announcement, publicly-available information. We also investigate which variables are most relevant in predicting these stock price reactions.

Our sample, constructed using standard screening criteria (Netter et al., 2011; Jaffe et al., 2013), consists of 9,517 M&A announcements by US public acquirers between 1992 and 2022. As independent variables, we use standard acquirer stock price reaction determinants considered by a range of previous studies (Moeller et al., 2004; Harford and Li, 2007; Ishii and Xuan, 2014; Becht et al., 2016; Elnahas and Kim, 2017; Adra et al., 2020). We evaluate three parametric methods (OLS, Ridge, and Lasso) and three nonparametric methods (RF, KNN, and LGBM) and follow state-of-the-art ML methodology for hyperparameter selection and cross-validation, making appropriate adjustments for the intertemporal nature of our dataset.

Our evidence suggests that the three nonparametric models have some ability to forecast acquirer stock price reactions, compared with simply using average acquirer CAR as a prediction. As expected given the high noise-to-signal ratio of daily abnormal stock returns, the forecasting power of even the best-performing method (LGBM) is modest. However, given the large average market value of acquirers, a relatively small out-of-sample accuracy may still yield significantly improved outcomes. We reiterate that the main contribution of our work does not lie in improving the forecastability achieved in previous studies, but rather in addressing the question of forecastability of acquirer stock price reactions—by any method—in and of itself, as this remains largely unexplored territory.

By contrast, we find that parametric models do not have any predictive power for acquirer stock price reactions. Although it is impossible to know ex ante what models work best for any given prediction problem, theory points to two ex post explanations for the superior performance of nonparametric methods in our research context. First, nonparametric approaches are more suited at ignoring irrelevant independent variables. Given the high noise-to-signal ratio in abnormal stock returns, we should expect a low number of relevant features in our research context, which gives nonparametric approaches an edge over their parametric counterparts (Athey and Imbens, 2019). Second, broadly speaking, nonparametric approaches are better at dealing with nonlinear effects of independent variables and with interaction effects between independent variables. It is quite plausible that there are indeed nonlinearities in the impacts of some of the features, notably size-related covariates, on acquirer announcement returns. For example, there may be economies and subsequent diseconomies of scale associated with M&A deals, leading to a nonlinear effect of size-related variables on M&A outcomes.

Another interesting and unexpected finding, which resulted from a SHAP variable importance analysis, is that only a handful of features are actually useful for predicting acquirer stock price reactions. Interestingly, most of these variables are "bread and butter" features commonly used as control variables in the literature, such as acquirer size and (relative) deal size. The small number of variables consistently deemed important for prediction purposes opens the possibility for more parsimonious prediction models for acquirer announcement returns.

7.2 Practical implications

The key question the reader may be asking is: *"Should managers use the present framework to decide on which M&A deals to pursue?"*. Given the modest forecastability of acquirer announcement returns uncovered by our tests, we would not foresee these forecasts to be the sole driver of M&A decisions. However, we do hope that our paper will encourage corporate decision makers, among which managers and their advisors, to at least consider forecasted acquirer *CAR* along with other indicators of deal suitability (Shi et al., 2017). From anecdotal evidence, we know that the selection of potential M&A targets consumes substantial corporate resources. Large publicly-quoted companies typically have a dedicated target selection team tasked with evaluating myriad potential target firms each month.[15] Given the substantial size of M&A deals and the nontrivial repercussions of engaging in bad deals, we would argue that even a modest forecastability of acquirer stock price reactions, as achieved through the nonparametric models studied in this paper, may lead to material improvements for various stakeholders, including not only the managers of the acquiring firm but also its supply chain partners and employees. Having some insight into the anticipated stock price reactions to an M&A deal may also be useful for target companies entering the negotiation process.

Of course, we acknowledge that acquiring-firm managers may still choose to pursue M&A deals that are predicted to result in a negative stock price reaction. Nonetheless, we believe it is valuable for them to have ex ante knowledge of the likely stock market reception of a deal. For example, managers who intend to proceed with a deal that has a negative forecasted *CAR* may wish to devote more effort to providing a comprehensive justification and quantification of the associated synergies, given that they have substantial discretion over the level of synergy detail disclosed to the market (Dutordoir et al., 2014).

We refrain from making strong claims about the trading implications of our findings, because our models estimate acquirer abnormal stock returns conditional on a given acquirer–target combination being announced. Since investors are unlikely to anticipate the exact timing of these announcements, they cannot directly trade on our forecasts. Nevertheless, our results are still relevant for investors, for the following two reasons. First, equity investors may wish to better understand the mechanisms driving stock price reactions to M&A announcements. Our findings show that nonparametric approaches

[15] As an illustration, please see `https://tinyurl.com/2s49ewae`.

outperform traditional parametric models, suggesting that market participants incorporate M&A-related information in a nonlinear, nonparametric fashion. This insight likely extends to other corporate events, such as seasoned equity offerings or dividend changes, where empirical studies still rely largely on in-sample OLS tests. Incorporating nonlinear specifications or ML methods may similarly enhance predictive accuracy in these settings.

Second, prior research documents that bidder–target combinations can be forecasted to some extent (Rodrigues and Stevenson, 2013; Song and Walkling, 2000; Futagami et al., 2021; Tunyi, 2021). If investors can anticipate forthcoming acquirer–target matches, our acquirer CAR forecasts could be used ex ante to take long (short) positions in prospective acquirers with predicted positive (negative) announcement returns that outweigh trading costs. Jointly modeling bidder–target combinations and the corresponding announcement returns is likely to present methodological and empirical challenges, but may nonetheless be within the capabilities of sophisticated investors.

7.3 Limitations and directions for future research

Our study has some limitations that may inspire future research. First, by design, we restricted the set of independent variables to publicly-available information from standard data sources. Consequently, our results may represent a lower bound on the forecastability of acquirer announcement returns. Future research could explore whether incorporating less accessible variables, such as CEO narcissism, overconfidence, or political orientation (Grinstein and Hribar, 2004; Billett and Qian, 2008; Elnahas and Kim, 2017; Ham et al., 2018), improves forecasting accuracy. However, a potential trade-off is that these non-standard variables are typically available only for a subset of firms, which could reduce the statistical power and generalizability of such forecasting tests (Petropoulos et al., 2022).

Second, beyond the acquirer stock returns analyzed in this paper, investor perceptions of M&A announcements could also be captured using alternative proxies, such as option-implied volatility, Google search volume, or social media activity related to the M&A deal. We encourage future research to examine the forecastability of these alternative metrics. Future work could also extend our analysis to non-US markets and explore different forecasting methodologies for predicting acquirer announcement returns.[16]

Finally, we encourage future academic research to investigate the predictability of stock price reactions to other major corporate decisions, with an emphasis on out-of-sample forecasting. Corporate events such as dividend changes, security offerings, and divestitures have been extensively examined using traditional event study methods. Assessing the out-of-sample forecastability of stock price reactions to these announcements, and identifying the key variables driving forecasting performance, would be a valuable extension.

Data availability statement: The data used for this study are under license from the SDC (LSEG Workspace) and WRDS databases.

[16] We welcome requests from researchers wishing to replicate or build upon our results. To facilitate reproducibility, we have deliberately used open-source software.

References

Abed, A. and W. Abdallah (2017). Does cross-listing increase managers' propensity to listen to the market in M&A deals? *Journal of Corporate Finance 46*, 97–120.

Adra, S., L. G. Barbopoulos, and A. Saunders (2020). The impact of monetary policy on M&A outcomes. *Journal of Corporate Finance 62*, 101529.

Ahern, K. R. and D. Sosyura (2014). Who writes the news? Corporate press releases during merger negotiations. *The Journal of Finance 69*(1), 241–291.

Alexandridis, G., N. Antypas, and N. Travlos (2017). Value creation from M&As: New evidence. *Journal of Corporate Finance 45*, 632–650.

Alexandridis, G., K. P. Fuller, L. Terhaar, and N. G. Travlos (2013). Deal size, acquisition premia and shareholder gains. *Journal of Corporate Finance 20*, 1–13.

Alexandridis, G., C. F. Mavrovitis, and N. G. Travlos (2012). How have M&As changed? Evidence from the sixth merger wave. *The European Journal of Finance 18*(8), 663–688.

Alonso-Robisco, A. and J. M. Carbó (2022). Can machine learning models save capital for banks? Evidence from a Spanish credit portfolio. *International Review of Financial Analysis 84*, 102372.

Altman, E. I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance 23*(4), 589–609.

Ambrosini, V., C. Bowman, and R. Schoenberg (2011). Should acquiring firms pursue more than one value creation strategy? An empirical test of acquisition performance. *British Journal of Management 22*(1), 173–185.

Andrade, G., M. Mitchell, and E. Stafford (2001). New evidence and perspectives on mergers. *Journal of Economic Perspectives 15*(2), 103–120.

Asquith, P., R. F. Bruner, and D. W. Mullins Jr (1983). The gains to bidding firms from merger. *Journal of Financial Economics 11*(1-4), 121–139.

Athey, S. and G. W. Imbens (2019). Machine learning methods that economists should know about. *Annual Review of Economics 11*, 685–725.

Avetisyan, E., Y. Baruch, P.-X. Meschi, E. Metais, and A. Norheim-Hansen (2020). Tying the acquirer's human resource management quality to cross-border acquisition divestment probability: Curvilinear connection with slacklining. *British Journal of Management 31*(3), 568–588.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Bayless, M. and S. Chaplinsky (1991). Expectations of security type and the information content of debt and equity offers. *Journal of Financial Intermediation 1*(3), 195–214.

Becht, M., A. Polo, and S. Rossi (2016). Does mandatory shareholder voting prevent bad acquisitions? *The Review of Financial Studies 29*(11), 3035–3067.

Ben-David, I., U. Bhattacharya, and S. E. Jacobsen (2025). The (missing) relation between announcement returns and value creation. *Journal of Finance, forthcoming*.

Bergstra, J. and Y. Bengio (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research 13*(2), 281–305.

Betton, S., B. E. Eckbo, and K. S. Thorburn (2008). Corporate takeovers. In B. E. Eckbo (Ed.), *Handbook of Empirical Corporate Finance*, pp. 291–429. Elsevier.

Bhagwat, V., R. Dam, and J. Harford (2016). The real effects of uncertainty on merger activity. *The Review of Financial Studies 29*(11), 3000–3034.

Billett, M. T., M. J. Flannery, and J. A. Garfinkel (2011). Frequent issuers' influence on long-run post-issuance returns. *Journal of Financial Economics 99*(2), 349–364.

Billett, M. T. and Y. Qian (2008). Are overconfident CEOs born or made? Evidence of self-attribution bias from frequent acquirers. *Management Science 54*(6), 1037–1051.

Bjarnadottir, M., D. Anderson, L. Zia, and K. Rhoads (2018). Predicting colorectal cancer mortality: Models to facilitate patient-physician conversations and inform operational decision making. *Production and Operations Management 27*(12), 2162–2183.

Bollaert, H. and M. Delanghe (2015). Securities Data Company and Zephyr, data sources for M&A research. *Journal of Corporate Finance 33*, 85–100.

Bonaime, A., H. Gulen, and M. Ion (2018). Does policy uncertainty affect mergers and acquisitions? *Journal of Financial Economics 129*(3), 531–558.

Bozos, K. and K. Nikolopoulos (2011). Forecasting the value effect of seasoned equity offering announcements. *European Journal of Operational Research 214*(2), 418–427.

Bradley, M., A. Desai, and E. H. Kim (1988). Synergistic gains from corporate acquisitions and their division between the stockholders of target and acquiring firms. *Journal of Financial Economics 21*(1), 3–40.

Brealey, R. A., S. C. Myers, F. Allen, and P. Mohanty (2018). *Principles of Corporate Finance, 12/e*, Volume 12. McGraw-Hill Education.

Breiman, L. (2001). Random forests. *Machine Learning 45*(1), 5–32.

Brown, L. D., G. D. Richardson, and S. J. Schwager (1987). An information interpretation of financial analyst superiority in forecasting earnings. *Journal of Accounting Research 25*(1), 49–67.

Brown, S. J. and J. B. Warner (1985). Using daily stock returns: The case of event studies. *Journal of Financial Economics 14*(1), 3–31.

Bruner, R. F. (2002). Does M&A pay? A survey of evidence for the decision-maker. *Journal of Applied Finance 12*(1), 48–68.

Buchanan, B. (2019). Artificial intelligence in Finance. Technical report, The Alan Turing Institute.

Cai, J., M. H. Song, and R. A. Walkling (2011). Anticipation, acquisitions, and bidder returns: Industry shocks and the transfer of information across rivals. *The Review of Financial Studies 24*(7), 2242–2285.

Cai, Y. and M. Sevilir (2012). Board connections and M&A transactions. *Journal of Financial Economics 103*(2), 327–349.

Çakmaklı, C. and D. van Dijk (2016). Getting the most out of macroeconomic information for predicting excess stock returns. *International Journal of Forecasting 32*(3), 650–668.

Campbell, J. L., E. Elfrink, F.-H. Huang, and H.-M. Lu (2024). What is the deal?: Predicting M&A outcomes with machine learning. SSRN Working Paper.

Campbell, J. T., D. G. Sirmon, and M. Schijven (2016). Fuzzy logic and the market: A configurational approach to investor perceptions of acquisition announcements. *Academy of Management Journal 59*(1), 163–187.

Campbell, J. Y. and S. B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies 21*(4), 1509–1531.

Cao, K. and H. You (2024). Fundamental analysis via machine learning. *Financial Analysts Journal 80*(2), 74–98.

Chacko, G. C., J. W. Jurek, and E. Stafford (2008). The price of immediacy. *The Journal of Finance 63*(3), 1253–1290.

Chen, S., Z. Guo, and X. Zhao (2021). Predicting mortgage early delinquency with machine learning methods. *European Journal of Operational Research 290*(1), 358–372.

Chen, X., Y. HA (TONY) CHO, Y. Dou, and B. Lev (2022). Predicting future earnings changes using machine learning and detailed financial data. *Journal of Accounting Research 60*(2), 467–515.

Couronné, R., P. Probst, and A.-L. Boulesteix (2018). Random forest versus logistic regression: A large-scale benchmark experiment. *BMC Bioinformatics 19*(1), 1–14.

Croci, E. and D. Petmezas (2015). Do risk-taking incentives induce CEOs to invest? Evidence from acquisitions. *Journal of Corporate Finance 32*, 1–23.

Cui, H., S. Rajagopalan, and A. R. Ward (2020). Predicting product return volume using machine learning methods. *European Journal of Operational Research 281*(3), 612–627.

Datta, D. K., G. E. Pinches, and V. K. Narayanan (1992). Factors influencing wealth creation from mergers and acquisitions: A meta-analysis. *Strategic Management Journal 13*(1), 67–84.

Deng, X., J.-k. Kang, and B. S. Low (2013). Corporate social responsibility and stakeholder value maximization: Evidence from mergers. *Journal of Financial Economics 110*(1), 87–109.

Dong, M., D. Hirshleifer, S. Richardson, and S. H. Teoh (2006). Does investor misvaluation drive the takeover market? *The Journal of Finance 61*(2), 725–762.

Dutordoir, M., P. Roosenboom, and M. Vasconcelos (2014). Synergy disclosures in mergers and acquisitions. *International Review of Financial Analysis 31*, 88–100.

Dutordoir, M., N. C. Strong, and P. Sun (2022). Does short-selling potential influence merger and acquisition payment choice? *Journal of Financial Economics 144*(3), 761–779.

Dye, R. A. and S. Sridhar (2002). Resource allocation effects of price reactions to disclosures. *Contemporary Accounting Research 19*(3), 385–410.

Eckbo, B. E., T. Makaew, and K. S. Thorburn (2018). Are stock-financed takeovers opportunistic? *Journal of Financial Economics 128*(3), 443–465.

Eden, L., S. R. Miller, S. Khan, R. J. Weiner, and D. Li (2022). The event study in international business research: Opportunities, challenges, and practical solutions. *Journal of International Business Studies 53*(5), 803–817.

Elnahas, A. M. and D. Kim (2017). CEO political ideology and mergers and acquisitions decisions. *Journal of Corporate Finance 45*, 162–175.

Erel, I., L. H. Stern, C. Tan, and M. S. Weisbach (2021). Selecting directors using machine learning. *The Review of Financial Studies 34*(7), 3226–3264.

Exterkate, P., P. J. Groenen, C. Heij, and D. van Dijk (2016). Nonlinear forecasting with many predictors using kernel ridge regression. *International Journal of Forecasting 32*(3), 736–753.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance 25*(2), 383–417.

Ferson, W., S. Nallareddy, and B. Xie (2013). The "out-of-sample" performance of long run risk models. *Journal of Financial Economics 107*(3), 537–556.

Fieberg, C., D. Metko, T. Poddig, and T. Loy (2023). Machine learning techniques for cross-sectional equity returns' prediction. *OR Spectrum 45*(1), 289–323.

Field, L. C. and A. Mkrtchyan (2017). The effect of director experience on acquisition performance. *Journal of Financial Economics 123*(3), 488–511.

Fitzpatrick, T. and C. Mues (2016). An empirical comparison of classification algorithms for mortgage default prediction: Evidence from a distressed mortgage market. *European Journal of Operational Research 249*(2), 427–439.

Fuller, K., J. Netter, and M. Stegemoller (2002). What do returns to acquiring firms tell us? Evidence from firms that make many acquisitions. *The Journal of Finance 57*(4), 1763–1793.

Futagami, K., Y. Fukazawa, N. Kapoor, and T. Kito (2021). Pairwise acquisition prediction with SHAP value interpretation. *The Journal of Finance and Data Science 7*, 22–44.

Goldstein, I. (2023). Information in financial markets and its real effects. *Review of Finance 27*(1), 1–32.

Grinstein, Y. and P. Hribar (2004). CEO compensation and incentives: Evidence from M&A bonuses. *Journal of Financial Economics 73*(1), 119–143.

Grønborg, N. S., A. Lunde, A. Timmermann, and R. Wermers (2021). Picking funds with confidence. *Journal of Financial Economics 139*(1), 1–28.

Gu, S., B. Kelly, and D. Xiu (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies 33*(5), 2223–2273.

Haase, F. and M. Neuenkirch (2023). Predictability of bull and bear markets: A new look at forecasting stock market regimes (and returns) in the US. *International Journal of Forecasting 39*(2), 587–605.

Ham, C., N. Seybert, and S. Wang (2018). Narcissism is a bad sign: CEO signature size, investment, and performance. *Review of Accounting Studies 23*(1), 234–264.

Harford, J. and K. Li (2007). Decoupling CEO wealth and firm performance: The case of acquiring CEOs. *The Journal of Finance 62*(2), 917–949.

Hoang, D. and K. Wiegratz (2023). Machine learning methods in Finance: Recent applications and prospects. *European Financial Management 29*(5), 1657–1701.

Houston, J. F., C. M. James, and M. D. Ryngaert (2001). Where do merger gains come from? Bank mergers from the perspective of insiders and outsiders. *Journal of Financial Economics 60*(2-3), 285–331.

Hu, N., L. Li, H. Li, and X. Wang (2020). Do mega-mergers create value? The acquisition experience and mega-deal outcomes. *Journal of Empirical Finance 55*, 119–142.

Hyndman, R. J. and G. Athanasopoulos (2018). *Forecasting: Principles and practice.* OTexts.

Ishii, J. and Y. Xuan (2014). Acquirer-target social ties and merger outcomes. *Journal of Financial Economics 112*(3), 344–363.

Jabeur, S. B., S. Mefteh-Wali, and J.-L. Viviani (2021). Forecasting gold price with the XGBoost algorithm and SHAP interaction values. *Annals of Operations Research 334*, 1–21.

Jaffe, J., D. Pedersen, and T. Voetmann (2013). Skill differences in corporate acquisitions. *Journal of Corporate Finance 23*, 166–181.

Jennings, R. H. and M. A. Mazzeo (1991). Stock price movements around acquisition announcements and management's response. *Journal of Business*, 139–163.

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *The American Economic Review 76*(2), 323–329.

Jones, S. and D. A. Hensher (2004). Predicting firm financial distress: A mixed logit model. *The Accounting Review 79*(4), 1011–1038.

Junni, P., R. M. Sarala, S. Y. Tarba, and Y. Weber (2015). The role of strategic agility in acquisitions. *British Journal of Management 26*(4), 596–616.

Kališ, R., M. Luptáčik, and D. Dujava (2023). Efficiency effects of mergers: Harmonising merged production. *Journal of the Operational Research Society 74*(5), 1261–1274.

Kanas, A. (2003). Non-linear forecasts of stock returns. *Journal of Forecasting 22*(4), 299–315.

Kau, J. B., J. S. Linck, and P. H. Rubin (2008). Do managers listen to the market? *Journal of Corporate Finance 14*(4), 347–362.

Ke, G., Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu (2017). Lightgbm: A highly efficient gradient boosting decision tree. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Volume 30. Curran Associates, Inc.

King, D. R., D. R. Dalton, C. M. Daily, and J. G. Covin (2004). Meta-analyses of post-acquisition performance: Indications of unidentified moderators. *Strategic Management Journal 25*(2), 187–200.

King, D. R., G. Wang, M. Samimi, and A. F. Cortes (2021). A meta-analytic integration of acquisition performance prediction. *Journal of Management Studies 58*(5), 1198–1236.

Kothari, S. P. and J. B. Warner (2007). Econometrics of event studies. In *Handbook of Empirical Corporate Finance*, pp. 3–36. Elsevier.

Lantzy, S. and D. Anderson (2020). Can consumers use online reviews to avoid unsuitable doctors? Evidence from RateMDs.com and the Federation of State Medical Boards. *Decision Sciences 51*(4), 962–984.

Lessmann, S. and S. Voß (2017). Car resale price forecasting: The impact of regression method, private information, and heterogeneity on forecast accuracy. *International Journal of Forecasting 33*(4), 864–877.

Lewis, C. M., R. J. Rogalski, and J. K. Seward (1999). Is convertible debt a substitute for straight debt or for common equity? *Financial Management 28*(3), 5–27.

Li, H., H. Liu, A. Siganos, and M. Zhou (2016). Bank regulation, financial crisis, and the announcement effects of seasoned equity offerings of US commercial banks. *Journal of Financial Stability 25*, 37–46.

Li, K., F. Mai, R. Shen, and X. Yan (2021). Measuring corporate culture using machine learning. *The Review of Financial Studies 34*(7), 3265–3315.

Li, Y., Z. Wu, and Y. Su (2023). Adaptive short-term wind power forecasting with concept drifts. *Renewable Energy 217*, 119146.

Liberti, J. M. and M. A. Petersen (2019). Information: Hard and soft. *Review of Corporate Finance Studies 8*(1), 1–41.

Liu, W., H. Fan, M. Xia, and M. Xia (2022). A focal-aware cost-sensitive boosted tree for imbalanced credit scoring. *Expert Systems with Applications 208*, 118158.

Liu, Y., F. Huang, L. Ma, Q. Zeng, and J. Shi (2024). Credit scoring prediction leveraging interpretable ensemble learning. *Journal of Forecasting 43*(2), 286–308.

Loughran, T. and A. M. Vijh (1997). Do long-term shareholders benefit from corporate acquisitions? *The Journal of Finance 52*(5), 1765–1790.

Louis, H. (2005). Acquirers' abnormal returns and the non-big 4 auditor clientele effect. *Journal of Accounting and Economics 40*(1-3), 75–99.

Lundberg, S. M. and S.-I. Lee (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems 30*.

Luo, Y. (2005). Do insiders learn from outsiders? Evidence from mergers and acquisitions. *The Journal of Finance 60*(4), 1951–1982.

Makridakis, S. (1990). Note—sliding simulation: A new approach to time series forecasting. *Management Science 36*(4), 505–512.

Makridakis, S., E. Spiliotis, and V. Assimakopoulos (2020). The M4 competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting 36*(1), 54–74.

Makridakis, S., E. Spiliotis, V. Assimakopoulos, A.-A. Semenoglou, G. Mulder, and K. Nikolopoulos (2023). Statistical, machine learning and deep learning forecasting methods: Comparisons and ways forward. *Journal of the Operational Research Society 74*(3), 840–859.

Malmendier, U., E. Moretti, and F. S. Peters (2018). Winning by losing: Evidence on the long-run effects of mergers. *The Review of Financial Studies 31*(8), 3212–3264.

Malmendier, U. and G. Tate (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics 89*(1), 20–43.

Mangalova, E. and E. Agafonov (2014). Wind power forecasting using the k-nearest neighbors algorithm. *International Journal of Forecasting 30*(2), 402–406.

Manne, H. G. (1965). Mergers and the market for corporate control. *Journal of Political Economy 73*(2), 110–120.

Martynova, M. and L. Renneboog (2008). A century of corporate takeovers: What have we learned and where do we stand? *Journal of Banking & Finance 32*(10), 2148–2177.

Martynova, M. and L. Renneboog (2009). What determines the financing decision in corporate takeovers: Cost of capital, agency problems, or the means of payment? *Journal of Corporate Finance 15*(3), 290–315.

Mercadier, M. and J.-P. Lardy (2019). Credit spread approximation and improvement using random forest regression. *European Journal of Operational Research 277*(1), 351–365.

Meschi, P.-X. and E. Métais (2015). Too big to learn: The effects of major acquisition failures on subsequent acquisition divestment. *British Journal of Management 26*(3), 408–423.

Mitchell, M., T. Pulvino, and E. Stafford (2004). Price pressure around mergers. *The Journal of Finance 59*(1), 31–63.

Moeller, S. B. and F. P. Schlingemann (2005). Global diversification and bidder gains: A comparison between cross-border and domestic acquisitions. *Journal of Banking & Finance 29*(3), 533–564.

Moeller, S. B., F. P. Schlingemann, and R. M. Stulz (2004). Firm size and the gains from acquisitions. *Journal of Financial Economics 73*(2), 201–228.

Morck, R., A. Shleifer, and R. W. Vishny (1990). Do managerial objectives drive bad acquisitions? *The Journal of Finance 45*(1), 31–48.

Mulherin, J. H., J. M. Netter, and A. B. Poulsen (2017). The evidence on mergers and acquisitions: A historical and modern report. In *The Handbook of the Economics of Corporate Governance*, Volume 1, pp. 235–290. Elsevier.

Myers, S. C. and N. S. Majluf (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics 13*(2), 187–221.

Netter, J., M. Stegemoller, and M. B. Wintoki (2011). Implications of data screens on merger and acquisition analysis: A large sample study of mergers and acquisitions from 1992 to 2009. *The Review of Financial Studies 24*(7), 2316–2357.

Ngai, E. W. and Y. Wu (2022). Machine learning in marketing: A literature review, conceptual framework, and research agenda. *Journal of Business Research 145*, 35–48.

Nguyen, N. H. and H. V. Phan (2017). Policy uncertainty and mergers and acquisitions. *Journal of Financial and Quantitative Analysis 52*(2), 613–644.

Obaid, K. and K. Pukthuanthong (2022). A picture is worth a thousand words: Measuring investor sentiment by combining machine learning and photos from news. *Journal of Financial Economics 144*(1), 273–297.

Officer, M. S. (2007). The price of corporate liquidity: Acquisition discounts for unlisted targets. *Journal of Financial Economics 83*(3), 571–598.

Ozbayoglu, A. M., M. U. Gudelek, and O. B. Sezer (2020). Deep learning for financial applications: A survey. *Applied Soft Computing 93*, 106–384.

Palepu, K. G. (1986). Predicting takeover targets: A methodological and empirical analysis. *Journal of Accounting and Economics 8*(1), 3–35.

Panagiotelis, A., G. Athanasopoulos, R. J. Hyndman, B. Jiang, and F. Vahid (2019). Macroeconomic forecasting for Australia using a large number of predictors. *International Journal of Forecasting 35*(2), 616–633.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research 12*, 2825–2830.

Pei, J., H. Wang, Q. Peng, and S. Liu (2024). Saving face: Leveraging artificial intelligence-based negative feedback to enhance employee job performance. *Human Resource Management 63*(5), 775–790.

Petropoulos, F., D. Apiletti, V. Assimakopoulos, M. Z. Babai, D. K. Barrow, S. B. Taieb, C. Bergmeir, R. J. Bessa, J. Bijak, J. E. Boylan, et al. (2022). Forecasting: theory and practice. *International Journal of Forecasting*.

Ranta, M., M. Ylinen, and M. Järvenpää (2023). Machine learning in management accounting research: Literature review and pathways for the future. *European Accounting Review 32*(3), 607–636.

Rau, P. R. and T. Vermaelen (1998). Glamour, value and the post-acquisition performance of acquiring firms. *Journal of Financial Economics 49*(2), 223–253.

Renneboog, L. and C. Vansteenkiste (2019). Failure and success in mergers and acquisitions. *Journal of Corporate Finance 58*, 650–699.

Reyes, T. (2018). Limited attention and M&A announcements. *Journal of Empirical Finance 49*, 201–222.

Rodrigues, B. D. and M. J. Stevenson (2013). Takeover prediction using forecast combinations. *International Journal of Forecasting 29*(4), 628–641.

Roll, R. (1986). The hubris hypothesis of corporate takeovers. *Journal of Business 59*(2), 197–216.

Schijven, M. and M. A. Hitt (2012). The vicarious wisdom of crowds: Toward a behavioral perspective on investor reactions to acquisition announcements. *Strategic Management Journal 33*(11), 1247–1268.

Schneider, C. and O. Spalt (2021). Bidder and target size effects in M&A are not driven by overconfidence or agency problems. *Critical Finance Review, forthcoming*.

Schoenberg, R. (2006). Measuring the performance of corporate acquisitions: An empirical comparison of alternative metrics. *British Journal of Management 17*(4), 361–370.

Servaes, H. (1991). Tobin's Q and the gains from takeovers. *The Journal of Finance 46*(1), 409–419.

Shapley, L. (1997). *A value for n-person games. Contributions to the Theory of Games II (1953)*, pp. 69–79. Princeton: Princeton University Press.

Shi, H.-L., Y.-M. Wang, S.-Q. Chen, and Y.-X. Lan (2017). An approach to two-sided M&A fits based on a cross-efficiency evaluation with contrasting attitudes. *Journal of the Operational Research Society 68*(1), 41–52.

Shleifer, A. and R. W. Vishny (1989). Management entrenchment: The case of manager-specific investments. *Journal of Financial Economics 25*(1), 123–139.

Shleifer, A. and R. W. Vishny (2003). Stock market driven acquisitions. *Journal of Financial Economics 70*(3), 295–311.

Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. *The Journal of Business 74*(1), 101–124.

Smeekes, S. and E. Wijler (2018). Macroeconomic forecasting using penalized regression methods. *International Journal of Forecasting 34*(3), 408–430.

Song, M. H. and R. A. Walkling (2000). Abnormal returns to rivals of acquisition targets: A test of theacquisition probability hypothesis'. *Journal of Financial Economics 55*(2), 143–171.

Strong, N. (1992). Modelling abnormal returns: A review article. *Journal of Business Finance & Accounting 19*(4), 533–553.

Tashman, L. J. (2000). Out-of-sample tests of forecasting accuracy: An analysis and review. *International Journal of Forecasting 16*(4), 437–450.

Teplova, T., A. Kurkin, and V. Baklanova (2023). Investor sentiment and the NFT market: Prediction and interpretation of daily NFT sales volume. *Annals of Operations Research*, 1–25.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological) 58*(1), 267–288.

Timmermann, A. (2018). Forecasting methods in Finance. *Annual Review of Financial Economics 10*, 449–479.

Timmermann, A. and C. W. Granger (2004). Efficient market hypothesis and forecasting. *International Journal of Forecasting 20*(1), 15–27.

Travlos, N. G. (1987). Corporate takeover bids, methods of payment, and bidding firms' stock returns. *The Journal of Finance 42*(4), 943–963.

Tunyi, A. A. (2021). Revisiting acquirer returns: Evidence from unanticipated deals. *Journal of Corporate Finance 66*, 101789.

Tunyi, A. A. and M. Machokoto (2021). The impact of weather-induced moods on M&A performance. *Economics Letters 207*, 110011.

Valizade, D., F. Schulz, and C. Nicoara (2024). Towards a paradigm shift: how can machine learning extend the boundaries of quantitative management scholarship? *British Journal of Management 35*(1), 99–114.

van Binsbergen, J. H., X. Han, and A. Lopez-Lira (2020). Man vs. machine learning: The term structure of earnings expectations and conditional biases. Technical report, National Bureau of Economic Research.

Wainer, J. and G. Cawley (2021). Nested cross-validation when selecting classifiers is overzealous for most practical applications. *Expert Systems with Applications 182*, 115–222.

Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural Computation 8*(7), 1341–1390.

Yankov, D., D. DeCoste, and E. Keogh (2006). Ensembles of nearest neighbor forecasts. In *European Conference on Machine Learning*, pp. 545–556. Springer.

Zaheer, S., M. Schomaker, and M. Genc (2003). Identity versus culture in mergers of equals. *European Management Journal 21*(2), 185–191.

Zhou, X., H. Zhou, and H. Long (2023). Forecasting the equity premium: Do deep neural network models work? *Modern Finance 1*(1), 1–11.

Zhu, B., C. Wan, P. Wang, and J. Chevallier (2023). Forecasting carbon market volatility with big data. *Annals of Operations Research*, 1–27.

Supplemental Online Material

Table I: Out-of-sample R^2 for the different forecasting methods tested. Subsample with All Cash = 0

	R^2_{OS}
LGBM	2.39%
RF	1.67%
KNN	1.24%
Lasso	-0.19%
Ridge	-9.97%
OLS	-14.88%

Notes: Out-of-sample R^2, from larger to smaller, for a subset of deals at least partly paid with stocks. OLS is a linear regression using Ordinary Least Squares as estimator. Ridge and Lasso are regularized linear regressions using Ridge and Lasso approaches, respectively. RF is a Random Forest model. KNN is a K-Nearest Neighbor model, while LGBM refers to Light Gradient-Boosting Machine.

Table II: Out-of-sample R^2 for the different forecasting methods tested. Subsample with All Cash = 1

	R^2_{OS}
LGBM	1.22%
KNN	0.55%
Lasso	-0.46%
RF	-0.75%
Ridge	-55.27%
OLS	-99.18%

Notes: Out-of-sample R^2, from larger to smaller, for a subset of deals fully paid in cash. OLS is a linear regression using Ordinary Least Squares as estimator. Ridge and Lasso are regularized linear regressions using Ridge and Lasso approaches, respectively. RF is a Random Forest model. KNN is a K-Nearest Neighbor model, while LGBM refers to Light Gradient-Boosting Machine.

Table III: Out-of-sample R^2 for the different forecasting methods tested. Subsample with private target companies only

	R^2_{OS}
LGBM	3.06%
RF	1.80%
KNN	0.63%
Lasso	-0.15%
Ridge	-13.42%
OLS	-34.75%

Notes: Out-of-sample R^2, from larger to smaller, for a subset of deals involving private targets only. OLS is a linear regression using Ordinary Least Squares as estimator. Ridge and Lasso are regularized linear regressions using Ridge and Lasso approaches, respectively. RF is a Random Forest model. KNN is a K-Nearest Neighbor model, while LGBM refers to Light Gradient-Boosting Machine.

Table IV: Out-of-sample R^2 for the different forecasting methods tested, with completed dummy included

	R^2_{OS}
LGBM	2.37%
RF	2.00%
KNN	1.53%
Lasso	-0.23%
Ridge	-18.05%
OLS	-31.18%

Notes: Out-of-sample R^2 for the different forecasting approaches, ranked from larger to smaller. OLS is Ordinary Least Squares. Ridge and Lasso are regularized linear regressions according to ridge and lasso approaches,respectively. RF is a random forest regression model. KNN is a K-nearest neighbor regression model, while LGBM refers to Light Gradient Boosting Machine.

Table V: Out-of-sample R^2 for the different forecasting methods tested, for CAR[+2,+60]

	R^2_{OS}
LGBM	0.22%
KNN	0.10%
RF	-0.01%
Lasso	-0.39%
Ridge	-47.90%
OLS	-66.37%

Notes: Out-of-sample R^2, from larger to smaller, using abnormal stock returns over an extended window following the deal as the dependent variable. OLS is a linear regression using Ordinary Least Squares as estimator. Ridge and Lasso are regularized linear regressions using Ridge and Lasso approaches, respectively. RF is a Random Forest model. KNN is a K-Nearest Neighbor model, while LGBM refers to Light Gradient-Boosting Machine.