Truth and Bias in M&A Target Fairness Valuations: Appraising the Appraisals

Truth and Bias in M&A Target Fairness Valuations: Appraising the Appraisals

Matthew Shaffer
Harvard Business School[*]

May 2019

Abstract

In U.S. M&A transactions, target directors are effectively required to seek and consider a "fairness opinion" and supporting valuations before accepting a takeover offer. Despite their legal status, critics have argued that these valuations are biased, and that even unbiased valuations would not be useful to public companies, which can use their pre-deal stock price as an appraisal of their value. I develop and implement tests for the *usefulness* and the *bias* of fairness opinion valuations, and find tight evidence for both. They impound information about fundamental mispricing in targets, and prospective deal synergies, which could make them useful to directors in exercising their duties. They also exhibit predictable bias: providers toggle their discount-rate assumptions *ex post* to rationalize negotiated deal prices. However, this bias has been constrained by judicial scrutiny in recent years. These findings suggest that third-party appraisal could have a useful role in M&A governance, and suggest avenues for reform.

Keywords: Fairness opinions; Valuation; Corporate governance; M&A

JEL: G30, G34

[*]First draft: November 2017. The author (mashaffer@hbs.edu) is a doctoral candidate at Harvard Business School. I thank Suraj Srinivasan, John Coates, Charles C.Y. Wang, Paul Healy, Krishna Palepu, Steven Davidoff Solomon, Eugene Soltes, Christopher Janssen, Matthew Cain, Michael Wildish, Rick Lacher, Gil Matthews, Chris Mercer, Rand Curtiss, Gerardo Perez Cavazos, Ethan Rouen, Akash Chattopadhyay, Jee Eun Shin, Aaron Yoon, Gennaro Ameno, Mike Educate, Jody Grewal, Jihwon Park, Elia Ferracuti, Alice Y. Wang, Wei Cai, Alexa Scherf, Liran Eliner, Yonah Meiselman, Nate Schwalb, Wilbur Chen, and Lynnie Fein-Schaffer for helpful comments, criticisms, and suggestions. Comments are most welcome.

1 Introduction

1.1 Overview

I develop and implement tests for the usefulness and bias of the valuations supporting target fairness opinions in mergers and acquisitions (M&A). In the U.S., target directors are required to seek and consider a fairness opinion—a third-party attestation that the terms of the transaction are "fair" to target shareholders, with supporting valuations—before accepting a takeover offer.[1] Fairness valuations[2] have important legal status. They are used to establish that directors have satisfied their duties of care and loyalty. And they frequently feature in shareholder litigation after the deal, as evidence that plaintiff target shareholders did or did not receive fair value. But they are also frequently criticized. Critics allege that fairness opinion valuations are biased, subjective, and sometimes low-quality—and that even unbiased appraisals would not improve the governance of M&A. Proponents argue that they provide useful information to target directors and shareholders. I test these claims empirically, and provide evidence that is relevant to this debate.

Accepting a takeover offer may be the most important decision a company can make. Target shareholders usually gain hefty returns, but their ownership rights are extinguished or significantly diluted (in a stock deal). More broadly, the allocation of ownership and control rights is vital for efficient investment and production (Grossman and Hart, 1986). Takeovers are also rife with conflicts of interest. Several governance mechanisms have arisen to mitigate them. U.S. M&A law and practice include a surprisingly large role for *third-party appraisals*,[3] such as fairness opinions, as one such mechanism.

This role is controversial. The most practical criticisms (see Section 2.2) center on the bias and ambiguity in the process as it is currently practiced. But there is also debate on a more fundamental question, of whether *any* third-party appraisal process—even if reformed—could improve M&A governance. Consider the following arguments: Managers' and shareholders' interests can be aligned through compensation contracts. Even if they are not fully aligned, if the target sells itself in a competitive

[1] Strictly speaking, it is more accurate to say that target directors are *effectively* required to seek and consider fairness opinions. There is no statutory law enforcing it, but prevailing interpretation of case law has rendered them ubiquitous.

[2] It is common to distinguish between the "fairness opinion"—the letter testifying that the terms of the deal are fair—and the "board book" of valuations that support it. For the sake of parsimony, I refer to these as "fairness valuations."

[3] This includes the fairness opinion process, but also includes Delaware appraisal. In Delaware appraisal, shareholders of bought-out firms can claim appraisal rights, have the court determine the value of their shares, and receive that value in lieu of the negotiated deal price. This is a controversial and live area of U.S. corporate law, and many of the arguments in the debate about the value of fairness opinions are mirrored in this controversy. See, e.g., the commentary around Vice Chancellor Laster's decision in the Aruba Networks case.

auction, its shareholders will receive the highest value feasible. If the value of a company is what the market will pay, then this price is fair by identity. For a public target, the pre-announcement stock price measures its standalone value, so the transaction premium is a summary statistic for a fair, well-priced offer. What value can a subjective valuation add to these "market tests"?

The case for third-party appraisal depends on the claim that these arguments fail. It may be impossible to fully align managers' and shareholders' interests through incentives. Managers are often reluctant to accept attractive takeover offers because they enjoy private benefits of control. Boards can try to counter this bias through change-in-control benefits ("golden parachutes"), but these can introduce new agency problems, increasing managers' "slack" (Bebchuk, Cohen, and Wang, 2014). Neither auctions nor the "premium test" fully resolve the problem. Targets fear the proprietary costs of exposing themselves to diligence by competitors, and acquirers fear a winner's curse, so both are loath to participate in competitive auctions (Boone and Mulherin, 2007). Target managers can exercise discretion over which prospective buyers are included and affect their willingness to pay by signaling their willingness to cooperate. Finally, the premium alone cannot be a summary statistic for an attractive offer if: (i.) the target was fundamentally mispriced; (ii.) the target was rumored or known to be seeking a buyer, so its stock price incorporated some expectation of a takeover; or, (iii.) the target could bring idiosyncratic synergies to the acquirer, for which it could negotiate. (See Section 2.1.)

Given this, there may be room for third-party appraisal as a governance mechanism in takeovers. *Ideally*, unbiased, accurate valuations could alert directors if their stocks had been temporarily undervalued—whether due to market inefficiency, or managerial manipulation. Or, they could inform managers or directors about the value of the synergies they could bring to prospective acquirers, to establish their negotiating position even *in lieu* of a competitive auction. Both industry insiders and my own empirical results indicate that the real fairness opinion process falls short of the ideal. But studying it is important for two reasons: First, U.S. courts have given fairness opinions legal status in establishing that directors have exercised their duties. It is worth assessing whether fairness opinions warrant this legal status. Second, studying this setting can provide evidence on how the current process might be reformed, and what role, if any, such a reformed process should play in governing takeovers.

1.2 Empirics

I obtain access to a proprietary, comprehensive database of fairness opinions disclosed in U.S. public company transactions from January 2006 through 2016. This database includes the discounted-cash-flow (DCF) valuations supporting the fairness opinion, and I collect the discount-rate and perpetuity-growth-rate (or terminal-multiple) assumptions for each. These DCF valuations and components are the focus of my empirical analysis.

1.2.1 Validation tests

My first set of tests centers on *validating* target-sought fairness valuations–that is, testing whether they impound useful, valid information about the value of the target. This poses an immediate research challenge. If we wished to validate an equity analyst's recommendations, we could test for whether the stocks that she rates a "buy" outperform, and those she rates a "sell" underperform, under the assumption that prices converge toward fundamental value over time (Bradshaw, Brown, and Huang, 2013). But the targets in my sample are almost always acquired. So their stock prices converge to the negotiated deal price until the effective date of the transaction, regardless of whether their "true" value differs from the deal price (which is precisely what the fairness valuation is supposed to determine). I have two major ways of getting around this problem.

First, I identify a subsample of cases in which targets had their offers withdrawn—and thus, continued to trade as independent firms—after they had received (and disclosed) their fairness valuations. For this subsample, I can implement a test much like the classic equity analyst test above. I find that the valuations of withdrawn targets forecast their subsequent abnormal returns (Figure 1 and Table 5). Companies that receive higher fairness valuations, but then have their offers withdrawn and continue to exist as standalone companies, perform better than companies that receive lower valuations, over the medium- to long-term. This difference persists after conditioning on the (foregone) premium, which suggests that fairness valuations can detect fundamental mispricing. This finding is consistent with the evidence on mispricing (Shleifer and Summers, 1990), and the ability of even fairly simple fundamental valuation procedures to detect and exploit it (see, e.g., Frankel and Lee (1998), Richardson, Tuna, and Wysocki (2010)).

Second, I use a toy model of an acquisition to calculate what the impact of the transaction on the *acquirer* would be if the fairness valuation of the target were correct. I call this quantity the *FOPRa*

(for "Fairness Opinion Predicted Return for the Acquirer"). For example, suppose that the acquirer has a market cap of $100, and it buys a target for $10, but the average fairness valuation indicates that the target is worth $9 to the acquirer. Under the assumption that the fairness opinion valuation is correct, the NPV to the acquirer from the transaction would be −$1, yielding a *FOPRa* of −1%. This allows me to validate the fairness opinion valuation of the *target* using the *ex post* outcomes for the *acquirer*, and use the vast majority of my sample (>95%), in which the deal is actually completed.

I find that the *FOPRa* can forecast the announcement-date abnormal returns for the acquirer (Figure 5 and Table 6). The fairness opinion is not disclosed on the date of the acquisition announcement, so the valuation cannot be *driving* the market reaction. Rather, this shows that the market's appraisal of the deal corresponds to, and validates, the fairness valuation. This holds even after conditioning on the (appropriately scaled) market premium, indicating that the fairness valuation impounds information that is incremental to the valuation represented by its pre-deal stock price. This result is driven by DCF valuations that explicitly include *synergies*.[4] Notably, while the coefficient in these regressions is significantly greater than zero, it is also significantly less than one, which can be interpreted to mean that the fairness valuation estimates the value of the target to the acquirer with significant measurement error. This finding is also reflected in the acquirer's future accounting: Deals in which the *acquirer* overpays according to the fairness opinion (i.e., where the deal price is high relative to the fairness valuation) are more likely to generate an impairment in the future (Table 7).

1.2.2 Bias tests

My next tests center on detecting *bias* in fairness valuations. I first test for how the *discount-rate* assumptions in fairness valuations change in response to the deal premium.[5] Critics allege that fairness opinion providers rationalize negotiated deal prices *ex post* in order to satisfy client management, rather than providing a truly independent appraisal. One way in which providers could rationalize deal terms

[4] I use the term 'synergy' in the broad sense, to mean any thing that makes two companies worth more together than separately, rather than in the narrow sense of specific cost savings. I code a DCF as being inclusive of synergies if the model includes the effects of any policy changes or scenario that is conditional on the acquisition.

[5] A natural question is why I test for bias in the *components* of the DCF valuations rather than the level of the ultimate valuation. There are two main reasons: First, it is not credible to test for bias in the *level* of the valuation of the target, e.g., by measuring how this quantity differs from the target's pre-deal market price, because there could be fundamental mispricing or deal-level synergies. Stated differently, there is no benchmark for the "true" value of the target, relative to which bias can be credibly measured by a researcher. Second, biased fairness opinion providers would not have monotonic incentives to increase or decrease the valuation of the target. Holding the deal price fixed, increasing the fairness valuation will make the deal price look *less attractive* and possibly *less fair* by comparison, and thus undermine support for the deal. Valuing it too low could cause acquirers (or their stakeholders) to balk, or lack credibility.

would be to decrease their discount rate assumptions (and thereby increase their valuations) to match richly-priced deals—and, symmetrically, increase their discount-rate assumptions for poorly priced deals.

I find strong evidence of such bias (Table 8). Providers appear to toggle their discount-rate assumptions *ex post* to rationalize deal prices. Economic theory provides guidance on the determinants of discount rates, which allows me to estimate a credible normative benchmark and get tight inference in this test. The coefficient from my most robust specification implies that in response to a 10 percentage point increase in the negotiated deal premium, fairness opinion providers decrease their discount rate assumptions by .19 percentage points on average. This is sizable relative to the relevant benchmark: Using in-sample mean parameters, this would generate a 3% change in the valuation, in terms of the implied premium (i.e., about 30% of the total). But this bias has, in recent years, been constrained by legal scrutiny: I find that it declined significantly after the *In re: Southern Peru Corporation* decision, especially for providers that were most exposed to the implications of the case (Table 9).[6]

I also test for bias via the perpetuity growth rate assumption. Here, economic theory does not provide as strong guidance as it does on discount rates, but it does provide a fundamental constraint: A company's perpetuity growth rate cannot logically exceed the perpetuity growth rate of GDP. I do *not* find that providers violate this constraint to rationalize deal prices (Table 10). This suggests that the bias and discretion in fairness valuations is somewhat *constrained* by best practices.

Finally, I explore how fairness opinion providers' valuations behave in especially conflicted transactions, management buyouts (Table 11). Consistent with prior research, I find that target shareholders receive lower premia in these transactions. Consistent with the *validity* of fairness valuations, the *average* valuation "pushes back" against these lower prices (i.e., assigns the target a higher fair valuation relative to the deal price). However, consistent with their *bias*, the providers dramatically expand their valuation ranges, to declare these transactions "within the range of fairness."

1.2.3 Reconciling

How can it be that fairness valuations have validity, even beyond that of market prices, given the biased incentives of the providers and the subjectivity of valuation? A theme in my conversations with practitioners—one that is echoed in my perpetuity growth rate test—is that they are highly

[6] In that decision, Chancellor Strine specifically criticized an investment bank fairness opinion provider's *post hoc* valuations, noting that they changed their discount-rate assumptions across multiple drafts of the board book, as the terms of the deal were revised. Industry insiders saw this case as embarrassing for the provider and investment banks reformed the control, documentation, and standardization of their valuation procedures thereafter.

constrained in producing their valuations. Some providers claimed that they were committed, out of ethics and professional pride, to upholding valuation best practices and objectivity; others perceived the fairness opinion as a difficult but necessary task to close the deal. But both types emphasized that their valuations would be disclosed in SEC filings and subject to scrutiny by their peers and M&A professionals, and—in the probable event of litigation—lawyers, economic consultants, and judges. As a result, they have limited discretion to deviate from best practices.

They also have access to private information and professional M&A experience. These facts, combined, can explain my results. Fairness valuations can detect fundamental mispricing, because the providers take free-cash-flow forecasts that impound some private information, and put them through a constrained fundamental valuation procedure. Their valuations can impound the value of expected synergies from the deal, because they can use their M&A expertise, access to target management, private deal information, and, in some cases, information from advising on the deal at earlier stages.

How, then, can the validation results be reconciled with the bias results? There is some tension here—the more that providers simply *ex post* rationalize the deal prices, the less valid their valuations are—but no strict contradiction. In Section 5, I provide a toy example that illustrates this.

1.3 Contribution

My paper is, to my knowledge, the most comprehensive empirical analysis of target fairness valuations to date. It is the first to document that these valuations impound information that is *incremental* to the target's pre-announcement market price, which could, consequently, render them useful to directors of public targets. It is also the first to document their bias.

This has implications for practitioners, regulators, and the courts. It suggests that directors could benefit from consulting third-party fundamental valuations before accepting a takeover offer, in order to avoid selling out when they have been fundamentally undervalued, and to ensure that they have negotiated for an adequate share of the synergies from the transaction. But the bias in the process attenuates this benefit. Boards, regulators (whether the SEC or FINRA), or the courts could restructure the process to mitigate this bias. It appears that judicial scrutiny in particular can be effective in reducing the bias. My findings are also relevant to the live debate over Delaware appraisal rights, which also centers on the question of whether and when an abstract valuation can serve as a superior measure of value to a negotiated transaction price.

My paper also contributes to the academic literature. While there is limited empirical research on this setting (see Section 2.3), it parallels the literature on financial analysis, intermediaries, and analysts, and their role in the stewardship and valuation functions in capital markets (Beyer, Cohen, Lys, and Walther, 2010). This literature has shown that analyst forecasts and valuations are informative (Francis and Soffer, 1997; Brav and Lehavy, 2003; Oehmichen, 1979), though also biased by analysts' incentives to curry favor with clients (Lim, 2001), career concerns (Trueman, 1990), and their employers' underwriting and other economic relationships (Dugar and Nathan, 1995; Dechow, Hutton, and Sloan, 2000), which can undermine their credibility and usefulness (Fischer and Stocken, 2001). But this bias (and its impact on accuracy) can be more subtle than critics allege (Groysberg, Healy, Serafeim, and Shanthikumar, 2013), as it is mediated by reputational incentives (Ljungqvist, Marston, and Wilhelm, 2006), complicated repeated-game dynamics (Kasznik and Mcnichols, 2002; Ke and Yu, 2006), and shaped by the changing regulatory and legal environment (Francis, Nanda, and Wang, 2006). I show how these themes in the research on financial intermediaries—the concurrence of, and tension between, information and bias, and the reputational and regulatory constraints (or lack thereof) on bias—emerge in my understudied setting.

My research is situated within the rich literature on corporate governance mechanisms (see, e.g., Armstrong, Guay, and Weber (2010) and Shleifer and Summers (1990)) that have developed in response to the separation of ownership and control (Jensen and Meckling, 1976). This literature has focused on the use of formal contracts, incentives, and structure (such as board composition) in aligning managers' delegated investment choices with the interests of their stakeholders. Third-party appraisal has received much less attention, despite its role and status in corporate law and practice. My findings call for further research on this alternative governance mechanism.

The remainder of the paper proceeds as follows. Section 2 provides further background on fairness opinions and valuations, and a review of prior literature on the topic. Section 3 describes my data sources and sample construction. Section 4 explains, reports, and discusses my empirical tests. Section 5 discusses implications and concludes.

2 Background

2.1 Theory Background: Role of Third-Party Appraisal

A third-party appraisal process could only be useful in M&A governance to the extent that the other, primary, governance mechanisms fail. In the following subsections, I consider arguments for why these mechanisms—CEO incentives, go-shops, and a market premium test for public targets—may fail, and thus what the complementary role of third-party appraisal could be.

2.1.1 Problems with the main governance mechanisms: auctions and incentives

Competitive auctions have not fully resolved target agency problems in M&A. The sale process for target firms is rarely fully competitive (Boone and Mulherin, 2007). Targets fear the proprietary costs of exposing themselves to diligence by competitors when a takeover offer is still uncertain (Verrecchia, 2001). Buyers are also reluctant to participate in very wide auctions, due to their fear of the winner's curse. Thus, targets are typically shopped to a fairly small set of prospective buyers. And the target's managers can exercise discretion over which prospective buyers are included and what private information is revealed to them. They can also affect buyers' willingness to pay by signaling their attitude toward the buyer and willingness to cooperate during integration.[7] As such, a go-shop process does not guarantee that the target has obtained a value-maximizing offer.

Nor can incentive contracts fully resolve the agency problem for takeover targets. There are at least two irresolvable problems: First, target managers may enjoy private benefits of control, which will partially cleave their interests from those of shareholders, no matter how much equity ownership they have. One solution to this problem is to grant managers change-in-control benefits ("golden parachutes," Lambert and Larcker (1985)), which can counterweight the manager's private interests, and encourage him to accept attractive takeover offers. But they introduce a new problem: Now, the manager may reduce his efforts to maximize the value of the firm on a standalone basis, preferring instead to "flip" it.[8] It is impossible to perfectly calibrate the manager's incentive contract without knowing his private preferences for control and effort.[9] Second, the acquirer will often want to retain the target manager

[7] Consider, for example, the 2013 management buyout of Dell Inc. While there was a formal "go-shop" process in this deal, prospective buyers of the company knew that Michael Dell wanted to take the company private himself, and thus might be reluctant to serve as a value-maximizing manager working with alternative outside ownership (Parloff, 2016).

[8] For example, Hartzell, Ofek, and Yermack (2004) find that managers with larger golden parachutes accept lower premia in takeovers; and Bebchuk et al. (2014) find that they reduce shareholder value, possibly by increasing managerial "slack."

[9] As an example, consider again the controversial management buyout of Dell Inc. in 2013. At the time, some observers

post-transaction. This gives the target manager an interest in the acquiring firm—and in pleasing its management.[10] This once again cleaves her interests from those of the firm's shareholders no matter how strong her equity incentives.[11]

2.1.2 Public companies: the premium test

Public companies have access to an easily observable, efficient appraisal of their value: their stock price. When a public company receives a takeover offer, the deal consideration is usually referenced in terms of its premium over the pre-deal market price. So it is reasonable to argue that the premium is a summary statistic for the attractiveness of the offer.

But there are at least three problems with this argument. First, there is evidence of fundamental mispricing in the stock market (Frankel and Lee, 1998; Shleifer and Summers, 1990). If a stock has been mispriced, then the market premium does not measure the premium to its fundamental value. This problem is especially pronounced in conflicted transactions—there is evidence that managers manipulate discretionary accruals to suppress the reported profitability, and thus the stock price, of their firms prior to management buy-outs (Perry and Williams, 1994) and that they time these transactions for periods in which their industries are fundamentally undervalued (Harford, Stanfield, and Zhang, 2018).

Second, the price that a target can negotiate for should be a function of the synergies that will be realized from the transaction.[12] Prospective deal synergies are not reflected in the target's pre-period market price, and so directors cannot infer the maximum price they can negotiate for using only their company's price history. Put plainly, a 15% premium may be the best option for a target that will not bring large synergies to an acquirer, but a poor option for one that will.

Third, and finally, if the target has been the subject of M&A rumors or actively announced that it is considering "strategic alternatives,"[13] and the expectation of a takeover has been impounded in its

argued that Michael Dell was underpaying for the company, as he would only buy the company if he had private information that it was undervalued. Others argued that he was willing to radically overpay for the company, due to his private, non-pecuniary interests in controlling the company that literally bore his name.

[10] Chen, Thomas, and Zhang (2016) document that target managers will suppress reported earnings after the deal is announced, but before it is consummated, to help "spring-load" the acquirer's post-integration profitability—an apparent effort to benefit, and curry favor with, the acquiring firm's managers.

[11] Note that both of these two problems persist even when the manager has 100% ownership of the firm.

[12] I use the term 'synergy' in the broad sense, to mean any thing that makes two companies worth more together than separately, rather than in the narrow sense of specific cost savings. Under this definition, it follows arithmetically that for the acquirer to yield positive NPV on the transaction, the value of the synergies must exceed the premium paid. Under the assumption that acquirers will not extend offers that are NPV negative, the deal-level synergy defines the upper bound of the premium the target can negotiate for.

[13] A well-known and boilerplate euphemism for seeking a buyer (Zha Giedt, 2018).

stock price, the market premium will not measure the premium relative to the company's standalone value. Once a company is known to be seeking a buyer, it can be impossible to disentangle what portion of its stock price reflects its standalone value, and what portion reflects the expected takeover offer. In this scenario, the current stock price cannot be used as a benchmark to assess the attractiveness of incoming offers.[14]

2.1.3 Room for third-party appraisal?

Given these failures, there could be a theoretically-grounded role for third-party appraisal in the governance of M&A. *Ideally*, third-party appraisers could inform boards if their firms were being fundamentally mispriced (whether due to random noise, or intentional managerial manipulation) and, thus, the deal price was not as attractive as the market premium would imply. Or, they could inform boards of the synergies that their firms could bring to deals, to establish their negotiating position with respect to the acquirer. Or, they could inform boards of the value their firms could fetch in a competitive process, even in lieu of fully "opening the kimono" in an open auction. As I discuss below, the real-world fairness opinion process differs from this ideal—but it provides a benchmark for understanding the value that fairness opinion valuations could provide.

2.2 Background on fairness opinions

2.2.1 Law and process

Fairness opinions are, broadly defined, *products*, in which a third-party provider, typically a financial services or valuation advisory firm, evaluates some proposed transaction and vouches that the terms of the transaction are "fair" from a financial point of view. While fairness opinions can be sought and offered for transactions including divestitures, private placements, and liquidations, they are most prominent in M&A.

Fairness opinions only acquired legal status—and became ubiquitous in M&A—as a result of the Delaware Supreme Court's decision in *Smith vs. Van Gorkom* (488 A.2d 858 [Del. 1985]).[15] The

[14]Suppose a company announces that it is seeking "strategic alternatives" and its stock price rises 20% on the announcement. Now suppose that over the subsequent weeks, as the company seeks a buyer, the S&P500 increases by 10%, and the company's industry peers have returns that are volatile, but average 20%. If the company receives a takeover offer, how should it determine how attractive it is, relative to its standalone value? Should it assume that its value would have increased by 20%, in line with its industry peers, and demand a premium relative to that implied valuation?

[15]It is impossible to calculate empirically what portion of target companies seek and receive a fairness opinion, because no public database collects fully comprehensive records on fairness opinions in M&A. Thus, in asserting that fairness

court held the board of TransUnion grossly negligent in accepting a buyout offer without sufficient consideration or evidence that its terms were adequate. The court "leaned heavily" on the board's failure to seek and receive a fairness opinion on the transaction. As a result, "academics and practitioners inferred and concluded...that a fairness opinion was a virtually mandatory requirement for an acquiree board" (Davidoff, 2006). Acquirers may voluntarily seek fairness opinions, but there is no legal precedent compelling them to do so, or to disclose the opinion if they do.

Historically, the fairness opinion provider was typically an investment bank, which sometimes had prior and ongoing relationships with the company—including, even, involvement in the financing of the deal itself. This gave the opinion providers obvious incentives to declare the deal fair, and was a major complaint among critics. Recently, there has been a trend toward the employment of "independent" providers and valuation advisory firms, and/or providers who had no financing role in the transaction.[16] (See Appendix C for a list of providers by their frequency in my sample.) As of 2016, the largest opinion provider was an independent advisory firm. Providers are paid upon accepting the engagement, so their fee is not contingent upon providing a "fair" opinion—however, if the provider has other roles in financing or advising upon the transaction, the fees from those may still be contingent upon deal "success." To mitigate the perception or reality of bias, some financial services firms have adopted structures to give their fairness-opinion review committees autonomy and independence. There have been some cases in which an investment bank's fairness opinion committee has refused to give a "fair" opinion on a deal in which the investment bank stood to provide financing and other advisory services.[17]

Fairness-opinion projects are considered somewhat desirable among financial services and advisory firms, given the reasonably high fees relative to the negligible capital commitment. Fees for fairness opinion providers in large public-company transactions typically exceed $1 million. But providers perceive a high level of reputational risk, due to the frequent shareholder litigation around deals and the public disclosure of their fairness opinions and analyses.

The target company engages the fairness opinion provider during the late stage of negotiations, once a firm offer is on the table.[18] The provider performs multiple valuations of the company, each with an

opinions are "ubiquitous" for U.S. targets, I rely upon the claims of industry insiders and inductive inference from legal precedent and I assume going forward that there is no significant selection issue for targets. As one senior dealmaker put it, "The law doesn't explicitly say you need one—but there's no law firm in the U.S. that would let their client *not* get one."

[16]Independent advisors have risen from 49% to 62% of my sample over my sample period, from 2006 through 2016.

[17]Since I cannot systematically observe cases where a provider refuses to declare a deal fair, I cannot quantify the frequency with which this happens—but I have had one such case credibly described and validated.

[18]If the fairness opinion provider was already previously serving as an advisor to the target, it will often already have

associated range. As long as the deal consideration falls within the range of valuations, the opinion provider can declare it fair on that basis. If the fairness opinion provider cannot provide a fair opinion, then, in principal, the transaction could be terminated or renegotiated, or the target management could hire a different provider.[19] If the provider is willing to declare the deal fair, it signs a formal letter to that effect, and presents the supporting valuation analyses—often referred to as the "board book"—to the board, during the meeting at which it votes on the takeover offer. If the board votes in favor, the deal is "announced" and becomes public knowledge. But target shareholders must vote to approve the transaction before it can be effected.

In preparation for the vote, the target distributes a merger proxy (form DEFM14A) to its shareholders.[20] The fairness opinion is considered material information, and the letter is disclosed in its entirety, while the supporting valuations are usually summarized in a high level of detail.[21] In my sample, the fairness opinion letter occupies an average of 1.7 pages of text in the HTML version of the proxy statement, while the description of the supporting valuation analyses occupies an average of 7.3 pages. See Appendix D for excerpts from a randomly-selected transaction.

2.2.2 Valuation methods

Fairness opinions in M&A transactions typically include multiple valuation approaches: over 97% of my sample include more than one approach, with a mode of 3 and a mean of 3.5. These valuation approaches include: (1.) **DCF valuation**, usually with explicit discount-rate and perpetuity-growth-rate assumptions; (2.) **Comparable company and/or transaction multiples** (whether multiples

performed several valuations of the target. In these cases, the valuations could be adapted and refined for the fairness opinion. This does not introduce any bias or inferential problems to my tests—my validation tests only require that the deal price and fairness valuation be fixed before the deal is publicly announced. But it does offer an extra interpretation of my findings: In these cases, the fairness opinion valuations can be seen as a publicly-observable distillation of the otherwise-unobservable valuation advisory that is provided to targets during the pre-announcement period.

[19] When this happens, it is unobservable, but industry insiders attest that it is rare in public company transactions that a failure to get a "fair" opinion is a first-order determinant of deal completion.

[20] In a small number of cases—86 in my original sample—the fairness opinion provider revises its opinion and valuation analyses between the board meeting and the distribution of the merger proxy, in order (they argue) to provide an analysis that is more up-to-date and informative to shareholders. In these cases, the valuations that are disclosed are not the same valuations considered by the board, and the providers could potentially use the announcement returns to update their valuations of the target, so I drop these from my analysis.

[21] This is ultimately how the fairness opinions and their supporting valuations become public and are selected into my sample. This means that although all fairness opinions in my sample declare the deal "fair," there remains substantial variation in the valuations, which I can exploit in my research design. Put differently, the fairness opinion letters and valuations in my sample imply that, cross-sectionally, *all deals are within the range of fairness, but some deals are more fair than others.*

of earnings, sales, EBITDA, or EBIT);[22] (3.) **Trading price-range analysis**, in which the provider compares the deal price to the target's price history, usually over the past year; (4.) **Liquidation analysis**; and miscellaneous other approaches.

All of the fairness opinions in my sample include some kind of comparables analysis, though they vary in which multiples are employed. 94% of the target-sought fairness opinions in my sample include a DCF valuation. I focus on these DCF valuations for two main reasons: First, DCF analyses are consistent and nearly ubiquitous across fairness valuation practice. In the 6% of cases in which a DCF analysis is *not* included, the fairness-opinion provider typically explicitly justifies its exclusion (often because the company does not have positive cash flow, and DCF valuation methodologies would be unusually subjective and unreliable—e.g., small pharma companies with a promising drug, but no cash flow history). Second, economic and financial theory (Modigliani and Miller, 1958) justify the use of DCF to estimate a company's value, and this theory provides constraints and normative guidance on how these analyses *should* be implemented.[23] This will be vital for my tests and inferences.

2.2.3 Timeline

From the perspective of a target firm and its fairness-opinion provider, a typical M&A deal would proceed as follows: (1.) Company management discusses and negotiates general terms of a possible transaction, and the acquirer makes an offer. (2.) Target company management engages a financial advisor to perform a fairness opinion valuation of the company. (3.) The board of directors meets to consider and vote on the acquisition offer. At this meeting, the board is presented with the "board book" of valuation analyses produced by the fairness-opinion provider. The board of directors votes to approve the deal and recommends that shareholders vote in its favor and the deal is "announced," pending approval. (4.) The target distributes the merger proxy in preparation for the shareholder vote.

[22] Multiples are ubiquitous in practice, and prior research (Liu, Nissim, and Thomas, 2002) finds that even simple multiples are surprisingly powerful in explaining firm value.

[23] DCF valuation is a technique in which a company is valued by forecasting and discounting its future free cash flows—that is, cash flows available to the financial claimants of the company. It is justified by the Modigliani-Miller dividend irrelevance theorem: Under the equilibrium assumption that the company's internal rate of return equals its cost of external capital, it is "irrelevant" for firm value whether a company distributes or reinvests available cash flows, and the company can be valued by discounting free cash flows rather than cash distributions (Modigliani and Miller, 1958). A DCF valuation essentially consists of three parts: (1.) Near- to medium-term forecasts for free-cash-flow realizations; (2.) A perpetuity growth rate (i.e., the rate of growth for all periods beyond those being explicitly forecast); and (3.) A discount rate. A DCF can be implemented as a *levered* DCF analysis, in which the free cash flows available to *equity owners* are forecast and discounted at the company's required rate of return on its equity; or an *unlevered* DCF analysis, in which the total free cash flows available to financial claimants are forecast and discounted at the company's weighted average cost of capital (WACC), to calculate the total financial value of the company, and the value of equity is inferred by backing out the value of its debt and other senior claims. In fairness opinion valuations, unlevered DCF analyses are the norm.

On this date, the fairness opinion becomes public for the first time. (5.) Target shareholders vote on the deal. (6.) The transaction is formally effected (if it is not withdrawn). See Appendix B for a visual representation. Note that in order for a fairness opinion to enter my sample, the proposed transaction must proceed at least as far as step (4.) above.

2.3 Prior literature and perspectives on fairness opinions

Empirical research on fairness opinions in M&A has been somewhat limited, likely due to lack of data availability. Kisgen, Qian, and Song (2009) examine 1,500 fairness opinions over the 1994-2003 period and their relation to takeover premia paid during the transaction. They infer that target-sought fairness valuations do not influence deal premia. Bowers and Latham (2006) examine drivers of the demand for fairness opinions among *acquirers*, including whether the deal is hostile, the consideration offered, litigation risk, information asymmetry, and uncertainty. They conclude that acquirers are more likely to voluntarily seek fairness opinions when there is significant information asymmetry.

Cain and Denis (2013) examine the valuations of targets in both target-sought and acquirer-sought fairness opinions over the 1998-2005 period. They find no evidence that potential sources of bias such as company-investment bank relationships drive errors in the fairness valuation, with errors defined as differences between fairness valuations and market prices. They find that differences between target fairness valuations and the offered takeover price are positively associated with acquirer returns around the announcement date, and also around the target's proxy disclosure date, which suggests that fairness opinions bring new information to the market.[24]

There is a more extensive legal literature on fairness opinions, focusing on interpretations of case law and normative analysis of policy. On the whole, the legal literature is fairly critical. For example, Bebchuk and Kahan (1989) identified three major categories of problems with fairness opinions in M&A: (1.) The courts have not provided an unambiguous definition of the "fair price" that valuators were asked to attest to (the company's value as a standalone entity, its value to the acquirer, or its potential sale price in a competitive auction?). (2.) Investment bankers had direct economic incentives (fees for their other roles in the transaction that were contingent upon deal completion) and other incentives

[24]My baseline result in Table 6 is similar to this result for their target-sought fairness opinion subset. My subsequent specifications build upon their work in key ways: First, I add a control for the *Scaled premium*, which allows me to test for the validity of the fairness valuation that is *incremental* to the information that is impounded and publicly available in the target's pre-announcement market price. Second, I use cuts that allow me to identify that it is specifically the ability to forecast *synergies* that drives this result. Additionally, I use a different approach to test for bias, and find evidence thereof.

(long-term relationships and social affinities) to complete transactions desired by management. (3.) Valuation is by its nature subjective and uncertain (which valuation should we use, and how should we forecast risky future cash flows?). The authors characterize these major problems as the problem of "definition," the problem of "conflicts of interest," and the problem of "measurement," respectively. They argue that, together, these problems mean that providers have nearly unlimited discretion to come up with any valuation, and pure incentives to use that discretion solely in the interests of completing the transaction.

The legal literature since that time has largely accepted Bebchuk and Kahan's description of these three major problems, and critical perspective. Fiflis (1992) argues that to make fairness valuations useful, the investment bankers should be responsible and liable to the transacting companies' shareholders, with the same duties of "care, candor, and loyalty that directors have, as well as a duty of skill." Davidoff (2006) argues that there is limited use for fairness opinions in public company transactions in particular (in which the target has an easily observable standalone valuation) and recommends that the process be "reformed through a quasi-public standard-setting body."

In the public sphere, criticisms of, and debates about, the quality of fairness opinions and their role in M&A often flare up around controversial transactions.[25] For example, commenting on the 2016 Tesla-Solar City deal, Aswath Damodaran wrote that "the fairness opinion, as practiced now, is more travesty than protection and an expensive one at that," adding that, "the first option is to remove it from the acquisition valuation process."[26] Matt Levine, a former Goldman Sachs lawyer, offered a tempered perspective in this debate, arguing that fairness opinions provided a "sanity check," but could not be expected to do much more. He wrote:

> "A fairness opinion is not a real valuation, not a pure effort.... A fairness opinion says: Given certain explicit assumptions that we've agreed on with management, and running those numbers through certain standard valuation techniques, you can get a pretty wide range of values, and the price of this deal is within that range... It is a sanity check on the deal price... That is not nothing! But it's not much more than that."[27]

[25] Skeptics also claim that boards do not take the fairness opinion seriously, and see it only as a source of legal protection. I cannot test this claim empirically, as I cannot observe transactions that are voted down by boards, or cases where a provider declines to provide a fair opinion. But this claim, whether or not it is true, has no bearing on my validation tests, which are about the properties and quality of the valuation, not about directors' response to them.

[26] A. Damodaran, "Fairness Opinions: Fix Them or Flush Them!" Wall Street Oasis, 9/25/2016

[27] M. Levine, "Merchant Banks and Bond Lies," Bloomberg, 9/9/2016

3 Data and Sample Construction

3.1 Data

Fairness opinions are not comprehensively collected by any public data source. I obtain access to a proprietary, hand-collected database constructed by a major fairness opinion provider ("PROVIDER"), spanning 4,011 unique fairness opinions from U.S. public-company transactions from 2006 through 2016. I first validate that this dataset is comprehensive and unbiased by querying Thomson One for all significant M&A transactions over the sample period which, based on legal precedent, should generate a fairness-opinion disclosure. (Specifically, cash deals in which the target is a U.S. public company, more than 50% of shares turn over, and the total transaction value is greater than $100 million.) I am able to locate 88% of these deals in this dataset. The PROVIDER originally collected the dataset for its own research on the valuation methodologies being used in the industry, and so it is implausible that their data collection was biased with respect to my research questions or design. As such, I assume that my dataset is unbiased and acceptably comprehensive.[28]

3.2 Sample construction

I begin with the PROVIDER dataset which codes the price-per-share DCF valuations for each fairness opinion and includes the merger proxy in which it is contained. There are several additional variables that are vital to my research design. I search each proxy for each DCF valuation and code: (1.) The discount rate (WACC) assumption (typically a range); (2.) The perpetuity growth rate or multiple assumption (also typically a range); (3.) An indicator for whether or not the DCF valuation was inclusive of synergies;[29] (4.) An indicator for whether the fairness opinion valuation was a valuation of the target firm as a whole and sought by the target. I only hand-collect data for transactions with a value of more than $25 million.

[28] Thomson One's SDC database is the only public database that to my knowledge collects information on fairness opinions. However, its coverage is neither comprehensive nor consistent. And the database is considered insufficiently reliable for transaction-specific data that are hand-collected (e.g., specific merger contract terms) (Coates, Palia, and Wu, 2017; Becher and Cai, 2010). It is impossible to cleanly execute some of my key tests using the Thomson One sample, without supplementing with additional data from the fairness opinion section of the merger proxy. For example, the Thomson One fairness opinion variables do not reliably distinguish whether the target or a component of the target (as in a "sum of parts" valuation) is being valued, which is critical for my research design; nor does it distinguish between valuations that include synergies vs. those that do not.

[29] Since I use 'synergy' in the broad/algebraic sense in this paper, I code a DCF as inclusive of synergies if the model includes the effects of any policy changes or scenario that is conditional on the acquisition.

I match this sample to Thomson One using the target CIK and filing date from the SEC filing. All of my tests require at least one data point or identifier from Thomson One, so no unmatched observations are used in any of my tests. My initial sample then consists of 2,820 fairness opinions matched to 1,822 unique Thomson One deals. I then match the target and acquirer to CRSP, for market data, and to Compustat, for fundamentals. I collect target shareholder voting data from ISS Voting Analytics, via WRDS. I drop 86 transactions in which the fairness opinion and supporting valuations are revised by the provider after the deal's announcement date. Because of the importance of the withdrawn-transaction test, I append this subsample with an additional 65 transactions from Thomson One. I identified these transactions by searching for cases in which the deal was withdrawn after a merger proxy was distributed, and confirmed the DCF valuations in these transactions by hand. The sample size for each of my tests depends on data availability, which can vary, given the discretionary nature of the disclosure of valuation inputs like discount rates and perpetuity growth rates.

4 Empirical Design and Analysis

4.1 Key Constructs

Before proceeding to my empirical results, I describe the construction of, and economic logic behind, variables that are essential to my tests.

4.1.1 The fairness-opinion predicted return for withdrawn targets (*FOPRtw*)

In a small subsample of my dataset, the acquisition offer is terminated after a fairness opinion has already been produced and published. For this subsample, the target continues to exist and trade as an independent entity, but no deal-level synergies are realized. This makes it possible to test for whether fairness opinion valuations detect fundamental mispricing. To implement this test, I regress the abnormal returns to the withdrawn target on the return that would be predicted under the assumption that the fairness valuation were true. The fairness-opinion predicted return for the target, upon withdrawal ($FOPRtw$) is intuitive and easy to calculate. Formally:

$$FOPRtw = \frac{F.O.\ DCF_{tw}}{Deal\ Price_i} - 1, \tag{1}$$

where *tw* indicates targets whose offers are withdrawn, and *F.O. DCF* indicates the mean of the fairness-opinion DCF valuation.

4.1.2 The "baseline model" and the importance of the *Scaled premium*

For the vast majority of my sample, the transaction is completed, and I cannot validate the target fairness valuation via the target's performance, and instead must validate it via the acquirer's. The first construct I use to implement these tests is what I will call the "baseline model" of an acquisition. In the baseline model, there is no mispricing in the target and no synergy, so an acquisition simply represents a cash transfer from the acquiring firm to target stockholders, and the acquirer's return is a simple function of the premium.

To understand this, consider the following toy example: Suppose that prior to an announced deal, the acquiring firm, A Corp, had a market capitalization of $100, and the target firm, T Corp, had a market capitalization of $10. Suppose that A Corp offers T Corp a premium of 20%—$12 in total—and the deal is completed. If there are no synergies between A Corp and T Corp, and the market is perfectly efficient, then A Corp should now be worth $98, since it has disbursed $12 in cash and received an asset worth $10 in return, generating returns to acquiring shareholders of -2%. The implication from this example is general: In a world without synergies or mispricing of target firms, the return to acquirers would be completely determined by the premium, scaled by the relative market capitalizations of the target and acquirer. This quantity, the *Scaled premium*, is important for my validation tests, because it provides a benchmark against which to compare the performance of the fairness opinion valuation, and will serve as a control variable that will determine whether the fairness valuation provides anything *incremental* to the valuation impounded in the target's pre-deal market price. Formally, the *Scaled premium* is calculated as:

$$Scaled\ premium = Premium\ (\frac{MCAP_t}{MCAP_a}), \tag{2}$$

where a indicates acquirer variables, and t indicates target variables. One way to understand the *Scaled premium* is that it is the negative of what the predicted return for the acquirer would be under the baseline model. As such, this variable could have alternatively been named the "Baseline Model Predicted Return for the Acquirer" (*BMPRa*). Henceforth, I will preferentially use *Scaled Premium* as the *variable name*, but make reference to the baseline model.

4.1.3 The fairness-opinion predicted return for the acquirer (*FOPRa*)

The second construct I use is the fairness-opinion predicted return for the acquirer (*FOPRa*). This construct measures what the return for the acquirer would be if the fairness valuation of the target were true. To understand it, revisit the model above. As before, A Corp pays $12 in cash, for T Corp, a company whose standalone market price had been $10. But now, abandon the assumption that there is no mispricing or synergy. Suppose that a fairness valuation of T Corp implies that it is worth $11 to the acquirer.[30] If the fairness opinion valuation of T Corp is correct, then A Corp would be worth $99 after the transaction—since it has disbursed $12 in cash and received an asset that is worth $11 *to it* in return—yielding a return of −1%. Generalizing this, the predicted return for the acquirer implied by the fairness valuation can be calculated as:

$$FOPRa = \frac{(F.O.\ DCF_t - Deal\ Price_i)}{MCAP_a}, \tag{3}$$

where i indicates transaction-level variables, a indicates acquirer variables, and t indicates target variables. The difference between the fairness opinion valuation of the target (*F.O. DCF_t*) and the deal price is the total NPV to the acquirer implied by the fairness valuation; scaling this quantity by the acquirer's market cap generates its predicted return. Put simply, the *FOPRa* measures the difference between the fairness opinion valuation and the deal price, scaled by size.

4.2 Validation results

The regressions in this section test whether fairness valuations impound information that could be useful for public targets: specifically, about fundamental mispricing and deal synergies. This naturally raises the question of whether the null hypothesis is plausible: Why would this practice exist if it did not provide useful information? I argue that this question is *ex ante* unclear and non-trivial, for three main reasons. First, fairness opinions became ubiquitous in M&A not as a result of evolutionary, economic process, but as a result of judicial fiat, issued by non-economist judges—as such, the practice could simply reflect historical and institutional contingency, rather than economic rationality. Second,

[30] The difference between this value and its standalone market value of $10 could be explained by any of three factors: (1.) acquisition synergies that make T Corp worth more to the acquirer than it is on a standalone basis; (2.) pre-announcement mispricing of T Corp; or (3.) fairness opinion error.

since my sample consists of public companies, "useful" information is inherently non-trivial. If market efficiency is the null hypothesis, the ability to detect fundamental mispricing is the alternative by definition. Synergies from transactions are intrinsically unobservable and difficult to forecast. Third, it is not clear that the fairness opinion providers have any economic incentives to provide valuations that pass these non-trivial tests.[31]

4.2.1 Validation test #1: Withdrawn targets

Figure 1 graphs cumulative abnormal returns for targets in the withdrawn-transaction subsample, around the date of the withdrawal. The portfolios are sorted by the fairness opinion-implied predicted return for the target, upon withdrawal (*FOPRtw*), defined above, and abnormal returns are calculated using the Fama-French 3-factor model. Withdrawn deals are a useful subset, because we can observe *ex post* outcomes for the target as a standalone company, and the synergies from the proposed transaction are lost, so the ability of the *FOPRtw* to predict target outcomes reflects an ability to detect fundamental mispricing. Notably, over the medium to long-term, the above-median-*FOPRtw* portfolio substantially outperforms the below-median portfolio. The long-term drift suggests that this is due to fundamental mispricing in the target being corrected over time.

Table 5 displays the results of regressions of the target's abnormal returns arond the withdrawal date on the *FOPRtw*, over the short-run (Panel A) and long-run (Panel B). In Columns 1 and 2 of Panel A, the horizon is *t-3* days to *t+3* days around the withdrawal, while in the next pairs of columns, the horizon extends from *t-3* to *t+10*, all the way to *t-3* to *t+180*. In each pair of columns, I use both, (1.) a binary above-vs.-below-median *FOPRtw* split, which offers an intuitive interpretation of the result in portfolio terms, and is more similar to conventional market mispricing tests, and (2.) a linear specification, because the *FOPRtw* constitutes a continuous linear prediction of the target's fundamental value. I start the event study at *t-3* days because I find evidence of significant price movement in the target in the two days before the announcement of the withdrawal. Throughout these tests, I control for the market premium that had been offered to the target, because the target loses the offered premium around the withdrawal and this should be a major determinant of its price performance.

[31] To my knowledge, there has only been one case, *In re Rural/Metro*, where a fairness opinion provider suffered significant financial liability. Since this case occurs at the very end of my sample period, I do not have the statistical power to perform credible tests around it.

As is evident, the *FOPRtw*—both in the linear specification and the binary sort—does not predict target performance from *t-3* days to *t+3* or the other short-term horizons, but it does predict target performance over all of the medium- to long-term horizons. This is consistent with the interpretation that the *FOPRtw* captures fundamental mispricing in the target, which may take time to resolve (Frankel and Lee, 1998).

4.2.2 Validation Test #2: Acquiror Announcement Returns

Figure 2 graphically displays announcement-date abnormal returns for acquirers, sorted by the fairness-opinion predicted return for the acquirer, as defined in Equation 3. The top line represents the abnormal returns for acquirers in the top 50% of the *FOPRa* distribution, while the bottom line represents the bottom 50%. Note that while the fairness valuation of the target has been fixed and presented to the board of directors, it is not public on this date. As such, it *cannot* be the case that investors are reacting to the content of the fairness valuation. Rather, on this date, the market appraises the deal based on public information, and re-prices the acquirer based on the expected NPV of the transaction (i.e., negative for deals in which the acquirer overpays, and positive for deals in which the target is purchased for less than it is worth to the acquirer). As such, this graph indicates that the fairness opinion providers, in valuing the target company pre-announcement, and the market, in valuing the deal upon announcement, are sorting in concert. If markets are efficient in quickly impounding the expected impact of the acquisition, this validates the fairness valuation.

Table 6 displays regressions of the acquirer's announcement returns on the *FOPRa*. Column 1 reports the results of a simple univariate regression of the *3-day CAR* on the *FOPRa*. As is evident, the coefficient on the *FOPRa* is positive, precisely estimated, and highly significant. It is also significantly less than 1, which could be interpreted to mean that it forecasts the market reaction with measurement error. The estimated intercept, which can be interpreted as a measure of the systematic forecasting bias,[32] is not distinguishable from zero.

My goal in these tests is to assess the validity of the fairness valuations via their ability to forecast the acquirer's announcement date returns. As such, it would be contradictory to my purpose to control for firm-level characteristics. Consider a simple example: Suppose that a target will bring large tax

[32]Note that this *forecasting bias* is distinct from the main notion of bias—economic bias—that I examine next.

benefits to its acquirer, and that the provider impounds these tax benefits into the valuation of the target. Suppose that the market, in re-pricing the acquirer in reaction to the deal, impounds these tax benefits into its valuation of the deal. This would generate an association between the *FOPRa* and the market reaction, but this would not be a *problem* for my inference—it would be the basis of that inference. (See Angrist and Pischke (2009) on the "bad control" problem.) If the fairness opinion provider's appraisal of the tax benefits is reflected in the market reaction to the deal, then that validates the fairness valuation, which is what I am testing.

There is one major exception to this, one variable that it is important to control for, in order to properly interpret these results: the *Scaled premium*. Because the targets in my sample are public, the directors of these firms (and the providers) have access to an easily available appraisal of the value of the target company: its stock price. This raises the possibility that there could be a trivial explanation for the result in Column 1: Suppose that the fairness opinion provider simply replicated the company's pre-deal market price, and thus provided no *incrementally* useful appraisal to the board. Suppose also that, as in the "baseline model" described in Section 4.1.2, the premium paid by the acquirer equaled overpayment. This would induce an association between the *FOPRa* and the acquirer's announcement return, because the fairness valuation, by remaining tethered to the stock price, would indicate the overpayment. In this case, the fairness valuation would be "validated" by the market reaction, as per the specification in Column 1, but it would provide no value to the board of directors. Generalizing this, if the fairness valuation cannot forecast acquirer returns *incrementally* to the *Scaled premium*, then it does not provide a signal about the target's value to its acquirer that is useful incremental to the target's pre-deal stock price.

As such, I first compare the relative performance of the *FOPRa* and the *Scaled Premium* in forecasting acquirer returns, then add the *Scaled Premium* as a control to my main specification. In Column 2, I report the results of a univariate regression of the *3-day CAR* on the *Scaled Premium*. Three things are noteworthy about this result: First, the coefficient on the *Scaled Premium* is negative and significant, as predicted. This indicates that the variation in deal premia partially capture acquirer overpayment as perceived by the market. Second, the intercept term in this regression is positive, highly significant, and reasonably large. In other words, acquirers do better on average than would be predicted by the

Scaled premium.[33] This indicates that the market perceives deal synergies to be real and significant. This is an empirical falsification of the assumptions of the "baseline model" in Section 4.1.2. Third, the coefficient on the *Scaled premium* in Column 2 is significantly smaller in absolute size than the coefficient on *FOPRa* in Column 1; and the R^2 in the *Scaled premium* regression in Column 2 is significantly lower than the R^2 in the univariate *FOPRa* specification in Column 1.

The fact that the R^2 and coefficient is smaller in Column 2 means that the *Scaled premium* is less reliable in forecasting acquirer market returns than the *FOPRa*. This means that the fairness valuation of the target provides a better indication of the target's value to the acquirer than does the target's own pre-deal market price. If directors had to choose between only one metric—their market value, or their fairness valuation—to assess the value of their company to the acquirer and thus appraise whether they had negotiated for the maximum NPV-positive offer they could fetch from the acquirer, they would do better to rely on the fairness valuation. Since the R^2 in a regression is insensitive to linear re-scalings of the underlying variables, this means that the directors would also fare better in using the fairness valuation than any fixed benchmark market premium. Importantly, the fairness opinion valuation has an advantage in that it can incorporate deal-level synergies which are not captured in the target's pre-deal market price. So this result *should not* be interpreted to imply that fairness valuations are better than market valuations in general. Rather, what is interesting, and surprising, about this result is that the fairness valuations' ability to forecast synergies more than compensates for their biases—and thereby renders them an overall superior measure of the target's value to the acquirer.

In reality, directors do not need to choose—they have access to both their standalone market price, and the fairness valuation. In Column 3, I repeat the regression from Column 1, but include the *Scaled premium* as a control variable. The coefficient on *FOPRa* remains positive and significant, indicating that the fairness valuation forecasts acquisition performance incremental to the *Scaled premium*. In Column 4, I replace the continuous *FOPRa* measure with a simple binary indicator for above-median *FOPRa*, to ease interpretation of my result in portfolio terms. In this binary specification, the intercept indicates the average *CAR* for acquirers in the below-median *FOPRa* portfolio (-.414%), while the statistically significant coefficient on *Above-med FOPRa* (1.044%) indicates the outperformance of the

[33] More precisely, this interpretation also relies on the fact that the domain of the *Scaled premium* is, by construction, positive, and the slope is greater than -1.

acquirers in the above-median *FOPRa* portfolio.[34]

In Columns 5 and 6, I separately estimate the regression from Column 1 for DCF valuations that are explicitly inclusive of synergies vs. those that are not. Together, they show that the association documented in Columns 1 and 3 is concentrated among valuations that explicitly include synergies, consistent with the discussion of the baseline model, described in Section 4.1.2. For ease of interpretation, the regression coefficients are reported in separate columns, but the F-test statistics for differences in the coefficient on *FOPRa* and the intercept are reported in Column 6.

In Columns 7 and 8, I compare the DCF valuations produced by independent vs. non-independent fairness opinion providers. I implement these tests in order to cast light on the policy debate and ongoing reform in fairness opinion practice, discussed in Section 2.2.1. For this, I subset down to DCFs that are inclusive of synergies, in order to draw an apples-to-apples comparison. Once again, for ease of comparison, the regressions are reported in separate columns, while F-test statistics are reported in Column 8. Two things are noteworthy about the results: First, for both the independent and non-independent providers, the coefficient on *FOPRa* is positive and statistically significant. While the point estimates differ, this difference is not statistically measurable, as indicated by the insignificant F-stat. Both independent and non-independent providers produce valuations that are validated by market outcomes. Second, the intercepts differ sharply and significantly. In Column 7 (the regression for the independent subset) the intercept is negative and significant. This indicates that they on average *overvalue* the target in terms of its value to the acquirer (because the acquirer does *less well* on average than would be predicted from their valuations). In Column 8 (the regression for the non-independent subset), in contrast, the intercept is starkly positive and significant. This indicates, symmetrically, that the non-independent providers on average *undervalue* the target in terms of its value to the acquirer.[35]

One explanation for this result is that non-independent providers are in the business of financing

[34] A natural question is whether the censoring of observations when the transaction does not occur undermines this inference. This is not an inferential problem for the hypothesis being tested here, because the coefficient of interest is a slope coefficient (i.e., the ability of the FOPRa to forecast announcement returns) in a forecasting regression. Suppose that some part of the range of FOPRa were fully censored–under the null hypothesis of zero informativeness, the slope coefficient would still be zero over the observable range. As an analogue, consider the classic sell-side equity analyst test: Sell-side analysts may choose not to cover stocks that they see as likely to be bankrupt or delisted, but this does not cause inferential problems for testing whether the stocks they rate a buy outperform those they rate a sell.

[35] Technically, these stylized interpretations also depend on the domain of the *FOPRa* and the slope—but I verify that these differences hold by directly calculating the announcement-return-implied valuation of the deal and comparing it to the synergy-inclusive fairness valuations. Note that this *forecasting bias* is distinct from the main notion of bias—economic bias—that I examine in my next tests.

and making deals and may downwardly bias their appraisals of the target firm in order to make the deal price appear more attractive to its board and shareholders by comparison, to help push the deal through. The finding that independent providers overvalue targets mirrors the analyst literature: Jacob, Rock, and Weber (2008) find that independent (i.e., non-investment bank) analysts' forecasts are more optimistic.

4.2.3 Validation Test #3: Goodwill forecasts

Table 7 reports regressions of forward goodwill writedowns on the *FOPRa*, as an accounting-based complement to the market-based validation tests. I implement these tests because the returns-based tests are susceptible to the criticism that the market may not be able to impound the value of future synergies from the deal on the announcement date, and, as a result, the announcement-date returns to the acquirer may not provide a dispositive measure of the value of the target to the acquirer.[36] A tradeoff is that because goodwill impairments are relatively rare, only measure the downside of deal performance, and may be suppressed by managerial discretion (Ramanna and Watts, 2012), proxying for deal performance with this measure sacrifices statistical power.

For these specifications, the dependent variable, *Impairment*, is an indicator that evaluates to 1 if the acquirer recognizes a goodwill impairment at any point up to 5 years after the transaction. I use fairly parsimonious controls for these tests: I wish to test for whether fairness opinion valuations are validated by their ability to forecast goodwill impairments, an *ex post* measure of deal performance. This test would be biased if I included variables that are causally downstream from the acquisition, such as post-transaction stock returns or profitability (Pearl, 2009). As such, I omit these measures and other variables that could be driven by the performance of the acquisition.

I do, however, control for major pre-deal background differences in acquiring firms that could predict future goodwill impairments: The most obvious and significant among these is lagged *Goodwill*. For the same reasons as those elaborated in Section 4.2.2, I also control for the *Scaled premium*. I also include time fixed effects to control for macroeconomic shocks and to account for the partial censoring of the

[36] Goodwill impairments are widely seen as providing a signal of acquisition overpayment and underperformance (see, e.g., Ghosh, Hovakimian, and Hu (2018), Gu and Lev (2011) and cites therein), though accounting standards this area are also an area of controversy and debate in accounting. To the extent that managers use their discretion in goodwill impairments opportunistically (Ramanna and Watts, 2012), and this reduces the level of goodwill impairments, this would bias against my result.

data for transactions that occur in later years in my sample. Finally, I control for generic acquirer background characteristics: *Sales*, *Assets*, and *ROE*. There is no obvious reason why these variables would, if omitted, cause a *bias* that would undermine the inferences from my main specification. But I add them to achieve some basic linear matching and adjustment for local differences across my sample.

Columns 1 through 3 estimate a linear probability model with controls sequentially added. Columns 4, 5, and 6 estimate logit models with the same pattern of controls. As is evident, the coefficient on *FOPRa* remains negative and statistically significant throughout all specifications and models. Also noteworthy is the stability of the coefficient on *FOPRa* across specifications within each model. This indicates that it less likely that some correlated omitted variable is driving the association (Oster, 2017). Stated plainly, fairness valuations of targets are validated by acquirers' subsequent goodwill impairments.

4.2.4 Untabulated validation results

In untabulated results, I find that the fairness valuation can forecast shareholder voting on transactions: When the deal price is higher relative to the fairness valuation, shareholders vote more strongly in favor of the transaction and are less likely to vote it down. These results hold after conditioning on the *Premium*, industry-time fixed effects, provider fixed effects, and firm-level controls. This result is relevant for two reasons: First, if target shareholders are informed and vote in their interests (Fischer, Gramlich, Miller, and White, 2009), then their voting behavior provides further validation. Second, it could be useful to directors to be able to forecast how their shareholders will vote on transactions, in order to prevent a costly unwinding. I also find that the fairness valuation of the target outperforms the baseline valuation model ($V_t{}^1$) from Frankel and Lee (1998) in forecasting abnormal returns for withdrawn targets. The coefficients and R^2s are higher in Table 5 than they are in a regression in which I substitute $V_t{}^1$ in place of the fairness valuation, and my result also holds when I include the implied return as a control. Finally I find that the loadings on *FOPRa* in Table 6 improve in size and significance when I used weighted least squares, weighting each observation by the ratio of the target's market capitalization to the acquirer's, though they remain significantly smaller than 1. Collectively, these results all indicate that fairness valuations have real economic content, incremental to the information impounded in the target's pre-deal market price.

4.3 Bias tests

4.3.1 Hypothesis

Next, I develop tests for provider bias. It is tempting to look for bias in the *level* of the DCF valuation, such as, e.g., examining how this quantity deviates from the target's pre-deal market price. But I argue that this is the wrong test, for two reasons. First, there could be fundamental mispricing or deal-level synergies, which would warrant giving a target a valuation that differed from its pre-deal market price. Stated differently, there is no benchmark for the "true" value of the target in this setting, relative to which bias can be credibly measured by a researcher. Second, biased fairness opinion providers would not have monotonic incentives to increase the valuation of the target. In fact, holding the deal price fixed, increasing the fairness valuation will make the deal price look less attractive and possibly less fair by comparison, and thus undermine support for the deal. Thus, I instead test for bias in the *components* of the DCF valuation—the discount-rate and perpetuity-growth rate assumptions.[37]

To understand these tests, consider the following stylized example. Suppose that a deal price is negotiated independently by acquirer and target management, and a fairness-opinion provider is brought in to value and testify to the fairness of this price. The analyst producing the DCF valuation model will face some constraints on her ability to toggle the FCF forecasts—whether because these FCF forecasts are produced in collaboration with management, or because she faces some outrage constraint or reputational cost from assuming FCF growth rates that differ dramatically from industry consensus. If FCF forecasts are somewhat constrained, then the analyst could assume lower discount rates and/or higher perpetuity growth rates in order to rationalize higher deal prices (and, symmetrically, higher discount rates, and/or lower perpetuity growth rates to rationalize lower deal prices). If the analyst were completely independent and unbiased by the negotiated transaction price, this would not be the case.

A memorable example that illustrates this is Morgan Stanley's coverage of Snap in the weeks after

[37] An unlevered DCF valuation model is composed of three essential components: Near-to-medium term free-cash-flow forecasts; a discount rate; and a perpetuity growth rate or terminal multiple. I focus on testing for bias in discount rates and perpetuity growth rates, because these components of the valuation are disclosed with sufficient regularity in the merger proxy, and economic and financial theory provide normative benchmarks for these assumptions. Near- to medium-term free cash flow forecasts, on the other hand, are not typically disclosed, are not easily operationalized (since they vary over time, and DCF modelers have discretion over how many periods they explicitly forecast), are not as frequently disclosed, and do not have a reliable, credible normative benchmark.

its IPO. Morgan Stanley had been one of the lead underwriters for the IPO, but once the quiet period ended, it initiated analyst coverage on Snap with a report that assumed a WACC of 9.7%, and generated a $28 per share price target, and, consequently, an overweight recommendation. But after this analysis was published, other analysts uncovered an error in its tax-rate forecast that had caused it to overstate forecasted EBITDA by as much as 40%. In its "correction" to its initial report, Morgan Stanley also, serendipitously, uncovered an "error" in the WACC assumption that, when "corrected" to 8%, precisely canceled out the effects of adjusting its tax rate assumption, allowing it to maintain its original price target and overweight recommendation (Di Maggio, Esty, and Saldutte, 2018).

4.3.2 Bias Test #1: Discount Rates

My goal in this test is to detect whether (and, if so, how much) providers systematically distort their discount-rate assumptions from a normative benchmark, in order to *ex post* rationalize negotiated deal prices. One approach to this problem would be to estimate a WACC for each target company in my sample. But this would depend on subjective researcher assumptions about factor risk prices, and discretionary choices about how to estimate factor risk exposures. So instead, I use economic and financial theory to build regression specifications that achieve the goal—capturing how much providers deviate their WACCs from a normative benchmark to rationalize the deal price—in a way that does not depend on researcher discretion, and maintains parsimony and interpretability.

Under Modigliani-Miller (1958) assumptions, a company's WACC is a function of its asset betas, i.e., the factor risk exposures of its operating cash flows, and thus insensitive to capital structure (in stark contrast to its equity beta), and predominantly a function of the company's line of business. Damodaran (2012) recommends estimating asset betas (also called unlevered betas) at the industry level (for firms in the same country), and this is common practice in the industry. The intuition behind this is that individual company beta estimates are noisy, and, as such, relative to the option of using the average industry asset beta, the cost of introducing measurement error would outweigh any benefit of trying to measure real idiosyncratic firm-level within-industry variation in asset betas.[38] Thus, under Modigliani-

[38]Damodaran publishes industry WACCs on his website, which are widely referenced in the finance industry. Several other financial services and valuation advisory firms publish industry-benchmark WACCs as well. This industry-level WACC will not usually be the *end point* of a DCF analysis—given the caveats to Modigilani-Miller I discuss (bankruptcy costs, liquidity premia, etc.) and given that firms can operate in multiple lines of business—nor the only discount rate for all firm investment decisions. I claim only that it is the major reference point, and accounts for most of the normative variation in tax-unaffected WACCs, for firms with similar size and liquidity.

Miller assumptions, a normative benchmark WACC would be one that is constant within the industry at any point in time, and, in a regression specification, we could test for *bias* in WACC assumptions simply by including an industry-time fixed effect, measuring how the residual in this regression reacts to deal pricing.

The major caveats to this Modigliani-Miller assumption, as established by a long stream of corporate finance literature (Stulz, 2000), include: (1.) Bankruptcy *costs*, (2.) the interest expense tax shield, (3.) liquidity, and (4.) idiosyncratic risk (under the assumption that investors are not fully diversified).[39] Thus, my basic research design is to estimate regressions in which I use industry-time fixed effects to do most of the work in establishing the normative benchmark WACC, while using either filters or linear controls to account for these caveats,[40] and regress the WACC assumption on the negotiated transaction *Premium*. Using the transaction *Premium* as the measure of the deal pricing makes this specification especially credible as a test for bias: If some important determinant of discount rates were omitted from my model (or, put differently, if there were error in my regression's implicit normative benchmark) that factor should affect the target's pre-deal market valuation, rather than the transaction premium *per se*.[41] Economic theory predicts that the premium should be a function of deal-level synergies and bargaining power, and there is no *a priori* reason to expect these to vary with measurement error in the normative discount rate. (I note also that, empirically, in my sample, when providers provide both a base case and a synergy-inclusive valuation, they do not vary their discount-rate assumptions between the two.)

Table 8 reports the results of these tests: The sequence of columns behave as economic and financial theory—and the hypothesis that providers are somewhat biased—would predict. Column 1 reports a

[39] Note that bankruptcy *costs* are different from bankruptcy *risk*, and the interest expense tax shield does not introduce bias into my tests, since the WACC used in unlevered DCF analysis is the tax-unaffected WACC charged by investors.

[40] This specification is thus a close analogue to the conventional practice of starting with industry WACCs, and then making adjustments for size premia, liquidity, and other idiosyncratic firm factors from there. A natural question is why I do not include disclosure and information-quality proxies in my discount-rate regression model. There is a long-running debate in accounting and finance research over whether disclosure and information quality are a priced risk factors, with both theoretical work (see, e.g., Easley and O'Hara (2004), Hughes, Liu, and Liu (2007), and cites therein), and empirical work (see, e.g., Francis, LaFond, Olsson, and Schipper (2005), Core, Guay, and Verdi (2008), and cites therein) taking each side. Reviewing the controversy as a whole, Beyer et al. (2010) conclude that "The relation between voluntary disclosures, financial reporting quality attributes, and cost of capital is still inconclusive. We cannot draw unambiguous conclusions [on] whether... [there is a link] ...between information quality and cost of capital." Additionally, my conversations with practitioners and reviews of business appraisal textbooks do not indicate that they consciously factor disclosure and information quality into their discount rates. As such, in the interest of parsimony and interpretability, I omit these factors from my model.

[41] Note also that the fairness opinion valuation is performed only once a firm offer is on the table and being considered, so the causal direction cannot be the reverse.

simple univariate regression of the average discount rate on the premium—note that the coefficient in this regression is *positive* but the R^2 is extremely low. In Column 2, I add an industry-time fixed effect. Consistent with "Modigliani-Miller benchmark" above, the R^2 increases to 61%, in this parsimonious specification.[42] In subsequent columns, I use parsimonious cuts and controls to account for the caveats to Modigliani-Miller, discussed above, and thereby implement my tests for bias.

In Column 3, I use a very spare cut, dropping target firms with market capitalizations below $100 million and/or negative profitability. As indicated by the negative, significant coefficient on *Premium*, I find evidence of bias in this specification. This indicates that for two (large, profitable) targets in the same industry at the same time, the one with the higher premium is assigned a lower discount rate on average. Interpreting the size of the coefficient, a 10% higher transaction premium leads to an average decrease in the discount rate of .28%. In Column 4, I instead use a robust set of linear controls. The controls for *Assets* and *MCAP* account for size premia; the controls for *Average Volume* and *Volatility* account for liquidity and idiosyncratic risk; the *Altman Z-score* accounts for bankruptcy risk; and the controls for *Sales Growth* and *Ave. Return* account for additional dimensions of corporate performance (I do not include ROA as a separate control, since the *Altman Z-score* loads heavily on it). In this specification, the coefficient on *Premium* is negative and statistically significant, as predicted, indicating that fairness opinion providers adjust their discount rates downward (upward) in response to higher (lower) negotiated deal prices.

In Column 5, I replace the industry-time fixed effect with an industry-time-provider fixed effect. This ensures that the effect is not driven by endogenous matching between providers (who may make idiosyncratic risk pricing assumptions) and targets, and facilitates the bias interpretation: With the provider fixed effect included, the regression results indicate that, on average, a single provider valuing two targets in the same industry in the same year will—after adjusting for size, distress, etc.—assign a lower WACC to the one with the higher negotiated premium. Also noteworthy is that the R^2 in this regression increases to 99%.

[42] In this specification, I do not include any controls for, or perform any matching on, size and bankrupty risk, and there is no measurable loading on *Premium*. This is not surprising *ex ante*, given the caveats to Modigliani-Miller discussed above. I inspect my data and find that the distribution of WACCs in my sample and the distribution of premia are right-tailed—there are a large number of transactions with both very high premia and discount rates, such as buyouts of small pharmaceutical companies, and acquirers "rescuing" targets from bankruptcy and severe distress. In these transactions, the target company's stock price—and thus measured premium—is highly volatile, and it appears that the target company's idiosyncratic risk appears to get factored into the discount-rate assumption. The cuts and controls in subsequent specifications account for these issues.

The estimated coefficient from my most robust specification, in Column 5, implies that in response to a 10 percentage point increase in the negotiated premium, providers adjust their discount rates downwards by .19 percentage points. This is sizeable relative to the relevant benchmark: Using a steady-state Gordon-Growth approximation, and in-sample average premia, discount rates, growth rates, and market value of equity to debt ratio, a .19 percentage point decrease in the WACC would generate a 3 percentage point increase in the equity valuation in terms of the implied premium. In other words, the estimated level of discount-rate bias is about 30% as large as it would be if providers were solely and mechanically rationalizing all of the variation in residual premia solely through discount-rate adjustments. Also notable is that the estimated coefficient is within a tight bound of plausibility for the bias interpretation: It would be impossible for the effect to be more than 3.3x as large as estimated (since this would imply that discount rates were being toggled *more* than necessary to fully rationalize deal prices). This gives additional credence to the bias interpretation of this result.

4.3.3 Discount Rate Bias Drivers

Table 9 reports the results of regressions estimating heterogeneity in, and drivers of, this bias. Column 1 repeats the main discount-rate bias specification, from Column 4 of Table 8, to ease comparison. Column 2 interacts the *Premium* with *Non-independent*—an indicator that evaluates to 1 if the provider is not independent (i.e., is an investment bank)—and thereby tests for whether independent providers exhibit less bias. As indicated by the coefficient on *Non-independent x Premium*, I do not find that independent providers are measurably less biased in the full sample.

Columns 3 and 4 test for whether the discount-rate bias was affected by a judicial "shock"—the Delaware Chancery court's decision in *In re: Southern Peru Copper Corporation*, decided on October 14, 2011. In that case, the court found in favor of plaintiff minority shareholders who alleged that they had not received fair value in a conflicted transaction. In his decision, Vice Chancellor Strine specifically criticized Goldman Sachs for a *post hoc* process in valuing the transaction for a fairness opinion, noting how Goldman increased its discount-rate assumptions between drafts of the "board book" as the terms were revised. Goldman's final, higher discount-rate assumption generated a valuation that came in more than 30% lower than Southern Peru's standalone market value at the time (Strine, 2011). This case was widely perceived as embarrassing for the fairness opinion provider and as a

signal of intensifying legal scrutiny. (In my conversations with practitioners, several independently identified it as the most significant case, in terms of its impact for providers, in my sample period.) Anecdotal evidence suggests that many large investment-bank providers revised the internal control of their fairness valuations, increasing the standardization of discount-rate estimation techniques (i.e., specifying a quantitative algorithm for how WACC would be determined, and only allowing adjustments with special justification), and documentation of the sequence of drafts of the valuations used in their M&A deals.

Column 3 interacts *Premium* with *Post*, an indicator that evaluates to 1 for transactions announced after the decision date. As indicated by the positive, significant coefficient on *Premium x Post*, the discount-rate bias is lower in the post-*Southern Peru* period. Interpreting, the coefficient on *Premium* implies that, in the pre-period, providers would adjust their discount rate assumptions upwards (downwards) by .30 percentage points in response to a 10 percentage point decrease (increase) in the premium; but in the post-period, this adjustment is mitigated to .12 percentage points.

Column 4 estimates a diff-in-diff (DID), further interacting with an indicator for *Non-Independent*. The assumption of this diff-in-diff is that the non-independent (investment bank) providers were differentially exposed to the scrutiny and consequences of the *Southern Peru* case. There are at least two reasons why this is likely to be the case: (1.) Goldman, who Chancellor Strine specifically criticized, was an investment bank; (2.) Investment bank providers are typically involved in the transaction from an earlier stage, as compared to independent providers, and thus have more drafts of valuations that could be discovered in shareholder litigation—as such, increased documentation and standardization would constrain them relatively more. As indicated by the positive, significant coefficient on *Non-independent x Post x Premium*, the investment bank providers *differentially* decreased their discount-rate bias in the post-*Southern Peru* period. Interpreting all of the interactions: The coefficient on *Premium* in Column 4 estimates the baseline discount-rate bias for independent providers in the pre-period; the negative, significant coefficient on *Non-independent x Premium* indicates that, in the pre-period, investment bank providers were *more biased* in their discount-rate assumptions; while the positive coefficient on *Non-independent x Premium x Post* indicates that they differentially improved in the post-period (and, indeed, the relative point estimates suggest that this difference fully reversed); and the insignificant coefficient on *Premium x Post* indicates that there was no measurable change in the amount of bias

in the post-period for the independent providers. Stated simply, the investment bank providers were more biased in their discount-rate assumptions in the pre-period, but they more-than compensated for this differential in the post-period.

In untabulated results, I test for parallel trends, and find no evidence that independent and non-independent providers exhibited non-parallel trends in any of the three years prior to the *Southern Peru* case or any combination of them. I also find that when I replace the indicator for *Non-Independent* with an indicator for *Goldman*, the coefficient of interest, *Goldman x Post x Premium* is even larger—i.e., Goldman differentially mitigated its discount-rate bias in the post-*Southern Peru* period. While I cannot conclusively rule out the possibility that some other contemporaneous factor could have driven this change, the differential response from Goldman, and the concentration of the change around the decision date, suggests that judicial scrutiny was likely the driver.

4.3.4 Bias test #2: Perpetuity growth rates

Table 10 reports the results of tests for bias in the perpetuity growth rate. In Columns 1 through 3, I use a linear specification, regressing the average perpetuity growth rate assumption from the DCF on the negotiated deal premium and relevant controls. In the first two columns, a univariate specification and one with industry-year fixed effects, respectively, the coefficient on *Premium* is not measurably different from zero. In Column 3, I add controls for target firm characteristics that could plausibly forecast future perpetuity growth rates. Here, the coefficient is positive and significant, indicating that residualized perpetuity growth rates increase with residualized deal premia. This is somewhat *suggestive* of bias, but not *dispositive*. While there is no obvious explanation for this result, aside from bias, economic theory does not rule out the possibility that firm-level perpetuity growth rates may idiosyncratically vary within industry and after controlling for observables. In other words, the normative benchmark implicit in this linear specification is not highly credible.

As such, in Columns 4 through 7, I switch to a specification where the dependent variable of interest is an indicator that evaluates to 1 whenever the lower bound of the perpetuity growth rate assumption is greater than 5%. Such an assumption would imply that the company is expected to grow perpetually *at least* at a rate greater than U.S. GDP has grown in any year since 1980. I characterize these perpetuity

growth rates as an upwardly-biased *error*.[43] Columns 4 and 5 report the results of a linear specification, while Columns 6 and 7 employ a logit specification, with controls sequentially added in each specification. As is evident, in Columns 4 through 7, the coefficient on *Premium* is never statistically significant. In other words, I find no evidence that providers bias their perpetuity growth rate assumptions beyond the constraints of economic theory to rationalize deal prices.

This could be a matter of statistical power, but it also illustrates a simple, but important fact: Providers are at least somewhat *constrained* in biasing their DCFs—they do not rationalize deal prices with transparently implausible assumptions. This helps to square the circle on how fairness valuations can be surprisingly valid, as indicated in Tables 5, 6, and 7, despite the apparently biased implicit incentives of the providers.

4.3.5 Bias Test #3: The Case of MBOs

In Table 11, I report the results of exploratory, descriptive regressions of the behavior of fairness opinion valuations in MBOs. In Column 1, I first establish, consistent with prior literature, that target shareholders in MBOs in my sample receive lower premia, even after conditioning on industry-time fixed effects and linear matching for general firm characteristics.[44] This raises the question of why fairness opinion providers declare these deals fair, and how they rationalize these lower prices. Thus, in Columns 2 through 5, I examine how the DCF valuations behave in MBO transactions, relative to non-MBO benchmarks.[45]

In Column 2, the dependent variable of interest is *DCF mean÷Deal Price*. Here, the coefficient on MBO is positive and highly significant, with a point estimate around 13%. It is most meaningful to

[43] As some practitioners have noted, they sometimes assume perpetuity growth rates above 5% not because they actually forecast perpetual growth at that rate, but because it allows them to achieve a plausible terminal value for companies that will still be growing quickly at the terminal date. In these cases, the perpetuity growth rate is in fact a weighted average of forward growth rates, with nearer-term growth rates receiving a much higher weighting than the actual perpetuity growth rate. This would be a problem for my inference if I found a positive result in these tests, but in fact I do not.

[44] The point estimate here is approximately -11%, whereas prior literature generally reports estimates closer to -5%. The point estimate in Column 1 is statistically measurably less than zero, but it is not estimated with much precision. As such, I emphasize the sign—and not the size—of the coefficient in this specification.

[45] I cannot estimate the regressions in Columns 2 through 5 with precision with an industry-time fixed effect, and I instead use separate industry- and time- fixed effects, and add a provider fixed effect. This introduces minor inconsistency across the Columns of this table, but this is not a problem for the inferences. Conceptually, the *Premium* regression tests what premia target shareholders fetch in MBOs relative to the counterfactual of a non-MBO transaction. Since there is a takeover market in each industry in each year, the industry-year FE is important to establish that counterfactual. However, it is less clear that the normative benchmark for, e.g., the *DCF Range* (or other dependent variables in Columns 2 through 5) should be an industry-time average. For these tests, it seems more important to include provider fixed effects, to understand how providers change their valuation practices in conflicted transactions.

interpret this result in relation to the result in Column 1, which showed that MBOs fetch an 11% lower premium. Taking the two columns together, this suggests that the *average* DCF valuation of the MBO "pushes back" against the lower negotiated deal price, or, in other words, it recognizes the relative undervaluation of these deals.

This might seem to only deepen the mystery of how the fairness opinion providers declare these deals fair. But it is easily resolved in Columns 3 through 5, which separately examine the upper bound of the DCF valuation relative to the deal price (*DCF high÷Deal Price*), the lower bound (*DCF low÷Deal Price*), and the range (*DCF Range*). Collectively, these show that in MBOs, providers dramatically expand their valuation ranges, so the *lower bound* of the range is comparably situated, relative to the deal price, to where it is in non-MBO transactions. Thus, by expanding the range, they can declare the deal within the range of fairness, even as the mean valuation suggests that the deal is relatively poorly priced. This setting illustrates the mixture of information and bias in the fairness-opinion valuation process. It also shows how the lack of constraints on the appropriate range can render any deal "fair."

5 Discussion and conclusion

I find evidence for truth and bias in M&A target fairness valuations. They appear to be able to detect fundamental mispricing, and impound expected synergies from the deal, both of which could make them incrementally useful even to directors of public targets. They also exhibit bias that attenuates their validity, and, thus, their usefulness.

While there is obvious tension in these findings, there is no strict contradiction. Consider the following example: Suppose that the negotiated deal price is $120, and the provider's DCF valuation gives a range of $72 to $128, for an average value of $100. Suppose that the true value of the target is, in fact, $100—so the average valuation is fully valid. Now, suppose that the negotiated deal price changes to $130. A biased provider would toggle the valuation in order to include the new deal price in the valuation range, and/or keep the average valuation reasonably close. Suppose the provider is biased and shifts the valuation range to $76 to $132, an average of $104. In this case, the new average valuation would still indicate (correctly, by construction) that the deal price was high, and point in the direction of the correct valuation—although the bias would attenuate the validity.

As this example shows, fairness valuations can provide a partially valid signal about the target's value, even as bias attenuates the validity and usefulness of this signal. The bias is created by the implicit and explicit economic incentives in the process as it is currently designed, and the discretion permitted in the valuation process. The validity in the signal is created by the constraints of the fundamental valuation process (providers do not have unlimited discretion in toggling discount rates or perpetuity growth rates, and have professional pride and reputational concerns) and the private information and expertise employed (providers have privileged access to managerial plans and forecasts that are relevant to understanding deal synergies).

This tension also has policy implications, suggesting that third-party appraisal has potential value, but the process should be reformed. Critics have argued that fairness-opinion valuations are least necessary in public company transactions—if the market is efficient, a public target's price is a summary statistic for its value, and so a premium is, *ipso facto*, more than fair. Under this view, all that a target director needs to know when considering a takeover offer is the premium. But my evidence shows that there are at least two flaws in this argument: First, markets may be imperfectly efficient, and may fundamentally misprice some stocks, particularly smaller companies, or those where managers have incentives to distort the information environment. If so, the target's market price is not a summary statistic for its fundamental value. Second, the target's market price does not include prospective synergies from the deal, which should affect the target firm's negotiating position and willingness to sell. Directors often do not have the expertise in forecasting, M&A, and valuation required to value these prospective synergies on their own, and may benefit from third-party valuators with this expertise. Given that fairness-opinion valuations have this potential value *even* in the "hardest" case, there could be a valuable role for third-party appraisal in other corporate transactions as well.

Additionally, my bias tests have obvious implications for reforming the fairness-opinion valuation process. Directors, regulators, or opinion providers could use regulation, professional standards, and commitment mechanisms to re-structure the process. For example, firms could have an independent committee choose a fairness-opinion provider for a term of several years, before any prospective deals are being considered and favored by management, and commit not to hire them for any other roles in transactions. Providers could be blinded to the negotiated deal price before they produce their valuations and ranges. These would remove both the incentive and the opportunity to simply *ex post*

rationalize the negotiated deal price. Or, an adversarial structure could be built into the sale process—boards could hire one team of advisors to make their best case against the deal, and thereby motivate them to uncover possible conflicts, or better alternatives. Fairness opinion providers could pre-commit to discount rate assumptions that are fixed by industry-size groupings each year, to limit discretion and bias in this component of the DCF valuation. Regulators or professional standards bodies could give more precise definition and guidance on "fairness," and constrain acceptable valuation ranges. The results in Table 9 suggest that even judicial attention and criticism—without explicit or financial penalties for the providers—can be effective in reforming practice.

Which, if any, of these reforms are adopted is up to regulators, boards, and investment professionals. My findings provide useful evidence for those who seek to improve the governance of the M&A process.

References

Angrist, J. D. and J. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion* (1 ed.). Princeton, NJ: Princeton University Press.

Armstrong, C. S., W. R. Guay, and J. P. Weber (2010). The role of information and financial reporting in corporate governance and debt contracting. *Journal of Accounting and Economics 50*(2-3), 179–234.

Bebchuk, L., A. Cohen, and C. C. Wang (2014). Golden parachutes and the wealth of shareholders. *Journal of Corporate Finance 25*, 140–154.

Bebchuk, L. A. and M. Kahan (1989). Fairness Opinions: How Fair Are They and What Can Be Done about It? *Duke Law Journal, 1989*(1), 27–53.

Becher, D. and J. Cai (2010). Do Shareholders Listen? M&A Advisor Opinions and Shareholder Voting. *2010 American Law and Economics Association meetings* (215).

Beyer, A., D. A. Cohen, T. Z. Lys, and B. R. Walther (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics 50*(2-3), 296–343.

Boone, A. L. and J. H. Mulherin (2007). How Are Firms Sold? *Journal of Finance 62*(2), 847–875.

Bowers, H. M. and W. R. Latham (2006). Information asymmetry, litigation risk, uncertainty and the demand for fairness opinions: Evidence from U.S. mergers and acquisitions, 1980-2002. *SSRN Working Paper Series* (2004), 1980–2002.

Bradshaw, M. T., L. D. Brown, and K. Huang (2013). Do sell-side analysts exhibit differential target price forecasting ability? *Review of Accounting Studies 18*(4), 930–955.

Brav, A. and R. Lehavy (2003). An Empirical Analysis of Analysts' Target Prices: Short Term Informativeness and Long Term Dynamics. *The Journal of Finance LVIII*(5).

Cain, M. D. and D. J. Denis (2013). Information Production by Investment Banks: Evidence from Fairness Opinions. *The Journal of Law and Economics 56*(1), 245–280.

Chen, S., J. Thomas, and F. Zhang (2016). Spring-loading future performance when no one is looking? Earnings and cash flow management around acquisitions. *Review of Accounting Studies 21*(4), 1081–1115.

Coates, J. C., D. Palia, and G. Wu (2017). Reverse termination fees in M&A: Design, Signals, and Bidder Returns.

Core, J. E., W. R. Guay, and R. Verdi (2008). Is accruals quality a priced risk factor? *Journal of Accounting and Economics 46*(1), 2–22.

Damodaran, A. (2012). *Investment Valuation* (3 ed.). New York: Wiley.

Davidoff, S. M. (2006). Fairness Opinions. *Wayne State University Law School Legal Studies Research Paper Series No. 07-07* (07), 1558–1623.

Dechow, P. M., A. P. Hutton, and R. G. Sloan (2000). The Relation between Analysts' Forecasts of Long-Term Earnings Growth and Stock Price Performance Following Equity Offerings. *Contemporary Accounting Research 17*(1), 1–32.

Di Maggio, M., B. C. Esty, and G. Saldutte (2018). Valuing Snap After the IPO Quiet Period (A).

Dugar, A. and S. Nathan (1995). The Effect of Investment Banking Relationships on Financial Analysts' Earnings Forecasts and Investment Recommendations. *Contemporary Accounting Research 12*(1), 131–160.

Easley, D. and M. O'Hara (2004). Information and the Cost of Capital. *LIX*(4), 1553–1583.

Fiflis, T. (1992). Responsibility of Investment Bankers to Shareholders. *Wash. ULQ 70*(2).

Fischer, P. E., J. D. Gramlich, B. P. Miller, and H. D. White (2009). Investor perceptions of board performance: Evidence from uncontested director elections. *Journal of Accounting and Economics 48*(2-3), 172–189.

Fischer, P. E. and P. C. Stocken (2001). Imperfect information and credible communication. *Journal of Accounting Research 39*(1), 119–134.

Francis, J., R. LaFond, P. Olsson, and K. Schipper (2005). The market pricing of accruals quality. *Journal of Accounting and Economics 39*(2), 295–327.

Francis, J., D. Nanda, and X. Wang (2006). Re-examining the effects of regulation fair disclosure using foreign listed firms to control for concurrent shocks. *Journal of Accounting and Economics 41*(3), 271–292.

Francis, J. and L. Soffer (1997). The Relative Informativeness of Analysts' Stock Recommendations and Earnings Forecast Revisions. *Journal of Accounting Research 35*(2), 193–211.

Frankel, R. and C. M. Lee (1998). Accounting valuation, market expectation, and cross-sectional stock returns. *Journal of Accounting and Economics 25*(3), 283–319.

Ghosh, A., A. Hovakimian, and H. Hu (2018). Do CEOs Pay a Price for Failed Acquisitions? Evidence From Goodwill Impairments. *AAA 2018 Annual Meeting*, 1–49.

Grossman, S. J. and O. D. Hart (1986). The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy 94*(4), 691–719.

Groysberg, B., P. M. Healy, G. Serafeim, and D. M. Shanthikumar (2013). The Stock Selection and Performance of Buy-Side Analysts. *Management Science 59*(5).

Gu, F. and B. Lev (2011). Overpriced shares, Ill-advised acquisitions, and goodwill impairment. *Accounting Review 86*(6), 1995–2022.

Harford, J., J. Stanfield, and F. Zhang (2018). Do insiders time management buyouts and freezeouts to buy undervalued targets? *Journal of Financial Economics*, 1–26.

Hartzell, J. C., E. Ofek, and D. Yermack (2004). What's in It for Me? CEOs Whose Firms Are Acquired. *Review of Financial Studies 17*(1), 37–61.

Hughes, J., J. Liu, and J. Liu (2007). Information Asymmetry, Diversification, and Cost of Capital. *The Accounting Review 82*(3), 705–729.

Jacob, J., S. Rock, and D. P. Weber (2008). Do non-investment bank analysts make better earnings forecasts? *Journal of Accounting, Auditing and Finance 23*(1), 23–61.

Jensen, M. C. and W. H. Meckling (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics 3*(4), 305–360.

Kasznik, R. and M. E. Mcnichols (2002). Does Meeting Earnings Expectations Matter? Evidence from Analyst Forecast Revisions and Share Prices. *Journal of Accounting Research 40*(3).

Ke, B. and Y. Yu (2006). The effect of issuing biased earnings forecasts on analysts' access to management and survival. *Journal of Accounting Research 44*(5), 965–999.

Kisgen, D. J., J. Q. Qian, and W. Song (2009). Are fairness opinions fair? The case of mergers and acquisitions. *Journal of Financial Economics 91*(2), 179–207.

Kothari, S. P. and J. B. Warner (1997). Measuring long-horizon security price performance. *Journal of Financial Economics 43*, 301–339.

Lambert, R. A. and D. F. Larcker (1985). Golden parachutes, executive decision-making, and shareholder wealth. *Journal of Accounting and Economics 7*(1-3), 179–203.

Lim, T. (2001). Rationality and Analysts' Forecasts Bias. *Journal of Finance 56*(1), 369–385.

Liu, J., D. Nissim, and J. Thomas (2002). Equity valuation using multiples. *Journal of Accounting Research 40*(1), 135–172.

Ljungqvist, A., F. Marston, and W. J. Wilhelm (2006). Competing for securities underwriting mandates: Banking relationships and analyst recommendations. *Journal of Finance 61*(1), 301–340.

Modigliani, F. and M. H. Miller (1958). The Cost of Capital, Corporation Finance and the Theory of Investment. *The American Economic Review 48*(3), 261–297.

Oehmichen, J. (1979). The Information Content of Financial Analysts Forecasts of Earnings. *Journal of Accounting and Economics 1*(3), 165–185.

Oster, E. (2017). Unobservable Selection and Coefficient Stability: Theory and Validation. *Journal of Business & Economic Statistics*.

Parloff, R. (2016). How Michael Dell Shortchanged Shareholders While Doing Nothing Wrong. pp. 1–8.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2 ed.). Cambridge, UK: Cambridge University Press.

Perry, S. E. and T. H. Williams (1994). Earnings management preceding management buyout offers. *Journal of Accounting and Economics 18*, 157–179.

Ramanna, K. and R. L. Watts (2012). Evidence on the use of unverifiable estimates in required goodwill impairment. *Review of Accounting Studies 17*(4), 749–780.

Richardson, S., I. Tuna, and P. Wysocki (2010). Accounting anomalies and fundamental analysis: A review of recent research advances. *Journal of Accounting and Economics 50*(2-3), 410–454.

Shleifer, A. and L. H. Summers (1990). The Noise Trader Approach to Finance. *Journal of Economic Perspectives 4*(2), 19–33.

Strine, L. (2011). In RE Southern Peru Copper Corporation Shareholder Derivative Litigation.

Stulz, R. M. (2000). Merton Miller and Modern Finance. *Financial Management 29*(4), 119.

Trueman, B. (1990). On the Incentives for Security Analysts to Revise Their Earnings Forecasts. *Contemporary Accounting Research 7*(1), 203–222.

Verrecchia, R. E. (2001). Essays on disclosure. *Journal of Accounting and Economics 32*(1-3), 97–180.

Zha Giedt, J. (2018). Economic Consequences of Announcing Strategic Alternatives. *SSRN 2018*(May).

Fig. 1.
Cumulative abnormal returns to targets after withdrawal

Figure 1 displays abnormal returns to targets subsequent to their withdrawal of their offers, sorted by the fairness-opinion implied predicted return, as defined in Equation 1, above. This figure is a graphical companion to Table 5 below. Abnormal returns are estimated using the Fama-French three-factor model, with the factor exposures estimated using a 70 day estimation window and a 50 day gap before the event date.

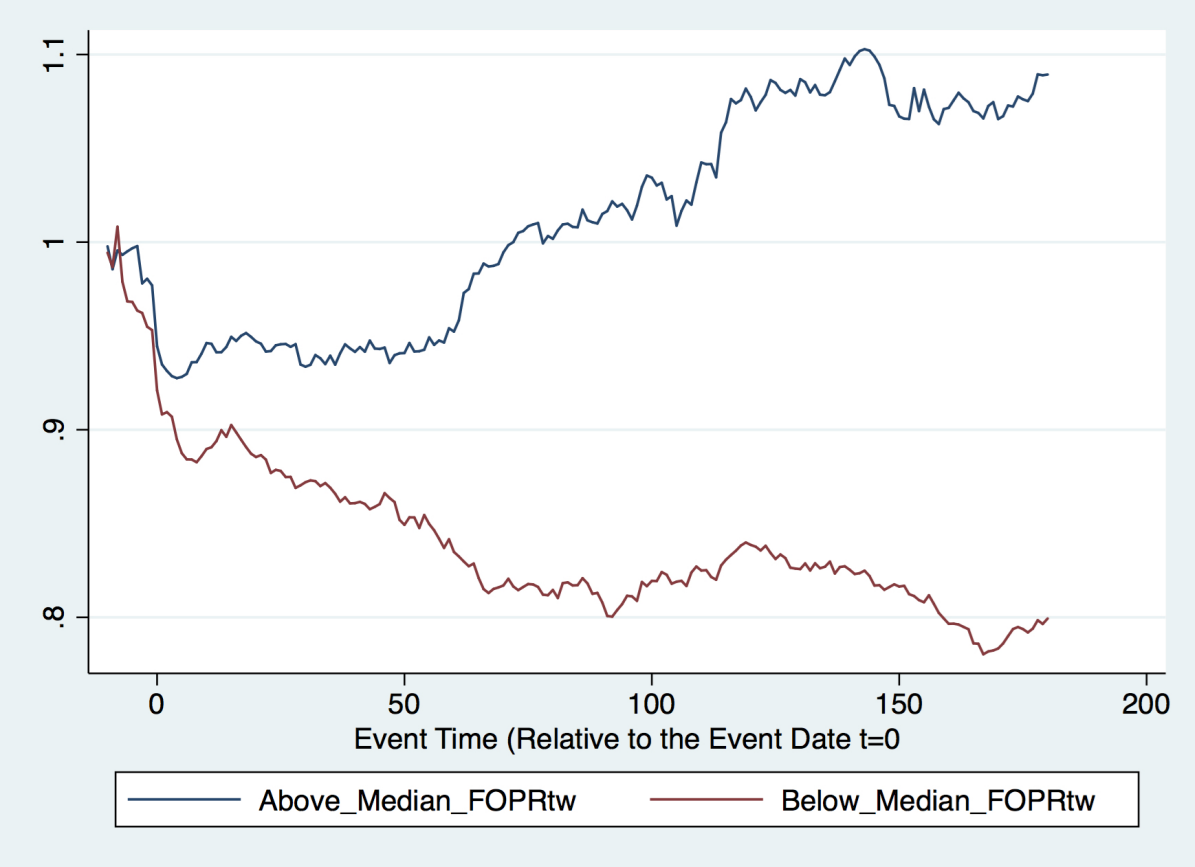

Fig. 2.
Cumulative abnormal returns around announcement date by above- vs. below-median FOPRa

Figure 5 displays gross cumulative abnormal returns to acquirers from *t-1* around the acquisition announcement date, sorted by the fairness-opinion implied predicted return for the acquirer, as defined in equation 3, above. Abnormal returns are estimated with the CAPM, with the factor exposure estimated using a 70 day estimation window and a 50 day gap before the event date.

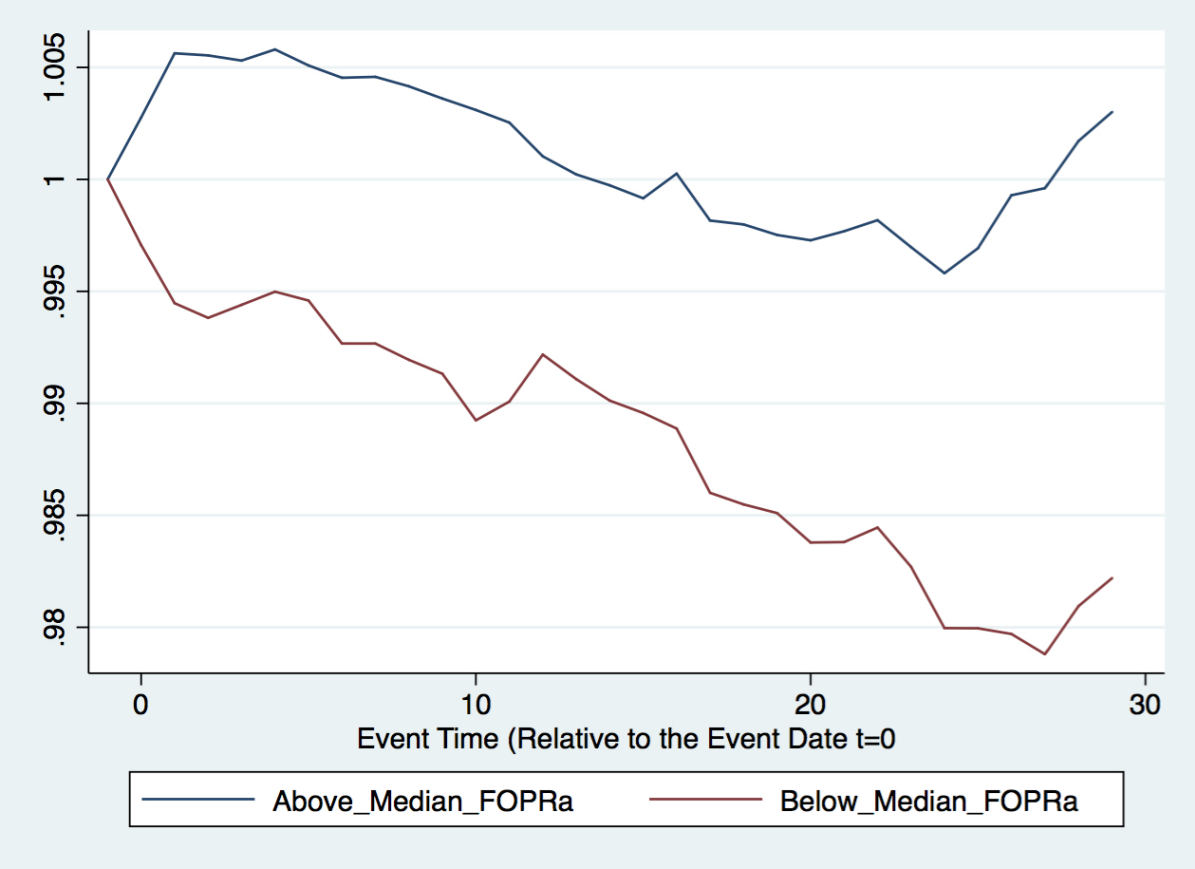

Table 1.
Frequency by year

Table 1 reports the frequency of transactions and fairness opinions in my sample by year.

Transactions

	freq	pct
2006	246	13.51
2007	252	13.84
2008	151	8.29
2009	94	5.16
2010	179	9.83
2011	142	7.80
2012	161	8.84
2013	148	8.13
2014	144	7.91
2015	169	9.28
2016	135	7.41
Total	1821	100.00

Fairness Opinions

	freq	pct
2006	352	11.79
2007	373	12.50
2008	209	7.00
2009	156	5.23
2010	258	8.64
2011	249	8.34
2012	244	8.17
2013	272	9.11
2014	316	10.59
2015	317	10.62
2016	239	8.01
Total	2985	100.00

Table 2.
Firm-level summary stats

Table 2 reports lagged summary statistics. Unscaled variables are in 000's of USD.

Acquirers

	mean	sd
Assets	58326.12	211096.09
Common Equity	10223.63	21274.97
Long-Term Debt	9457.84	35013.39
Acq. Sales Contr.	670.99	3316.02
Employees	36.68	69.49
Net Income	1714.03	5060.83
ROE	0.16	1.00
Sales	15367.79	31306.27
Sales Growth (%)	1.86	19.54
Goodwill	4557.18	10553.60
Observations	892	

Targets

	mean	sd
Assets	4311.87	33867.34
Common Equity	1035.97	5307.38
Long-term Debt	947.77	7786.53
Acq. Sales Contr.	51.63	352.35
Employees	5.24	16.64
Net Income	102.13	672.16
ROE	0.02	2.03
Sales	1652.82	6354.72
Sales Growth (%)	1.27	3.35
Lagged Return (%)	0.96	4.13
Observations	1636	

Table 3.
Deal-level summary stats

Table 3 reports deal-level summary statistics. The number of observations in this sample is higher than the number of observations in Table 2 because some fairness opinions can be matched to Thomson One but not to Compustat.

	mean	sd	min	max
One month premium (%)	38.66	27.19	-1.42	104.82
One-week premium (%)	34.96	23.74	0.41	88.68
One-day premium (%)	33.82	24.81	-1.10	94.69
Management buyout	0.02	0.14	0.00	1.00
Acquirer three-day CAR	0.00	0.07	-0.36	0.45
Offer eventually withdrawn	0.04	0.19	0.00	1.00
Hostile deal	0.00	0.02	0.00	1.00
Target votes in favor (%)	0.96	0.08	0.27	1.00
Observations	1821			

Table 4.
Fairness opinion DCF valuation-level summary stats

Table 4 reports summary statistics on the key DCF-valuation-level variables and constructs that form the core of my analysis. The number of observations is larger than the number of transactions, as in Table 1 because there can be multiple fairness opinions per transaction, but smaller than the total number of fairness opinions, because not all fairness opinions include DCF valuations.

	mean	sd	min	max
FOPRt	0.01	0.36	-0.53	1.09
FOPRa	-0.00	0.04	-0.08	0.10
Deal Price ÷Mean DCF	1.07	0.24	0.71	1.53
DCF Range	0.33	0.21	0.09	0.89
Mean WACC	0.12	0.04	0.06	0.26
WACC low	0.11	0.03	0.07	0.18
WACC high	0.13	0.04	0.08	0.23
Mean perpetuity growth rate	0.00	0.05	-0.11	0.07
Perpetuity growth rate, low	-0.01	0.05	-0.14	0.06
Perpetuity growth rate, high	0.02	0.04	-0.08	0.09
Observations	2820			

Table 5.
Cumulative Abnormal Returns for Withdrawn Targets

Table 5 reports the results of regressions of cumulative abnormal returns (CAR) to withdrawn targets on the fairness-opinion implied predicted return for the withdrawn target (*FOPRtw*), as defined above. I start the event study three days prior to the announced withdrawal of the transaction, because I find evidence of significant price movement in the target in the two days before the formal announcement of the withdrawal of the offer. Column 1 reports the results using the abnormal return from *t-3* to *t+3*, while controlling for the premium. Column 2 replaces the continuous *FOPRtw* measure with a binary indicator of above- vs. below-median *FOPRtw*, to offer a portfolio interpretation of the result. Each subsequent pair of columns repeats this pattern, but substituting horizons of *t+10*, *t+30*, *t+60*, *t+90*, *t+120*, *t+150*, and *t+180* days. Standard errors are bootsrapped, consistent with the recommendation of (Kothari and Warner, 1997), and reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

Panel A: Short- to medium-term CAR

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	t-3→t+3	*t-3→t+3*	*t-3→t+10*	*t-3→t+10*	*t-3→t+30*	*t-3→t+30*	*t-3→t+60*	*t-3→t+60*
FOPRtw	-0.019		0.013		0.009		0.148***	
	(0.03)		(0.03)		(0.04)		(0.05)	
Above-median FOPRtw		0.003		0.031		0.014		0.075*
		(0.02)		(0.02)		(0.03)		(0.04)
One month premium	0.002	0.007	-0.014	-0.010	-0.035	-0.034	-0.082	-0.096
	(0.03)	(0.03)	(0.03)	(0.03)	(0.05)	(0.04)	(0.06)	(0.07)
Constant	-0.055***	-0.057***	-0.041**	-0.060***	-0.042**	-0.051*	-0.018	-0.064*
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Observations	180	180	180	180	180	180	180	180
R^2	-0.0080	-0.0107	-0.0086	0.0031	-0.0070	-0.0060	0.0492	0.0348

Panel B: Medium- to long-term CAR

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	t-3→t+90	*t-3→t+90*	*t-3→t+120*	*t-3→t+120*	*t-3→t+150*	*t-3→t+150*	*t-3→t+180*	*t-3→t+180*
FOPRtw	0.221***		0.267***		0.374***		0.403***	
	(0.07)		(0.09)		(0.13)		(0.15)	
Above-median FOPRtw		0.166***		0.184***		0.212***		0.239***
		(0.04)		(0.06)		(0.07)		(0.07)
One month premium	-0.108	-0.116**	-0.149*	-0.162*	-0.052	-0.081	-0.107	-0.137
	(0.07)	(0.06)	(0.09)	(0.08)	(0.13)	(0.12)	(0.11)	(0.11)
Constant	0.014	-0.086**	0.095*	-0.016	0.052	-0.077	0.085	-0.061
	(0.04)	(0.04)	(0.05)	(0.05)	(0.06)	(0.07)	(0.06)	(0.07)
Observations	180	180	180	180	180	180	180	180
R^2	0.0669	0.0827	0.0683	0.0731	0.0759	0.0565	0.0704	0.0588

Table 6.
Announcement date returns

Table 6 reports the results of regressions of the announcement-date cumulative abnormal returns (CAR) to acquirers. Column 1 reports the results of a simple univariate regression of the CAR on the fairness-opinion implied predicted return for the acquirer (*FOPRa*), as defined above. Column 2 replaces the *FOPRa* with the *Scaled premium*, to provide a benchmark comparison (see discussion of the "baseline model" above). Column 3 regresses the CAR on the *FOPRa* while controlling for the *Scaled premium*. Column 4 replaces the continuous *FOPRa* measure with a binary indicator for an above- (vs. below) median *FOPRa*, to offer a portfolio interpretation of the result. Columns 5 and 6 are estimated on DCF valuations of targets that explicitly include synergies vs. those that do not, respectively. Columns 7 and 8 subset down to DCF valuations that explicitly include synergies and compare those produced by *Independent* fairness opinion providers (as defined above) vs. non-independent providers, respectively. For ease of comparison, the regressions for the subsamples in Columns 5 vs. 6, and 7 vs. 8, are estimated separately, but F-stats and P-values for a test of differences in the coefficient on *FOPRa* and the intercept are reported in the latter of each pair. Standard errors are clustered at the transaction level and reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

	(1) *3-day CAR*	(2) *3-day CAR*	(3) *3-day CAR*	(4) *3-day CAR*	(5) *3-day CAR Synergy=0*	(6) *3-day CAR Synergy=1*	(7) *3-day CAR Synergy=1 Independent=1*	(8) *3-day CAR Synergy = 1 Independent = 0*
FOPRa	0.222***		0.216***		0.081	0.382***	0.327**	0.417***
	(0.05)		(0.06)		(0.07)	(0.09)	(0.13)	(0.11)
Scaled premium		-0.044**	-0.039**					
		(0.02)	(0.02)					
Above-med. FOPRa				0.871**				
				(0.40)				
Constant	0.104	0.465*	0.438*	-0.343	0.258	-0.143	-1.213**	0.846*
	(0.20)	(0.26)	(0.26)	(0.28)	(0.23)	(0.36)	(0.53)	(0.48)
Observations	770	752	752	770	491	279	135	144
R^2	0.0207	0.0052	0.0222	0.0049	0.0009	0.0624	0.0387	0.0799
F-stat (FOPRa)						7.91		0.28
P-value						0.0050		0.6004
F-stat (Intercept)						0.08		8.29
P-value						0.7779		0.0043

Table 7.
Future Goodwill Impairments

Table 7 reports the results of regressions of *Impairment*—an indicator for whether the acquirer recognized a goodwill impairment in the 5 years after the transaction—on the *FOPRa*, as defined the above. Column 1 reports the results of a simple univariate linear probability model with year fixed effects. Column 2 adds a control for the *Scaled premium*, which would be a summary statistic for material overpayment under the "baseline model" discussed above, as well as the firm's lagged goodwill. Column 3 adds controls for size and performance. Columns 4, 5 and 6 report the results from logit specifications, with controls following the same pattern. Standard errors are clustered at the transaction level and reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Impairment	*Impairment*	*Impairment*	*Impairment*	*Impairment*	*Impairment*
FOPRa	-0.675*	-0.697*	-0.678*	-0.039*	-0.042*	-0.041*
	(0.38)	(0.39)	(0.39)	(0.02)	(0.02)	(0.02)
Scaled premium		0.462***	0.469***		0.025***	0.025***
		(0.13)	(0.13)		(0.01)	(0.01)
Goodwill		1.393***	1.610***		0.077***	0.090***
		(0.20)	(0.25)		(0.01)	(0.02)
Assets			-0.000			-0.000
			(0.00)			(0.00)
Sales			-0.000			-0.000
			(0.00)			(0.00)
ROE			-0.763			-0.041
			(2.49)			(0.13)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	973	953	953	973	953	953
R^2	0.0572	0.1078	0.1073			
Pseudo R^2				0.0611	0.1080	0.1105

Table 8.
Discount Rate Bias

Table 8 reports the estimates of regressions of the average WACC discount rate in the DCF valuations on the negotiated deal premium. Column 1 reports the results of the simple univariate regression. Column 2 adds industry-year fixed effects. Column 3 estimates the same industry-year fixed effect regression on a sample that excludes loss firms (prior year *Net Income* <0) and small firms (with market capitalizations below $100m), as indicated by Thomson One, in order to mitigate the outlier effect of distressed and small firms. Column 4 reverts to the original sample (i.e., does *not* drop loss and small firms), but adds linear controls for a variety of firm-level fundamental and market variables. Column 5 replaces the industry-year fixed effect with an industry-year-provider fixed effect. Standard errors are clustered at the transaction level and reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

	(1) *WACC*	(2) *WACC*	(3) *WACC*	(4) *WACC*	(5) *WACC*
Premium	0.012***	0.000	-0.027***	-0.019***	-0.019***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
MCAP				0.000**	-0.000***
				(0.00)	(0.00)
Assets				-0.000***	0.000***
				(0.00)	(0.00)
Altman z-score				-0.152***	-0.131***
				(0.05)	(0.05)
Sales Growth (%)				0.019	-0.913***
				(0.09)	(0.15)
Ave. Volume				-0.000***	0.000***
				(0.00)	(0.00)
Ave. Return				-0.197***	-0.429***
				(0.04)	(0.06)
Ave. Volatility				0.201***	0.269***
				(0.03)	(0.02)
Industry-Year FE	No	Yes	Yes	Yes	No
Industry-Year-Provider FE	No	No	No	No	Yes
Observations	1,817	1,817	1,339	1,188	1,188
R^2	0.0040	0.6079	0.6503	0.7604	0.9900

Table 9.
Discount Rate Bias: Determinants and Constraints

Table 9 reports regressions estimating the determinants of the discount-rate bias identified in Table 8, above. Column 1 repeats the specification from Column 4 of Table 8 as a benchmark. Column 2 interacts *Premium* with an indicator for *Non-independent* providers, to test whether investment banks (vs. independent providers) exhibit more or less discount rate bias. Columns 3 and 4 exploit the *In re: Southern Peru Copper* court decision, discussed in the text. *Post* is an indicator for that evaluates to 1 for transactions announced after the decision date (October 14, 2011). Column 3 interacts *Premium* with *Post*, to test whether the discount-rate bias decreased in the period after the decision. Column 4 implements a diff-in-diff, further adding indicators for *Non-independent x Premium* and *Non-independent x Premium x Post* to test whether investment-bank providers, which were differentially exposed to scrutiny from this decision, differentially decreased their discount-rate bias. Standard errors are clustered at the transaction level and reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

	(1) *WACC*	(2) *WACC*	(3) *WACC*	(4) *WACC*
Premium	-0.019***	-0.026***	-0.030***	-0.017*
	(0.01)	(0.01)	(0.01)	(0.01)
Non-independent x Premium		0.016*		-0.026***
		(0.01)		(0.01)
Premium x Post			0.018*	0.006
			(0.01)	(0.01)
Non-independent x Premium x Post				0.028***
				(0.01)
MCAP	0.000**	0.000*	0.000*	0.000
	(0.00)	(0.00)	(0.00)	(0.00)
Assets	-0.000***	-0.000***	-0.000***	-0.000***
	(0.00)	(0.00)	(0.00)	(0.00)
Altman Z-score	-0.155***	-0.145***	-0.150***	-0.130**
	(0.05)	(0.05)	(0.05)	(0.05)
Sales Growth (%)	0.032	0.008	0.069	0.114
	(0.10)	(0.10)	(0.10)	(0.10)
Ave. Volume	-0.000***	-0.000***	-0.000***	-0.000***
	(0.00)	(0.00)	(0.00)	(0.00)
Ave. Return	-0.201***	-0.213***	-0.203***	-0.194***
	(0.04)	(0.04)	(0.04)	(0.04)
Ave. Volatility	0.202***	0.200***	0.203***	0.189***
	(0.03)	(0.03)	(0.03)	(0.03)
Industry-Year FE	Yes	Yes	Yes	Yes
Industry-Year-Provider FE	No	No	No	No
Observations	1,188	1,188	1,188	1,188
R^2	0.7608	0.7663	0.7614	0.7653

Table 10.
Perpetuity Growth Rate Bias

Table 10 reports the estimates of regressions of the assumed perpetuity growth rate on the negotiated deal premium. The perpetuity growth rate is either hand-collected directly or imputed from the terminal multiple and discount rate. In Columns 1 through 3, the dependent variable of interest is the average perpetuity growth rate. Column 1 reports the results of a univariate regression of this quantity on the transaction premium. Column 2 adds industry-time fixed effects. Column 3 adds linear controls for firm characteristics that relate to future growth potential and provider fixed effects. In Columns 4 through 7, the dependent variable of interest is an indicator that evaluates to 1 when the lower bound of the perpetuity growth rate assumption is greater than 5%. (I characterize this as a mistake, as the perpetuity growth rate of a firm cannot exceed the perpetuity growth rate of GDP, and U.S. GDP growth never exceeded 5% for two decades before my sample period.) Standard errors are clustered at the transaction level and reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mean PGR	*Mean PGR*	*Mean PGR*	*PGR Low >.05*	*PGR Low >.05*	*PGR Low >.05*	*PGR Low >.05*
				Linear	*Linear*	*Logit*	*Logit*
Premium	0.007	0.009	0.029***	0.000	0.000	0.002	-0.004
	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)
Sales Growth			-0.040		0.001		2.538
			(0.22)		(0.00)		(1.57)
ROE			-0.110		0.000		0.211
			(0.12)		(0.01)		(0.45)
Assets			-0.000*		-0.000		-0.001
			(0.00)		(0.00)		(0.00)
Sales			0.000*		0.000		-0.002
			(0.00)		(0.00)		(0.00)
Industry-Year FE	No	Yes	Yes	Yes	Yes	No	Yes
Provider FE	No	No	Yes	No	No	No	No
Observations	1,797	1,797	1,217	1,797	1,218	1,797	158
R^2	0.0004	0.3880	0.5166	0.1817	0.3385		
Pseudo R^2						0.0002	0.2424

Table 11.
Management Buyouts

Table 11 reports the results of regressions exploring the behavior of fairness-opinion valuations in MBOs. Column 1 first regresses the transaction premium on an indicator for *MBO* as well as miscellaneous target firm-level characteristics and industry-year dummies, in order to establish that target shareholders receive lower premia in MBOs relative to relevant non-MBO counterfactuals. The regressions in Column 2 through 5 then explore how the DCF valuations of target respond in MBOs. I drop industry-year fixed effects due to lack of power and instead include separate time- and industry- fixed effects and add provider fixed effects, consistent with the interpretation of these tests. In Column 2, the dependent variable of interest is *DCF_mean÷Deal Price*, i.e., the ratio of average of the DCF valuation of the target to the takeover price. (As such, higher values of *DCF_mean÷Deal Price* indicate a deal that is more undervalued according to the fairness opinion.) In Column 3, the dependent variable of interest switches to *DCF_high÷Deal Price*, i.e., the ratio of the upper bound of the DCF valuation of the target to the takeover price. In Column 4, the dependent variable of interest switches to *DCF_low÷Deal Price*, i.e., the ratio of the lower bound of the DCF valuation of the target to the takeover price. And in Column 5, the dependent variable of interest switches to *DCF_range*, the difference between the two. All dependent variables are scaled for ease of interpretation in percentage point terms. Standard errors are reported in parentheses. Significance levels are indicated by *, **, *** for 10%, 5%, and 1% respectively.

	(1) *Premium*	(2) *DCF÷Deal Price*	(3) *DCF_high÷Deal Price*	(4) *DCF_low÷Deal Price*	(5) *DCF range*
MBO	-11.312**	13.315**	20.392***	3.694	18.777***
	(5.10)	(5.67)	(7.11)	(6.19)	(5.70)
Assets	0.001***	0.000	-0.000	0.000	-0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Sales	-0.000*	-0.000	-0.000	0.000	-0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Sales Growth (%)	-1.383**	-0.772	-1.107	-0.678	-0.425
	(0.61)	(0.78)	(0.96)	(0.70)	(0.64)
ROE	-1.628*	-0.656*	-0.660	-0.577*	-0.087
	(0.93)	(0.38)	(0.46)	(0.33)	(0.31)
Ave. Volume	-0.000**	0.000*	0.000***	0.000	0.000***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Ave. Return	-0.975***	-0.346	-0.938***	0.127	-1.067***
	(0.23)	(0.24)	(0.30)	(0.22)	(0.20)
Ave. Volatility	1.179***	0.199	0.599***	-0.200	0.797***
	(0.15)	(0.18)	(0.22)	(0.16)	(0.14)
Industry FE	No	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	No	No	No	No
Provider FE	No	Yes	Yes	Yes	Yes
Observations	1,234	1,234	1,234	1,234	1,234
R^2	0.6614	0.0835	0.1419	0.1296	0.3932

Appendix A Description of Variables

This table defines variables used in my analysis. Fairness-opinion level data is from a comprehensive sample of merger proxies spanning 2006 through 2016, originally collected by a major provider. The provider hand-collected and coded the DCF valuations (typically in price per share terms) and the name of the fairness opinion provider. Discount rates and perpetuity growth rates and multiples are hand-collected from these filings and matched to their DCF valuations. Transaction-level data are collected from Thomson One. Firm-level fundamentals and market data are collected from Compustat and CRSP, respectively. Abnormal-return estimations are automated using Event Study by WRDS. Shareholder voting data are collected from Voting Analytics, from WRDS Institutional Shareholder Services.

Variable	Description	Computation
3-day CAR	Acquirer abnormal returns from *t-1* to *t+1* around the acquisition announcement date	Estimated using the CAPM and the default estimation parameters for WRDS Event Study (a 100-day estimation window, requiring 70 valid daily returns, with a 50 day gap before the event) and cumulated from *t-1* to *t+1*
FOPRa	Fairness opinion-predicted return for the *acquirer*	(*Mean DCF−Deal Price*) ÷*Acquirer MCAP*
Above-med. FOPRa	An indicator that evaluates to 1 if the *FOPRa* is above the in-sample median	
Scaled premium	The one-month premium scaled to capture its predicted impact on the acquirer under the "baseline model" of no synergies or fundamental mispricing	(*One Month Premium*) * (*Target MCAP* ÷*Acquirer MCAP*)
Independent	An indicator that evaluates to 1 if the fairness opinion provider is "independent," i.e., provides advisory but not financing	Set to 0 if the provider is indicated as having "Provided Financing," "Arranged Financing," or served as "Underwriter" in any deal in my sample
Synergy	An indicator that evaluates to 1 if the DCF valuation is explicitly inclusive of synergies	Hand-collected
CAR, t-3 →t+i	Target abnormal returns from *t-3* to *t+i* around the date of the withdrawal of the transaction.	Estimated using the Fama-French 3-factor model and the default estimation parameters for WRDS Event Study (a 100-day estimation window, requiring 70 valid daily returns, with a 50 day gap before the event) and cumulated from *t-3* to *t+i*.

Appendix A Continued

Variable	Description	Computation
FOPRtw	The fairness-opinion predicted return for the *target* in *withdrawn* transactions	(*Mean DCF*÷*Deal Price*)−1
Above-med. FOPRtw	An indicator that evaluates to 1 if the *FOPRtw* is above the in-sample median	
Impairment	An indicator that evaluates to 1 if the acquirer impairs goodwill in any of the 5 years subsequent to the transaction	
WACC	The mean weighted average cost of capital used in the corresponding DCF valuation, hand-collected from merger proxy	*(WACC low + WACC high)*÷2
Mean PGR	The mean perpetuity growth rate assumption used in the corresponding DCF valuation, hand-collected from merger proxy	For most cases, this is simply calculated as *PGR low + PGR high* / 2. In cases where the valuation uses a terminal FCF multiple, the implied perpetuity growth rate is imputed using the mean WACC and the perpetuity identity, calculated as WACC−(1 ÷*Terminal Multiple*)
Management Buyout	An indicator that evaluates to 1 if the management will have an interest in the acquiring entity, from Thomson One	
DCF range		(*DCF high*−*DCF low*)÷*Deal Price*

Appendix B Stylized timeline of fairness opinion process

This figure is a stylized visual representation of the general timeline of the fairness opinion process. From the perspective of a target firm and its fairness-opinion provider, a typical M&A deal would proceed in roughly this order: (1.) Company management discusses and negotiates general terms of a possible transaction, and the acquiror makes an offer (approx. *t−20*). (2.) Target company management engages a financial advisor to perform a fairness opinion valuation of the company (approx. *t−15*). (3.) The board of directors meets to consider and vote on the acquisition offer. At this meeting, the board is presented with the "board book" of valuation analyses produced by the fairness-opinion provider (*t−1*, in-sample median). The board of directors votes to approve the deal and recommends that shareholders vote in its favor and the deal will be "announced," pending approval (t). (5.) The target distributes the merger proxy in preparation for the shareholder vote. On this date, the fairness opinion becomes public for the first time (*t+69*, in-sample median). (6.) The target company shareholders vote on the deal (*t+109*, in-sample median). (7.) The transaction is formally effected, if it is not withdrawn (*t+122*, in-sample median). (In cases where the transaction is not effected, it is at some point formally withdrawn.)

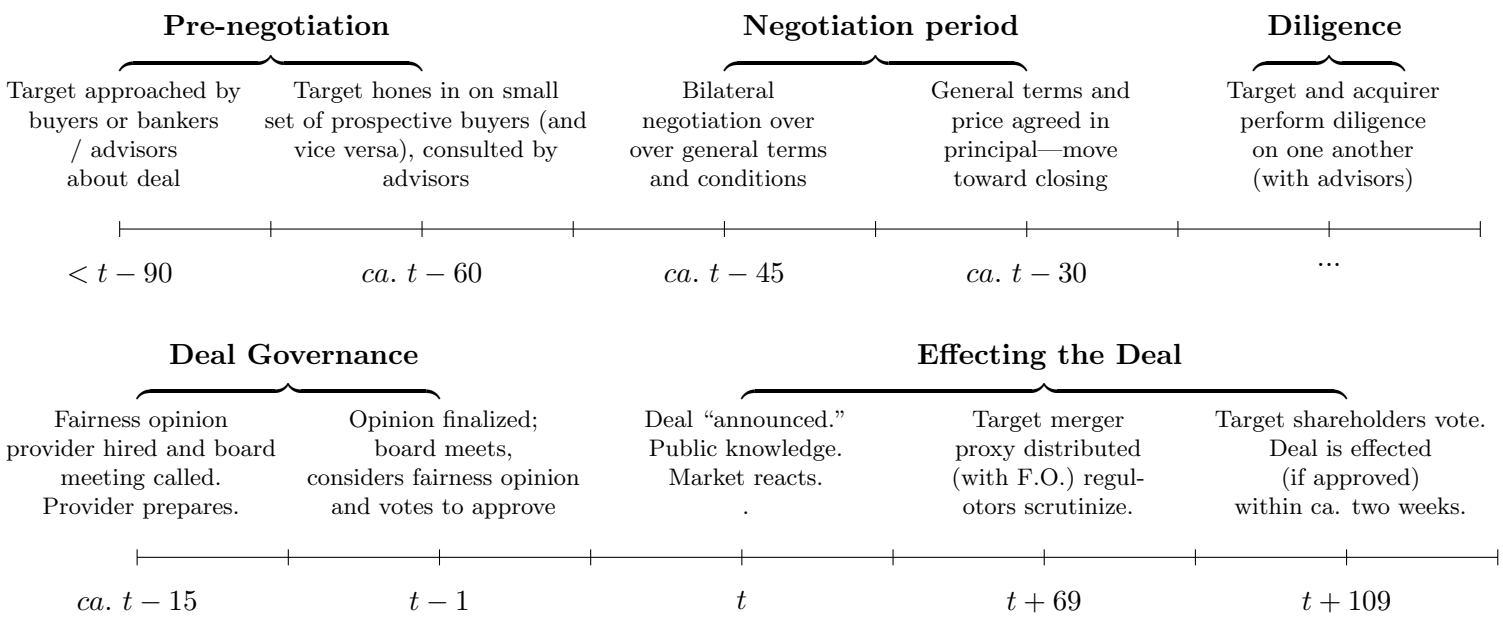

Appendix C Frequency of Opinions by Providers

Appendix C reports the frequency of fairness opinions in my sample by provider. Due to the large number of boutique investment banks and consultancies, all providers with fewer than 30 fairness opinions in my sample are collected under "Other" in this table. (They remain distinct when I use provider fixed effects in my empirical tests.)

	freq	pct	cumpct
Goldman Sachs	313	10.56	10.56
Morgan Stanley	199	6.72	17.28
JP Morgan	192	6.48	23.76
Sandler ONeill	129	4.35	28.11
Merrill Lynch	126	4.25	32.37
Lazard Freres	104	3.51	35.88
Credit Suisse	102	3.44	39.32
Citigroup	98	3.31	42.63
Barclays	89	3.00	45.63
Keefe Bruyette Woods	89	3.00	48.63
UBS Securities	82	2.77	51.40
Evercore	79	2.67	54.07
Jefferies Company	74	2.50	56.56
Deutsche Bank	70	2.36	58.93
Houlihan Lokey	61	2.06	60.99
Lehman Brothers	47	1.59	62.57
Centerview Partners	44	1.48	64.06
Piper Jaffray Co	44	1.48	65.54
Raymond James	40	1.35	66.89
Qatalyst Partners	36	1.21	68.11
Stifel Nicolaus and Company	36	1.21	69.32
RBC Capital Markets	35	1.18	70.50
William Blair and Company	35	1.18	71.68
Moelis and Company	34	1.15	72.83
Greenhill Partners	32	1.08	73.91
Perella Weinberg Partners	31	1.05	74.96
Other	744	25.04	100
Total	2963	100.00	100

Appendix D Selections from proxy in IntercontinentalExchange NYSE-Euronext Merger

1. Relevant opinion-related disclosures in the proxy statement:

Opinion of Perella Weinberg, Financial Advisor to NYSE Euronext[46]

The NYSE Euronext board of directors retained Perella Weinberg to act as its financial advisor in connection with the proposed combination. The board of directors selected Perella Weinberg based on Perella Weinberg's qualifications, expertise and reputation and its knowledge of the business and affairs of NYSE Euronext and the industries in which NYSE Euronext conducts its business. Perella Weinberg, as part of its investment banking business, is continually engaged in performing financial analyses with respect to businesses and their

[46]https://www.sec.gov/Archives/edgar/data/1174746/000119312513187058/d461946ddefm14a.htm

securities in connection with mergers and acquisitions...On December 20, 2012, Perella Weinberg rendered its oral opinion...that the aggregate consideration to be received by the holders of NYSE Euronext common stock...was fair, from a financial point of view, to such holders...

Summary of Material Financial Analyses

Illustrative Premium Paid Analysis. Perella Weinberg took the sum of (a) the product of (1) the exchange ratio of 0.1703 (which exchange ratio was used to calculate the amount of stock included in the standard election amount) and (2) the per share closing price of ICE common stock on December 19, 2012..., $128.31 and (b) $11.27 (the amount of cash included in the standard election amount), and compared the resulting $33.12 implied value per share of NYSE Euronext common stock...to the following: the closing market price per share of NYSE Euronext common stock on the reference date; ... one-week prior to the reference date; and ... one-month prior to the reference date... Perella Weinberg also calculated (based on data provided by FactSet) the median premia paid in selected transactions since January 1, 2007 with values between $5 billion and $15 billion... Based on the premia calculated above...Perella Weinberg selected a representative range of premia paid of 20% to 40%...and compared that to the $33.12 implied per share merger consideration to be received by the holders of NYSE Euronext common stock pursuant to the original merger.

Historical Stock Trading. Perella Weinberg reviewed the historical trading price per share of NYSE Euronext common stock and ICE common stock for the 52-week period ending on the reference date. Perella Weinberg noted that, in the 52-week period ending on the reference date, the range of trading market prices per share of NYSE Euronext common stock was $22.25 to $31.25 compared to the $33.12 implied value per share merger consideration to be received by the holders of NYSE Euronext common stock pursuant to the original merger...

Selected Publicly Traded Companies Analysis. Perella Weinberg reviewed and compared certain financial information for NYSE Euronext and ICE to corresponding financial information, ratios and public market multiples for certain publicly held companies in the exchange operator industry. Although none of the following companies is identical to NYSE Euronext or to ICE, Perella Weinberg selected these companies because they had publicly traded equity securities and were deemed to be similar to NYSE Euronext and ICE in one or more respects including being operators of exchanges...: CBOE Holdings, Inc., CME Group Inc., Deutsche Brse AG, London Stock Exchange plc, The NASDAQ OMX Group, Inc. For each of the selected companies, Perella Weinberg calculated and compared financial information and various financial market multiples and ratios...

Discounted Cash Flow Analysis of NYSE Euronext. Perella Weinberg conducted a discounted cash flow analysis for NYSE Euronext based on the Company budget case and volume growth case forecasts by first calculating, in each case, the present value as of December 31, 2012 of the estimated standalone unlevered free cash flows (calculated as adjusted earnings before interest payments after taxes (assuming tax rates as provided in NYSE Euronext forecasts) plus depreciation and amortization, minus capital expenditures, and adjusting for changes in net working capital and other cash flows) that NYSE Euronext could generate for fiscal years 2013 and 2014 utilizing discount rates ranging from 9.5% to 10.5% based on estimates of the weighted average cost of capital of NYSE Euronext... Estimates of unlevered free cash flows used for these analyses utilized the financial forecasts prepared by NYSE Euronext management and the Goldman Sachs Research forecasts for 2014. Perella Weinberg also calculated, in each case, a range of terminal values utilizing terminal year multiples of last twelve months, or LTM, EBITDA ranging from 7.75x to 8.25x (which range was determined by Perella Weinberg in the exercise of its professional judgment) and discount rates ranging from 9.5% to 10.5% based on estimates of the weighted average cost of capital of NYSE Euronext... The present values of unlevered free cash flows generated over the period described above were then added, in each case, to the present values of terminal values resulting in a range of implied enterprise values for NYSE Euronext. From the range of implied enterprise values, Perella Weinberg derived ranges of implied equity values for NYSE Euronext. These analyses resulted in the following reference ranges of implied enterprise values and implied equity values per share of NYSE Euronext common stock: [$30.65 - 33.65]

Selected Transactions Analysis: Perella Weinberg analyzed certain information relating to selected announced or proposed transactions in the exchange operator industry (transactions analyzed occurred between October 2010 and June 2012, some of which ultimately did not close). Perella Weinberg selected the transactions because, in

the exercise of its professional judgment, Perella Weinberg determined the targets in such transactions to be relevant companies having operations similar to NYSE Euronext. For each of the selected transactions, Perella Weinberg calculated and compared the resulting enterprise value in the transaction as a multiple of reported LTM EBITDA. The multiples for the selected transactions were based on publicly available information at the time of the relevant transaction...

...Based on the multiples calculated above...Perella Weinberg selected a representative range of multiples of 9.5x - 10.5x to apply to reported LTM EBITDA of NYSE Euronext. Perella Weinberg applied such ranges to reported LTM EBITDA to derive an implied per share equity reference range for NYSE Euronext common stock of approximately $32.35 to $36.60 and compared that to the $33.12 implied value per share merger consideration...

2. Fairness opinion letter:

Members of the Board of Directors:

You have requested our opinion as to the fairness from a financial point of view to the holders of Company Common Stock.. of the Aggregate Consideration to be received by such holders in the proposed Merger.

For purposes of the opinion set forth herein, we have, among other things:

- reviewed certain publicly available financial statements and other business and financial information with respect to the Company and Parent, including research analyst reports;

-reviewed certain internal financial statements, analyses, forecasts (the Company Forecasts), and other financial and operating data relating to the business of the Company, in each case, prepared by management of the Company;

-reviewed certain publicly available financial forecasts relating to the Company;

-reviewed estimates of synergies anticipated from the Merger (collectively, the"Anticipated Synergies"), prepared by managements of the Company and Parent;

...With respect to the Company Forecasts, we have been advised by the management of the Company and have assumed, with your consent, that they have been reasonably prepared on bases reflecting the best currently available estimates and good faith judgments of the management of the Company as to the future financial performance of the Company...

This opinion addresses only the fairness from a financial point of view, as of the date hereof, of the Aggregate Consideration to be received by the holders of the Company Common Stock... We express no opinion as to the fairness of the amount or nature of any compensation to be received by any officers, directors or employees of any parties to the Merger, or any class of such persons... Our opinion does not address the underlying business decision of the Company to enter into the Merger or the relative merits of the Merger as compared with any other strategic alternative which may be available to the Company.

We have acted as financial advisor to the Board of Directors of the Company in connection with the Merger and will receive a fee for our services, a portion of which is payable upon the rendering of this opinion and a significant portion of which is contingent upon the consummation of the Merger... During the two year period prior to the date hereof, Perella Weinberg Partners LP and its affiliates provided certain investment banking services to the Company and its affiliates..., including having acted as financial advisor to the Company in a proposed combination transaction with Deutsche Brse AG; [and] may in the future provide investment banking and other financial services to the Company and Parent and their respective affiliates... The issuance of this opinion was approved by a fairness opinion committee of Perella Weinberg Partners LP.

This opinion is for the information and assistance of the Board of Directors of the Company in connection with, and for the purposes of its evaluation of, the Merger. This opinion is not intended to be and does not constitute a recommendation to any holder of Company Common Stock as to how such holder should vote...

Based upon and subject to the foregoing, including the various assumptions and limitations set forth herein, we are of the opinion that, as of the date hereof, the Aggregate Consideration to be received by the holders of Company Common Stock (other than Parent or any of its affiliates) pursuant to the Merger Agreement is fair from a financial point of view to such holders.

Very truly yours,

PERELLA WEINBERG PARTNERS LP