Biased is best: Merger advisors and dialectical Bayesian persuasion

7th August, 2025

Abstract

Advisors to bidding firms are rewarded if bids for target firms are successful irrespective of the value created by the transactions. The use of this success-fee compensation is puzzling since it biases bidder advisors toward pitching deals even when they are value destroying. In a Bayesian persuasion setting, we show that when bidder advisors and their incentives are viewed as one part of a M&A ecology that also includes bidding firms, target firms, and target advisors, success-fee compensation for bidder advisors' will, in fact, yield desirable outcomes for target and bidder firms. We show that the dialectic between biased bidder and biased target advisors produces information that provides a fairly accurate characterization of the value-add potential of the merger. Compensating target advisors simply for blocking acquisitions without producing viable alternatives can be optimal. Making target-advisor compensation sensitive to the bidder's offer can reduce welfare.

Keywords: Advisors, mergers, acquisitions, Bayesian persuasion

JEL classification: D83, D86, G34

1 Introduction

In North America alone, the annual count of M&A transactions averaged over 20,000 between 2014 and 2023, and the annual average aggregate value of these transactions was well over $2 trillion. Advisors, typically investment banks, play important roles in these transactions.[1] Despite the ubiquity of advisors in M&A transactions, serious concerns have been raised about whether the advisors are motivated to provide value-increasing advice. At the root of the concerns is the structure of advisors' contracts, and the most impugned feature of the contracts being *success fee compensation*: the nearly universal practice of making bidder advisor compensation contingent only on whether the acquisition bid succeeds (Edmans, 2024).

Because success fee compensation gives bidder advisors an incentive to pitch deals even if they believe that the deals will destroy value, some finance researchers believe that the use of this fee structure is "rather odd" (Edmans, 2024). Moreover, offsetting factors such as the long-run reputation costs associated with bidder advisors pitching value-destroying mergers appears to be negligible (Bao and Edmans, 2011). However, despite concerns about the structure of advisor contracts, empirical studies (e.g., McLaughlin, 1992; Hunter and Jagtiani, 2003; Calomiris and Hitscherich, 2007; Cain and Denis, 2013) have been unable to confirm the hypothesis that advisory fee structures actually generate significant inefficiencies.

The critique of success fee compensation is based on the principal-agent perspective: optimal compensation for bidder advisors should align the incentives of the principal, the bidding firm, with the incentives of the agent, the bidder advisor. Implicitly, this framing assumes that bidder advisor compensation does not affect the behavior of targets or target advisors. A somewhat smaller literature on target advisor compensation (e.g., Cain and Denis, 2013; McLaughlin, 1992) considers its effects on targets while holding constant the behavior of bidders and bidder advisors.

We step out of the principal-agent framing by modeling M&A transactions as a game with four players, two principals, the bidder and target firms, and two agents, the bidder advisor and the target advisor. Outcomes are determined by the dialectic between bidder and target advisors' attempts to sway the bidder and target firms. In this setting, one principal's change to its advisor's compensation can spill over and change the actions of the second principal and its advisor. We show that, frequently, the spillover has a greater effect on the welfare of a principal than the change in its own advisor's actions. For this reason, "odd" and misaligned advisor compensation, if it fosters the advisor dialectic, can produce more informative advisor

[1] Advising activity is, in turn, an important revenue source for bulge bracket banks, For example, from 2021 to 2023, advisory fees accounted for 9% of total revenue of Goodman Sachs. Boutique banks like Lazar and Evermore are even more dependent on advisory fees revenue.

reports and larger takeover gains than incentive-aligning compensation.

In our model, a firm (the *bidder*) attempts to acquire another firm (the *target*). In some states of nature, the acquisition destroys value and in others it increases value. Absent information production, the expected value-add of the acquisition is negative. So, in order for the acquisition to occur or succeed, credible information must be produced. The bidder and target engage advisory firms (called the *advisors*) and fix their compensation contracts. The bidder's advisor produces information related to the target's value if acquired. Consistent with boards' legal obligations, the target advisor gathers information relating to the value of alternatives to acquisition by the bidder, e.g., merger with another firm or corporate restructuring.[2]

In our baseline setting, the bidder advisor makes a publicly observable report about its information. The target advisor then makes a publicly observable report about its information. After observing both advisors' reports, the bidder makes an offer and the target decides whether to accept the offer, accept the alternative, or reject both the offer and the alternative, i.e., maintain the status-quo.

Mergers, and the valuations/opinions that advisors supply to support them, are frequently scrutinized by courts.[3] Hence, advisors are constrained by the fact that their "talk is not cheap" and the information they produce must stand up to scrutiny. Therefore, we model advisor information production using the Bayesian persuasion framework (e.g., Kamenica and Gentzkow, 2011; Dworczak and Martini, 2019; Kamenica, 2019): Advisor reports consist of distributions over signals, where each signal represents an estimate of the expected value of the takeover (bidder advisor) or alternative (target advisor) value. Signal distributions satisfy the expectancy condition, so the unconditional expectation of each advisor's signal distribution equals the ex ante expected value of the advisor's option.

Each advisor strategically chooses the informativeness of its report.[4] A fully informative report reveals all the advisor's information. A totally uninformative report simply reports the unconditional expectation of value. A partially informative report might simply reveal whether or not value exceeds some threshold.

Using this framework we consider the question of whether the advisory contracts typically employed by firms lead to significant inefficiencies. So, consistent with practice, bidder advisor compensation always takes the form of a *success fee* paid if and only if the acquisition attempt succeeds. In the baseline setting, but not in some extensions, target advisors are paid if and only if the takeover attempt fails. Hence, neither fee structure is aligned with the interests of either advisor's principal. Despite the obvious principal-agent

[2]See Pederson (2024) for a discussion of target firms' legal obligations.

[3]Over 2009–20018, between 71% and 90% of public company M&A deals were litigated (Research, 2018).

[4]Report A is more informative than report B if the signal distribution produced by A is a mean-preserving spread of the signal distribution produced by B, i.e., A dominates B in the convex order (Chapter 3.A: Shaked and Shanthikumar, 2007).

conflicts engendered by these fee structures, we find that the dialectic between the rival advisors can lead to M&A outcomes that approach first-best outcomes.

To understand this result, first consider takeover outcomes when there is no dialectic, i.e., when the target *does not* have an advisor. The bidder advisor wants to maximize the probability that the takeover succeeds. In order for the takeover to succeed, the bidder advisor must send a *persuasive signal*, i.e., a signal that convinces the target and bidder that acquisition value at least equals the target's status-quo value, w_o.

The expectancy condition implies that (a) increasing the value signaled by the persuasive signal reduces the probability of sending a persuasive signal and that (b) the bidder advisor cannot always send a persuasive signal (because ex ante expected takeover value is less than the status-quo value). Thus, a bidder advisor aiming to maximize the probability of receiving the success fee, will never send a persuasive signal greater than the status-quo value.

When the bidder advisor sends the *status-quo signal*, i.e., the signal that equals the status-quo value, a takeover, if attempted, will succeed. When the bidder advisor sends any signal $s_B < w_o$, the attempt fails. The target receives its status-quo value after both signals, and the bidders profit is zero. Thus, the M&A transaction creates no value for the bidder or target and the bidder has no incentive to launch a takeover attempt. This outcome, which we will call the "pitch the deal" outcome, matches the results of success-fee bidder advisor compensation predicted by its critics.

Now consider the effect of introducing a target advisor. In our baseline setting, the bidder advisor knows that if it sends a status-quo signal the target advisor will try to "top" it, i.e., the target advisor will try to send, with positive probability, a signal that the alternative's value exceeds the status-quo value. Will this change the bidder advisor's signaling strategy? Yes, provided that the target advisor is capable, i.e., either the odds that the alternative is better than the status quo are at least even *or* the variability of the alternative's payoff is low relative to its mean value.

The threat of being topped by a capable target advisor leads the bidder advisor to signal acquisition values higher than the status-quo value. The dialectic between advisors leads the target advisor to send a fully revealing signal, and leads the bidder advisor, when sending a persuasive signal, to signal that acquisition value equals the alternative value. Thus, the target's payoff in the Bayesian persuasion game equals its first-best payoff when information is common knowledge. As the target advisor's capability increases, the bidder's payoff also approaches its common-knowledge upper bound. Hence, not only does the target firm benefit from hiring a capable target advisor, the bidding firm benefits as well because a capable target advisor forces the bidder advisor to provide more informative signals that yield better acquisition decisions. As one might expect, if the target advisor is sufficiently incapable, the bidder advisor ignores the target advisor's feeble attempts to persuade and pitches the deal.

Alternative target-advisor compensation schemes that, at first glance, appear to better align interests of target advisors with targets do not dominate the baseline target advisor compensation. Conditioning the target advisor's fee on an alternative being accepted does increase the bidder's offer by an amount equal to the target advisor's fee. However conditioning fee payments on acceptance of the alternative motivates target advisors to just "pitch the alternative", i.e., signal an alternative value equal to the status-quo, when the bidder fails to persuade. Hence, because advisory fees, in reality, are almost always a small fraction of deal value, improved advice should dominates the relatively small increase in the offer price.

Making target advisor compensation more sensitive to the value produced by the acquisition attempt by, for example, making target compensation depend entirely on the increase in target value engendered by the acquisition attempt, also does not dominate baseline target advisor compensation. The problem with target-value sensitive compensation is that it makes target advisor payoffs insensitive to a bidder advisor's signaling strategy. The signal insensitivity can lead the bidder advisor to switch to the pitch-the-deal strategy, thereby making the acquisition attempt untenable and eliminating target and bidder acquisition gains.

To assess the effects of our baseline assumption that advisors' signals are publicly disclosed we consider a private-signaling setting: each principal only observes the signal produced by its advisor.[5] In the private signaling setting, the bidder advisor must produce more informative persuasive signals to induce the bidder to set a price high enough to ensure that the target advisor cannot propose a viable alternative to the acquisition. If the bidder advisor opts to send these high signals, the private information setting leads to more information production than the baseline public-signaling setting. However, in many cases, bidder advisors who would have used informative signaling strategies in the baseline setting opt for status-quo signaling in the private signaling setting, which results in less information production and lower target and bidder payoffs. Thus, informational efficiency, and the payoffs to the target and bidder in the private signaling setting, are neither uniformly greater or smaller than their values in the baseline public-signaling setting. The equilibrium configurations are the same under public and private signaling but the parameter regions that support these configurations vary significantly.

[5]We also show that the private signaling setting is outcome equivalent to a setting where signals are public but the bidder advisor makes offers before observing the target advisor's signal. So, the analysis of this section also addresses the effects of changes in the sequence of moves on the informational efficiency of the M&A market.

Next, we consider the effect of our assumption that the bidder advisor reports first by considering a setting where advisors simultaneously submit reports. We show that, in many cases, the baseline and simultaneous signaling settings produce identical equilibrium advisor behavior. However, in some cases, profoundly different equilibrium configurations emerge. The advisors' equilibrium signal distributions are quite complex and resemble equilibrium distributions in all-pay auction models more than typical signal distributions in Bayesian persuasion games. Although, advisors' strategies sometimes very significantly from strategies in the baseline model, the payoffs to the bidder and target and the efficiency of information production in the simultaneous setting fairly closely track the results of the baseline model.

Finally, our baseline model assumes that bidders make first-and-final offers to targets and thus assigns all bargaining power to bidders. To examine the effects of this assumption, we consider mechanisms that permit the target to capture gains in excess of its reservation demands. We find that our results about advisor behavior are not affected by increasing the bargaining power of the target.

Related literature

As discussed in the introduction, our research is motivated by the extensive body of empirical research on the role of advisors in M&A transactions.[6] Our model is built on the Bayesian persuasion paradigm developed by Kamenica and Gentzkow (2011), which has been applied in many branches of economics (cf. Kamenica, 2019). In finance contexts Bayesian persuasion has been used to model disclosure by regulators (Goldstein and Leitner, 2018; Pavan and Inostroza, 2021), over-the-counter market benchmarks (Duffie, Dworczak, and Zhu, 2017), and financing policy (Szydlowski, 2021).

Our application of the Bayesian persuasion paradigm is novel in two respects. First, it is the only application of Bayesian persuasion to M&A transactions, a setting in which persuasion plays an important role. Second, it embeds Bayesian persuasion in a rather novel market environment: multiple signal senders (the two advisors) and multiple signal receivers (the bidder and target). This contrasts with the basic Bayesian persuasion model, which has one (signal) sender and one receiver. Extensions of the basic model typically allow for either multiple senders and one receiver (e.g., Gentzkow and Kamenica, 2017b,a), or multiple receivers and one sender (e.g., Bergemann and Morris, 2019).[7]

Our analysis is also related to Gentzkow and Kamenica (2017b), which shows that, if the information

[6]In addition to the works sited in the introduction, see Kale, Kini, and Ryan (2003); Golubov, Petmezas, and Travlos (2012); Rau (2000).

[7]Gul and Pesendorfer (2012) model multiple senders and receivers in a completely different economic setting: competing political parties influencing voter coordination.

environment is Blackwell-connected, competition between senders is a necessary and sufficient condition for full revelation. In our model, advisors produce statistically independent signals and thus our information environment is not Blackwell-connected. Nevertheless, the dialectic between advisors greatly stimulates information revelation. [8]

2 Baseline Model

There are two firms, a *bidder* and a *target*. The bidder seeks to acquire the target. We have nothing new to add to the rich and well-developed literature on the firms choosing whether to initiate acquisitions. So to keep the paper focused, we start at the point where the bidder has decided to attempt the acquisition. In Section 2.5, we identify equilibrium configurations of our model in which, under standard assumptions in the takeover literature, initiating the acquisition is not tenable.

The status-quo value of the target is $w_o > 0$. If the target is acquired by the bidder, its value is $\tilde{w}_B$, where $\tilde{w}_B$ is a random variable that can take one of two possible values, 0, or $\bar{w}_B > w_o$. The target's value under an alternative to acquisition by the bidder, $\tilde{w}_T$, is also random variable that can take one of two values, 0, or $\bar{w}_T > w_o$. The alternative to the acquisition may be the acquisition of the target by a different firm, the acquisition by the target of another firm, a joint venture involving the target, or a restructured target firm. The value under acquisition, $\tilde{w}_B$, is independent of value under the alternative, $\tilde{w}_T$. We denote the expectations of $\tilde{w}_B$ and $\tilde{w}_T$ by w_B and w_T respectively.

To help it evaluate the target, the bidder retains an advisor, the *bidder advisor*. The target observes this action and also hires an advisor, the *target advisor*. The advisors and all other agents are risk neutral and maximize their expected payoffs. They are patient and do not discount future payoffs.

Each advisor undertakes an investigation that results in a publicly observable signal. The bidder advisor's signal s_B contains information about $\tilde{w}_B$, the post-acquisition value of the target. The target advisor's signal, s_T, contains information about $\tilde{w}_T$, the target's value under an alternative to the acquisition. The two firms act after the advisor's signals are revealed. In Section 6, we demonstrate the robustness of our results in alternative information settings.

The agents' actions are presented below in the sequence with which they act.

[8] Dewatripont and Tirole (1999) show that competition among advisors, who they refer to as advocates, stimulates information production. The tension in their model is that information production depends on advisors' efforts. They do not model strategic information disclosure by advisors, which is our focus. Our setting is more congruent with the concerns raised by empirical literature on advisor compensation. This literature never expresses concerns about the whether M&A advisors receive compensation sufficient to motivate effective information production but does express significant doubts about whether advisor compensation contracts lead advisors to disclose sufficient information to increase the value of M&A transactions, as opposed to just disclosing enough information to ensure that deals are consummated.

1. *Bidder Advisor*: The bidder advisor conducts an investigation into the target's post-acquisition value that results in the signal s_B about $\tilde{w}_B$. The signal realization is publicly observed.

2. *Target Advisor*: The target advisor conducts an investigation into the target's value under alternatives to acquisition and produces the signal s_T about $\tilde{w}_T$. The signal realization is publicly observed.

3. *Bidder*: The bidder chooses an offer price, $P \geq 0$ which, if accepted, requires the target shareholders to exchange all of their shares for a cash payment of P. The price is publicly observed.

4. *Target:* The target chooses a response, $r \in \{o, T, B\}$. Response o represents choosing the status-quo; T represents accepting the alternative put forth by the target advisor, and B represents accepting the bidder's offer.

5. *Payoffs realized*: Payoffs to the advisors, the target, and the bidder are realized.

Note that, in some situations, any offer, P, that the target might accept generates a loss for the bidder. In practice, bidders would simply refrain from making offers in such situations. In our setting, offers that will be rejected with probability 1 are payoff equivalent to refraining to make an offer. So, to reduce the need for introducing new notation, we adopt the convention that the bidder always makes an offer.[9]

2.1 Advisor compensation and incentives

The advisors are compensated for their services. The bidder pays its advisor a fee that is contingent on its offer being accepted and the target pays its advisor a fee that is contingent on the bidder's offer being rejected. Thus, the advisors' objectives are completely antithetical: the bidder advisor aims to maximize the probability that the bidder's offer succeeds and the target advisor aims to minimize this probability.

In practice, in the vast majority of acquisitions, bidder advisors receive a success fee that is paid if and only if the target firm accepts an offer from the bidder. This success-fee based structure for bidder advisors is the source of academic and practitioner critiques of advisor compensation detailed in the introduction. In all scenarios we consider, bidder advisor compensation is in the form of success fees.

The baseline fee structure maximizes the target advisor's incentive to investigate alternatives, and is a reflection of the legal obligation of target boards when entertaining cash merger offers, of "acting reasonably to seek the transaction offering the best value reasonably available to the stockholders".[10] In practice, target

[9]Note also that, because all information is public information, the medium of exchange, e.g., stock vs. cash, is not informative. Thus, one can interpret the bidder's offer of a cash payment of P as offering financial assets with a market value of P.

[10]Arnold v. Society for Savings Bancorp, Inc. 650 A.2d 1270, 1289-90 (Del. 1994) referenced in Pederson (2024).

advisor compensation varies and they can be paid fees that are contingent on the outcome of a bidder offer, whether the target chooses an alternative, and the price of an offer it accepts. In Section 5, we consider alternatives to the baseline target advisor fee structure and demonstrate the robustness of our baseline results.

Payments that are not contingent on the acquisition outcome have no effect on the advisors' incentives, so such payments are irrelevant in our setting. Moreover, only differences between outcome-contingent payoffs determine advisors' incentives. Hence, we normalize fixed payments to the advisors to zero and, without loss of generality, we assume that minimum outcome-contingent payment equals zero.

To simplify the exposition of the baseline model, we assume that the fee is the same for both bidder and target advisors. The fee is represented by $f > 0$. Heterogeneous fees will affect the equilibrium outcomes but these effects are proportional to the size of fees relative to firm value. Thus, under the reasonable assumption that fees are small relative to firm value, fee effects are second order and have no qualitative effects. We defer further discussion of fees to Section 5.

2.2 The advisors' signals

We adopt the Bayesian persuasion paradigm (e.g., Kamenica and Gentzkow, 2011; Dworczak and Martini, 2019) to model the advisors' information production. Specifically, the advisors, who are ex ante uninformed, launch investigations. To investigate $\tilde{w}_j$, $j = B, T$, the advisor chooses a data generation function that, for each possible realization of $\tilde{w}_j$ assigns a probability distribution over some set of reports. The realized report is then verifiably conveyed to other agents. Every report realization generates a posterior probability distribution for $\tilde{w}_j$, which agents use to form a conditional expectation.

For example, suppose that the value of the target under bidder control is either 0 and $\bar{w}_B$ with equal probability and the that advisor chooses the following reporting strategy: when $\tilde{w}_B = \bar{w}_B$, report ρ_1 with probability 1. When $\tilde{w}_B = 0$ report ρ_0 with probability $1/2$ and report ρ_1 with probability $1/2$. After observing report ρ_0, Bayes rule implies that agents will know that $\tilde{w}_B = 0$. After observing report ρ_1, Bayes rule implies that

$$\mathbb{P}[\tilde{w}_B = \bar{w}_B | \tilde{\rho} = \rho_1] = \frac{\mathbb{P}[\tilde{\rho} = \rho_1 | \tilde{w}_B = \bar{w}_B]\,\mathbb{P}[\tilde{w}_B = \bar{w}_B]}{\mathbb{P}[\tilde{\rho} = \rho_1 | \tilde{w}_B = \bar{w}_B]\,\mathbb{P}[\tilde{w}_B = \bar{w}_B] + \mathbb{P}[\tilde{\rho} = \rho_1 | \tilde{w}_B = 0]\,\mathbb{P}[\tilde{w}_B = 0]} = \frac{2}{3}.$$

So, after observing ρ_1, agents will believe that $\tilde{w}_B = \bar{w}_B$ with probability $2/3$ and 0 with probability $1/3$.

Thus, agents' posterior distribution of conditional expectations is

$$\mathbb{E}[\tilde{w}_B|\tilde{\rho}] \stackrel{\text{dist}}{=} \begin{cases} \frac{2}{3}\bar{w}_B & \text{with probability } \frac{3}{4}, \\ 0 & \text{with probability } \frac{1}{4}. \end{cases}$$

This example provides two important insights. First, it shows that, because the advisors choose the investigation they will conduct before collecting information and they send verifiable reports of the investigations' results, the advisors' choices of data generation functions and their subsequent reports convey no private information. This stands in stark contrast to cheap talk models where agents are ex ante informed and send unverifiable messages (e.g., Crawford and Sobel, 1982).

Second, the example shows that the specific form the advisors' reports take is, *per se*, irrelevant. For instance, ρ_0 and ρ_1 in the example could well represent two distinct imaginary numbers. All that matters is the posterior distribution of $\tilde{w}_B$ conditioned on the report. In fact, as we shall see, the only relevant statistics produced by the investigations are the distributions of conditional expectations $\mathbb{E}[\tilde{w}_j|\tilde{\rho}_j]$, $j = B, T$.

We use $\tilde{s}_j$ to denote these conditional expectations. We denote the realized value of $\tilde{s}_j$ by s_j, $j = B, T$ and refer to it as the *signal* or the *signal sent*. We use F_j to represent the distribution of $\tilde{s}_j$, and call F_j the *signal distribution* and $\tilde{s}_j$ the *signaling strategy*. An *advisor strategy* at a given history of the game is the signal distribution it chooses at that history. We refer to $s_j < w_o$, $j = B, T$ as an *unpersuasive* signal since it cannot result in the advisors' favored outcome being realized, and refer to a signal greater than or equal to the status-quo value, w_o, as a *persuasive* signal. Note that a persuasive signal makes the advisors' favored outcome possible, but does not guarantee it will be realized.

2.3 Feasible advisor signals

The feasible set of signaling distributions for advisor $j \in \{B, T\}$ equals the set of all distributions that can be produced from $\tilde{w}_j$ by a data generating process. Although, at first glance, identifying the set of feasible signaling distributions seems challenging, Blackwell (1953) and Kamenica and Gentzkow (2011) provide a necessary and sufficient condition for identifying the feasible set: The set of feasible signal distributions given random variable $\tilde{x}$ equals the set of all distributions that are mean-preserving contractions of the distribution of $\tilde{x}$.[11] Because we assume that $\tilde{w}_j$, $j = B, T$, places probability weight on just two points 0 and

[11]Mean-preserving contractions of the distribution of a random variable are sometimes called "garblings" of the distribution in the economics literature (Kleiner, Moldovanu, and Strack, 2021), and "fusions" of the distribution in the mathematics literature (Elton and Hill, 1992). There are many equivalent definitions of mean preserving contractions. One such definition is as follows: F_X is a mean preserving contraction of F_Y if whenever $\tilde{X} \stackrel{d}{\sim} F_X$ and $\tilde{Y} \stackrel{d}{\sim} F_Y$, $\mathbb{E}[g(\tilde{X})] \leq \mathbb{E}[g(\tilde{Y})]$, for all convex functions g, i.e., F_X is dominated by F_Y in the convex order (Definition 3.A.1 Shaked and Shanthikumar, 2007).

$\bar{w}_j$, the sets of feasible signal distributions for the advisors is defined as follows:[12]

Definition 1. The *feasible signal distribution set* for advisor $j = B, T$ equals the set of all distributions, F, satisfying the

(a) *support condition:* The support of F_j is a subset of $[0, \bar{w}_j]$, and the

(b) *expectancy condition:* $\int_0^\infty s\, dF_j(s) = w_j$.

It's obvious that, for any mean-preserving contraction of a random variable, its support is going to be a subset of the original distribution, and thus satisfy condition (a) of Definition 1. Expectancy condition (b) is also always satisfied by a mean-preserving contraction. Thus, the restrictions on the signal distribution imposed by our value distributions, both of which have two-point supports, are the weakest possible restrictions. The rich set of possible signals these weak restriction permit can generate all of the equilibrium configurations that could be produced by other value distributions.

As is common in Bayesian persuasion models, in our baseline model and in most of model extensions, optimal signal distributions have two-point supports. For this reason, the following definition economizes on notation and simplifies our discussions:

Definition 2. We refer to all signal distributions for advisor $j = B, T$ that are supported by two points, a and b, and are feasible for j as *simple signal distributions* and use $F_j^{a,b}$ to represent such distributions, i.e.,

$$F_j^{a,b}(s) := \left(\frac{b - w_j}{b - a}\right) \mathbf{1}_{s \geq a}(s) + \left(\frac{w_j - a}{b - a}\right) \mathbf{1}_{s \geq b}(s), \quad 0 \leq a < w_j < b \leq \bar{w}_j,\ j \in \{B, T\},$$

where $\mathbf{1}$ is an indicator function.[13]

2.4 Assumptions and equilibrium

The importance of advisors depends on the ex ante expected value of a control change. If the target's highest possible value under bidder control, $\bar{w}_B$, were less than its status-quo value, w_o, then the bidder advisor could never persuade the target to accept a bidder offer. Similarly, if $\bar{w}_T$ were less than the status-quo value, the target advisor could never persuade the target to choose the alternative. In our baseline setting, we impose the following parameter restrictions that maximize the importance of advisors:

Assumption 1.

$$0 < w_T < w_B < w_o < \bar{w}_T < \bar{w}_B, \quad f < w_o < \bar{w}_T - f.$$

[12] For a derivation, see, for example Proposition 3.13 in Elton and Hill (1992).

[13] An indicator function, $\mathbf{1}_K(s)$ for a set K equals 1 if $s \in K$ and 0 otherwise.

The restrictions $\bar{w}_T < \bar{w}_B$ and $w_T < w_B$ reflect the notion that the bidder advisor's opportunities first-order stochastically dominate those of the target advisor, as is likely to be the case when target management wants to retain control because this objective will restrict the set of acceptable alternatives the target advisor can consider. Reversing this ordering does not affect the qualitative implications of our analysis. Since $w_j < w_o$ the expectancy constraint on feasible signals (Definition 1) implies that some signals sent by the two advisors must be less than w_o, and thus unpersuasive. The restriction $w_o < \bar{w}_T - f$ ensures that advisor fees do not preclude the possibility of successful persuasion by either advisor in some extensions of the model. Finally, $f < w_o$ ensures that the target firm constrained by limited liability always has enough resources to pay the advisor fee.

Our assumptions result in a sequential game of perfect information (with Nature being one of the players). So we employ the standard equilibrium concept in such settings: subgame perfect Nash equilibrium (SPNE) described by the following set of equilibrium strategies, (F_B^*, F_T^*, P^*, r^*). In the few situations where applying a refinement is required to identify the outcome of the game, we refine the equilibrium set using Pareto dominance criterion, frequently termed in the game-theory literature as the *payoff dominance criterion*.[14] An SPNE is payoff dominated, if there exists another SPNE that yields a weakly higher payoff to all agents (the advisors, the bidder, and the target) and a strictly higher payoff to at least one of the agents. Payoff dominance is sometimes called the Pareto criterion.

2.5 *Initiation tenability*

As the subsequent analysis shows, the bidder's payoff is never negative in any SPNE. However, under some parameter specifications of our baseline model and its extensions, the bidder's payoff from initiating an acquisition attempt in the SPNE of the acquisition game can be zero. As discussed at the start of this section, we do not model the decision to initiate an acquisition of a target and thus hire an advisor. However, much of the literature on acquisitions, including Grossman and Hart's (1980) classic paper, has focused on market imperfections that deter the initiation of acquisitions. In such papers, fixing a target for acquisition is assumed, quite reasonably, to entail a positive costs unrelated to advisor compensation. If such costs are indeed positive, then, when the anticipated SPNE of the acquisition game following initiation yields no

[14]Payoff dominance is the most plausible equilibrium refinement in our setting. For games of complete information like ours, risk dominance and symmetry are the most commonly advanced rival refinements. Since we model a perfect information game, i.e., the information sets of the agent at nodes at which they make decisions, are singletons, there is no strategic uncertainty and thus risk dominance is irrelevant (See Carlsson and Van Damme, 1993, for a discussion of risk dominance). Symmetry is irrelevant because our game has no symmetries.

bidder gains, bidders will not initiate. So, as in Fishman (1989, cf. footnote 3), we assume that bidders do not initiate zero payoff acquisitions. In such cases, we will say that initiating the acquisition is *untenable*. If anticipated bidder gains are positive, we will say that initiating the acquisition is *tenable*.

3 Analysis of the baseline model

To benchmark the effects of endogenous signal generation by advisors on bidder and target welfare, we first consider the two limiting exogenous information environments: no advisors to generate information and advisors who generate fully informative signals. Welfare without advisors to generate information is easy to analyze since Assumption 1 ensures that the unconditional expected value of the firm under both bidder control and under the alternative is less than the status-quo value. Thus, the status-quo will prevail, the payoff to the bidder is zero, and the payoff to the target is w_o.

At the other extreme, both advisors' signaling strategies are fixed at the full-information strategy, i.e., $\tilde{s}_j = \tilde{w}_j,\ j = B, T$, and the expectancy condition ensures that the signal $\bar{w}_j,\ j \in \{B, T\}$, is produced with probability $w_j/\bar{w}_j$. The model reduces to a simple first-and-final offer model of acquisitions: The acquisition is viable if and only if $\tilde{w}_B = \bar{w}_B$. In any equilibrium, the bidder offers the target a price, P, that equals the target's full information reservation value, which is $\bar{w}_T - f$ if $\tilde{w}_T = \bar{w}_T$ and $w_o - f$ if $\tilde{w}_T = 0$.

The bidder's payoff equals the expected value of the firm under bidder control, $\bar{w}_B$, net of the bidder advisor's fee, f, and the target's reservation value. The target advisor's fee lowers the target's reservation value, and thus the offer price, by f, and the bidder advisor's fee imposes a cost of f on the bidder. So, for the bidder, advisor fees net out and the bidder's expected payoff equals

$$\text{Bidder}^{\text{FI}} = w_B - \frac{w_B}{\bar{w}_B}\left(\frac{w_T}{\bar{w}_T}\,\bar{w}_T + \left(1 - \frac{w_T}{\bar{w}_T}\right) w_o\right). \tag{1}$$

The target's payoff equals its expected reservation value, i.e.,

$$\text{Target}^{\text{FI}} = \overbrace{\frac{w_T}{\bar{w}_T}\,\bar{w}_T + \left(1 - \frac{w_T}{\bar{w}_T}\right) w_o}^{\text{Target's Gross Payoff - FI}} - f. \tag{2}$$

The advisory fee lowers the target's payoff by f regardless of the outcome. The target pays the fee directly when the offer is rejected and indirectly, when the offer is accepted, through the reduction in reservation value engendered by the fee. The total payoff to the bidder and the target is given by

$$\text{Total}^{\text{FI}} := \text{Target}^{\text{FI}} + \text{Bidder}^{\text{FI}} = w_B + \left(1 - \frac{w_B}{\bar{w}_B}\right)\left(\frac{w_T}{\bar{w}_T}\,\bar{w}_T + \left(1 - \frac{w_T}{\bar{w}_T}\right) w_o\right) - f.$$

Note that the target's full information gross payoff can be expressed as $\mathbb{E}[\max[\tilde{w}_T, w_o]]$. Because all signaling strategies are mean-preserving contractions of the distribution of $\tilde{w}_T$ and the map $s_T \mapsto \max[s_T, w_o]$ is convex, $\mathbb{E}[\max[\tilde{s}_T, w_o]] \leq \mathbb{E}[\max[\tilde{w}_T, w_o]]$. Because $\tilde{w}_T$ is distributed $F_T^{0,\bar{w}_T}$, the target's payoff is maximized when the target advisor adopts the *fully-revealing signaling strategy* of either sending the unpersuasive signal, $s_T = 0$, or the persuasive signal, $s_T = \bar{w}_T$. [15]

3.1 Preliminary results

Now consider endogenous information production. We have grafted an M&A bidding problem into a Bayesian persuasion model. As in many other persuasion settings, simple signaling distributions are optimal in our setting. Actions spaces are continua and there arise discontinuity points in agents' payoffs functions that depend on the actions of the other agents. Since an SPNE requires non-empty best reply sets for all agents in all subgames, the need to ensure non-empty best reply sets, as in bidding models, constrains the set of SPNEs. Our first three results, Lemmas 1–3, verify these properties of SPNEs in our setting.

Lemma 1 characterizes bidder and target behavior in Stages 3 and 4 of the model (stages detailed in Section 2), conditioned on the signals sent by the advisors in Stages 1 and 2.

Lemma 1. *In any SPNE* (F_B^*, F_T^*, P^*, r^*),

(a) if $s_B > \max[s_T, w_o]$, *then* $r^* = B$, *i.e., the target accepts the bidder's offer,*

(b) if $s_B < \max[s_T, w_o]$, *the target rejects the bidder's offer, i.e.,* $r^* \neq B$,

(c) if $r^* = B$, *then the bidder offer is* $P^* = \max[s_T, w_o] - f$.

The target's reservation value conditioned on the advisors' signals is $\max[s_T, w_o] - f$. Parts (a) and (b) of Lemma 1 follow because, when $s_B - \max[s_T, w_o] > (<)\, 0$, offering a price the target will accept produces a larger (smaller) payoff to the bidder than offering an unacceptable price, $P < \max[s_T, w_o] - f$. Thus, the acquisition will occur (not occur) when $s_B - \max[s_T, w_o] > (<)\, 0$. Part (c) simply states that when the target is acquired, the bidder offers the target its reservation value.

Now consider Stage 2. The target advisor has observed the bidder advisor's signal and thus knows the firm's value after the acquisition. If this value is higher than the status-quo, to block the acquisition and earn its fee, the target advisor must signal an even higher alternative value. Otherwise, it receives nothing. Lemma 2 describes the target advisor's equilibrium behavior and its implications.

[15] Because $\max[\tilde{s}_T, w_o] = w_o + \max[\tilde{s}_T - w_o, 0]$, the target's gross payoff under signaling strategy $\tilde{s}_T$ can be interpreted as the status-quo value of the firm, w_o, plus s call option on the alternative. Less than fully informative signaling strategies reduce the call option's volatility and thereby reduce the target's payoff.

Lemma 2. *Suppose that (F_B^*, F_T^*, P^*, r^*) is an SPNE. Then in any subgame with history s_B, when the target advisor replies to s_B with s_T*

(a) *If $s_B < w_o$, $w_o \le s_T \Longrightarrow r^* = T$ and $w_o > s_T \Longrightarrow r^* = o$.*

(b) *If $w_o \le s_B < \bar{w}_T$, $s_B \le s_T \Longrightarrow r^* = T$ and $s_B > s_T \Longrightarrow r^* = B$.*

(c) *Whenever, $s_B \in [w_o, \bar{w}_T)$, the target advisor's persuasive signal matches the bidder advisor's persuasive signal, and the equilibrium signaling distribution of the target advisor thus satisfies $F_T^* = F_T^{0,s_B}$.*

Part (a) of Lemma 2 simply states that, when bidder advisor produces an unpersuasive signal, i.e., $s_B < w_o$, the acquisition will not occur regardless of the signal the target advisor produces. The lemma's other two parts describe the outcome conditional on a persuasive bidder advisor signal $s_B \in [w_o, \bar{w}_T)$. In this case, the target advisor's choice is the solution to a capacity allocation problem: The expectancy condition, which ensures that the unconditional expectation under a signal distribution equals w_T, acts as the capacity constraint. Capacity is allocated across signals by the signal distribution. The capacity used by a given signal, s_T, equals the probability of sending the signal times the signal. The target advisor's objective is to allocate capacity to minimize the probability that the acquisition occurs.

The target advisor can send a "topping signal", $s_T > s_B$, with positive probability to block the acquisition and receive the fee, f. Among all topping signals, ever smaller signals use up ever less capacity while achieving the same fee. Thus, as s_T approaches s_B from above, the target advisor's expected payoff increases. In any SPNE, the target advisor must have a best response, so the supremum must be attained by a feasible strategy. This "closure-from-above" condition implies that, when $s_T = s_B < \bar{w}_T$, the acquisition is blocked in an SPNE. Any signal $s_T < s_B$ produces the same outcome as sending $s_T = 0$: the acquisition occurs and the target advisor is not paid a fee. Since any signal $s_T \in (0, s_B)$ uses up more capacity than the signal $s_T = 0$ and also yields no fee payment, sending a signal $s_T \in (0, s_B)$ is never optimal. Hence, as stated in Lemma 2, the target advisor responds to $s_B \in (w_o, \bar{w}_T)$ with the simple signal distribution F_T^{0,s_B}.

In other words, the target advisor tries to match a persuasive signal less than $\bar{w}_T$ from the bidder advisor, and succeeds in blocking the acquisition when the advisor is able to send a persuasive signal. Matching by the target advisor is robust to several model extensions that we consider below, as is the use of the simple signaling strategy $F_j^{0,\bar{w}_T}$ that delivers a persuasive signal equal to the target's "upside", its maximum value without being acquired. Hence, for easier exposition, we define the following terms:

Definition 3. In any subgame with history s_B, we will refer to the target advisor's choice of the simple signal distribution F_T^{0,s_B} in response to s_B as the *matching strategy*. In any history, we will refer to the simple signal distribution $F_j^{0,\bar{w}_T}$, as the *target-upside strategy*. We will refer to simple bidder signal distributions F_B^{0,s_B}, $s_B \in (w_o, \bar{w}_T)$ as *toppable* strategies.

Now consider Stage 1 where the bidder advisor chooses a signaling strategy. Lemma 3 describes the bidder advisor's optimal strategy and equilibrium payoff.

Lemma 3. *In any SPNE (F_B^*, F_T^*, P^*, r^*), (i) The bidder advisor's equilibrium signaling distribution is a simple signaling distribution $F_B^* = F_B^{0,s_B}$, for some $s_B \in [w_o, \bar{w}_T]$. (ii) The payoff to the bidder if the bidder sends persuasive signal s_B given that the target, bidder and target advisor follow the equilibrium strategies specified in Lemmas 1 and 2, which we denote by $u_{BA}^*(s_B)$, is given by*

$$u_{BA}^*(s_B) := f \times \begin{cases} 0 & s_B < w_o, \\ 1 - w_T/s_B & s_B \in [w_o, \bar{w}_T), \\ 1 & s_B \geq \bar{w}_T. \end{cases} \tag{3}$$

and the equilibrium payoff to the bidder equals $\max\{\frac{w_B}{s_B} u_{BA}(s_B) : s_B \in [w_o, \bar{w}_B]\}$.

Because of the expectancy condition, like the target advisor, the bidder advisor also faces a capacity allocation problem though it must also contend with the threat of being matched by the target advisor and the acquisition being blocked: A higher persuasive signal reduces or may even eliminate the probability of being matched. However, a higher persuasive signal uses more capacity. There is always a persuasive signal that achieves an optimal balance between these two opposing forces. Hence, as stated in Part (i) of Lemma 3, the bidder advisor adopts a simple signaling strategy F_B^{0,s_B} for some $s_B \in [w_o, \bar{w}_T]$.

Part (ii) shows that $u_{\mathrm{BA}}^*(\bar{w}_T) = f$, which implies that the acquisition succeeds with probability 1 when the bidder advisor sends signal $\bar{w}_T$. This follows from the fact that the bidder advisor can guarantee itself the fee f with a signal $s_B > \bar{w}_T$. A closure-from-above condition ensures that the bidder advisor must be able to attain the same payoff if $s_B = \bar{w}_T$ in an SPNE. It also ensures that the target firm accepts the bidder's offer when $s_B = \bar{w}_T$ regardless of the signal sent by the target advisor. Hence, the acquisition occurs even when both bidder and target advisors signal $\bar{w}_T$.

In summary Lemmas 1–3 show that, in any SPNE, the bidder advisor adopts a simple signal distribution that is supported by the unpersuasive signal, 0 and some persuasive signal, $s_B \in [w_o, \bar{w}_T]$. When the bidder advisor sends the unpersuasive signal 0, the acquisition attempt fails, the target advisor adopts a simple signal distribution, and the alternative is adopted if and only the signal sent by the target advisor at least equals status-quo value, w_o. When the bidder advisor sends any persuasive signal $s_B < \bar{w}_T$, the target advisor plays the matching strategy and the acquisition occurs only when the target advisor fails to produce the matching signal. When the bidder advisor sends the persuasive signal $\bar{w}_T$, the acquisition occurs with probability 1.

3.2 Equilibrium information production

The bidder advisor must choose between two types of signaling strategies: a toppable strategy and the target-upside strategy. In the following proposition, we describe the bidder advisor's optimal choice and its implications when model parameters satisfy the hypothesis

$$4\,w_T > \bar{w}_T. \tag{Hyp-P1}$$

In essence, this hypothesis limits the coefficient of variation for the target's value under the alternative, $\tilde{w}_T$.[16] We defer further interpretation of this condition to the next section, Section 4, where we examine all equilibrium configurations, including those where hypothesis Hyp-P1 is not satisfied.

Proposition 1. *When hypothesis Hyp-P1 is satisfied, in any subgame perfect Nash equilibrium (SPNE)*

(a) The bidder advisor's equilibrium signaling strategy is the target-upside strategy.

(b) The acquisition occurs if and only if the bidder advisor signals $s_B = \bar{w}_T$; otherwise, the bidder advisor sends signal $s_B = 0$, and the acquisition attempt fails.

(c) In any payoff dominant SPNE, the sum of bidder and target payoffs equals

$$\begin{aligned} Total^{BP} &:= w_B + \left(1 - \frac{w_B}{\bar{w}_T}\right)\left(\frac{w_T}{\bar{w}_T}\,\bar{w}_T + \left(1 - \frac{w_T}{\bar{w}_T}\right) w_o\right) - f \\ &= Total^{FI} - \left(\frac{w_B}{\bar{w}_T} - \frac{w_B}{\bar{w}_B}\right)\left(\frac{w_T}{\bar{w}_T}\,\bar{w}_T + \left(1 - \frac{w_T}{\bar{w}_T}\right) w_o\right). \end{aligned}$$

(d) Given the bidder advisor's equilibrium strategy, the fully revealing target advisor strategy, is always a best reply for the target advisor.

(e) Initiation of the acquisition is always tenable.

We refer to equilibria in which the bidder advisor adopts the target-upside strategy as *target-upside equilibria*. Proposition 1 shows that all equilibria are target-upside equilibria when hypothesis Hyp-P1 is

[16] A simple calculation shows that $\bar{w}_T/w_T = 1 + (\mathrm{CV}[\tilde{w}_T])^2$, where CV represent the coefficient of variation, i.e., the standard deviation divided by the mean.

satisfied. In these equilibria, the acquisition occurs along the equilibrium path if and only if the bidder advisor's signal $s_B = \bar{w}_T$. The resulting total payoff (bidder + target) equals the signal sent by the bidder advisor less the advisory fee, i.e., $\bar{w}_T - f$. The price paid is $\mathbb{E}[\max[\tilde{s}_T, w_o]] - f$, the target's reservation.

When the bidder advisor sends the signal $s_B = \bar{w}_T$, since the acquisition is always completed, the target advisor does not receive a fee. Hence, all feasible signal distributions are best responses for the target advisor. Since the total payoff from the acquisition is unaffected by target advisor's signal, any target advisor signal distribution is supported by a payoff dominant SPNE. However, increasing the informativeness of the target advisor's signal distribution increases the target's reservation, so the most informative signaling strategy from the target advisor, $\tilde{s}_T = \tilde{w}_T$, maximizes the target's gains. This is very much in line with solutions to mechanism design problems, which are pervasive and only require that agents to have a weak preference for actions being implemented.

When the bidder advisor sends the signal $s_B = 0$, the acquisition does not occur. In this event, the bidder's and bidder advisor's payoffs equals zero, and the target advisor is paid f regardless the signal it sends. The most informative target-advisor signaling strategy, $\tilde{s}_T = \tilde{w}_T$, maximizes the payoff to the target and, along the equilibrium path, the signal sent by the target advisor does not affect the payoffs of the target advisor, bidder, or bidder advisor. Thus, in any payoff dominant SPNE, the target advisor sends the most informative signal, $s_T = \bar{w}_T$.

The efficiency loss in a target-upside equilibrium relative to the full-information environment is proportional to the gap between the maximum value under bidder control and the alternative, $w_B/\bar{w}_T - w_B/\bar{w}_B$, times the target's gross payoff under full information. As the gap shrinks to zero, i.e., $\bar{w}_T \to \bar{w}_B$, total value approaches value under full information. Thus, even though the advisors' incentives are profoundly misaligned with the interests of the target and the bidder, the dialectic between the advisors forces information revelation and results in outcomes that approach those under full information.

4 Advisor capability and other equilibrium configurations.

In this section, we extend the analysis of our baseline model to include situations in which hypothesis Hyp-P1 is not satisfied. This yields two new equilibrium configurations: *status-quo equilibria* and *intermediate equilibria*. In status-quo equilibria, the bidder and target advisors send the persuasive signals $s_B^* = s_T^* = w_o$; in intermediate equilibria, the advisors send the persuasive signals $s_B^* = s_T^* = 2\,w_T \in (w_o, \bar{w}_T)$. The following proposition describes the complete set of equilibrium outcomes for our baseline model:

Proposition 2. *In any subgame perfect Nash equilibrium (SPNE),* *(a) the bidder advisor employs the simple signaling strategy F_B^{0,s_B^*} of sending either $s_B = 0$ or $s_B = s_B^*$, where*

$$s_B^* = \begin{cases} w_o & \textit{if} \quad \dfrac{w_T}{w_o} < \dfrac{1}{2} \textit{ and } \dfrac{w_T}{\bar{w}_T} < \dfrac{w_T}{w_o}\left(1 - \dfrac{w_T}{w_o}\right) & \text{status quo,} \\ 2w_T & \textit{if} \quad \dfrac{w_T}{w_o} > \dfrac{1}{2} \textit{ and } \dfrac{w_T}{\bar{w}_T} < \dfrac{1}{4} & \text{intermediate,} \\ \bar{w}_T & \dfrac{w_T}{\bar{w}_T} > \dfrac{1}{4} \textit{ or } \left(\dfrac{w_T}{w_o} < \dfrac{1}{2} \textit{ and } \dfrac{w_T}{\bar{w}_T} > \dfrac{w_T}{w_o}\left(1 - \dfrac{w_T}{w_o}\right)\right) & \text{target upside.} \end{cases}$$

(b) In intermediate equilibria, where $s_B^* = 2w_T$, the sum of bidder and target payoffs is given by

$$Total^{BP} = Total^{FI} - \left(\frac{w_B}{2w_T} - \frac{w_B}{\bar{w}_B}\right)\left(\frac{w_T}{\bar{w}_T}\bar{w}_T + \left(1 - \frac{w_T}{\bar{w}_T}\right)w_o\right).$$

(c) Initiation of the acquisition is untenable if and only if the resulting equilibrium is the status-quo equilibrium, i.e., $s_B^* = w_o$.

Part (a) of Proposition 2 shows that the equilibrium configuration depends on the ratio $w_T/\bar{w}_T$, which reflects the coefficient of variation of the target's value under the alternative. The configuration also depends on the ratio w_T/w_o, which reflects the gap between the target's expected value under the alternative and its status-quo value, and is an upper bound of the probability that the target advisor sends a persuasive signal. Together, these two ratios determine two dimensions of the *target advisor's capability*, i.e., the target advisor's capacity to match persuasive signals from the bidder advisor.

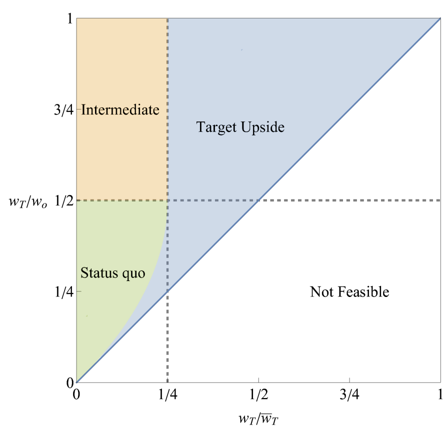

Figure 1: The figure depicts the equilibrium configuration as determined by two factors, $w_T/\bar{w}_T$ (horizontal axis) and w_T/w_o (vertical axis). Because $\bar{w}_T > w_o$, it is not possible for $w_T/w_o > w_T/\bar{w}_T$. So points to the right of the diagonal are not feasible.

Figure 1 graphs how the equilibrium configuration varies with the target advisor's capability. Points to the right of the vertical line satisfy Hyp-P1 and the outcome is a target-upside equilibrium. When Hyp-P1 is not satisfied, intermediate equilibria obtain when $w_T/w_o > 1/2$ so the target advisor is quite capable of sending a persuasive signal. Status-quo equilibria exist when this inequality is violated and the target advisor has little capacity to send a persuasive signal. Hence, status-quo equilibria are confined to cases where *both* the variation in the value of the target under the alternative is large and *and* the target advisor has less than even odds of proposing an alternative that tops the status quo. Thus, target advisor incapability can lead to the sort of "pitch the deal" outcomes pointed out in the literature, but for this to occur, the target advisor's incapability must be profound. However, even when the target advisor is quite incapable along one dimension, target-upside equilibria can exist.

To provide insight into the variation in equilibrium configurations we turn to the concave envelope approach from Aumann and Maschler (1966), which is commonly used to analyze persuasion games (e.g., Lipnowski and Ravid, 2020).[17] We only use the following well known facts: (a) the concave envelope evaluated at w_j, $j = B, T$, represents the maximum advisor payoff attainable by any random signal $\tilde{s}_j$ supported by $[0, \bar{w}_j]$ and satisfying the expectancy condition, (b) the concave envelope is majorized by all of it support lines, and (c) the concave envelope of a function majorizes that function.[18] As well as providing insights into the tensions that determine equilibrium configurations in a standard framework, we will show that, in our baseline setting and its extensions, equilibrium advisor behavior can be determined simply by identifying the persuasive bidder advisor signal that produces the largest reward-to-signal ratio.[19]

Remark 1 (Concave envelope). Let $\hat{u}_{\text{BA}}$ represent the concave envelope of the bidder advisor's payoff function, u^*_{BA} described in equation (3). Let ℓ denote a linear function. We refer to any line satisfying $\ell(0) = 0$ and $\ell(s') = u^*_{\text{BA}}(s')$ for some $s \in (0, \bar{w}_B]$ as an s'*-line*. The 0-*support line* for u^*_{BA} is an s'*-line* with the property that $\ell(s) \geq u^*_{\text{BA}}(s)$ for all $s \in [0, \bar{w}_B]$. Thus, a 0-support line is an s'-line that majorizes u^*_{BA}. If an s'-line is a 0-support line we will refer to the points $(0,0)$ and $(s', \ell(s'))$ as support points for the s'-line. Note that,

[17] The concave envelope of a function, say g, denoted by $\hat{g}$ is the *least* upper semicontinuous concave function that majorizes g, i.e., $\hat{g} \geq g$, and if h is any other upper semicontinuous concave function that majorizes g, $\hat{g} \leq h$. The concave envelope approach is also utilized in financial economics (e.g., Carpenter, 2000), operations research (e.g., Hochbaum, 2009), and, in mathematics, is a key ingredient in establishing Choquet representations of measures (Phelps, 2001) .

[18] A proof of these results, in a much more general context than ours, is provided by Proposition 11.8 in Simon (2011).

[19] Identifying optima by maximizing ratios is an intuitive and commonly used approach in finance and economics research. For example, in portfolio theory, the tangent portfolio is the line emanating from the risk-return profile of the riskless asset that supports the set of efficient portfolios. This line, called the capital market line, is in turn determined by Sharpe ratio.

if an s'-line is 0-support line then

$$\hat{u}_{\mathrm{BA}}(s') \le \ell(s') = u^*_{\mathrm{BA}}(s') \le \hat{u}_{\mathrm{BA}}(s'), \tag{4}$$

where the first inequality follows because the concave envelope is majorized by its support lines and the second inequality follows because the concave envelope of a function majorizes that function. An identical argument establishes the same inequalities for $s = 0$. Therefore, if ℓ is a 0-support line

$$\ell(s') = \hat{u}_{\mathrm{BA}}(s') = u^*_{\mathrm{BA}}(s') \quad \text{and} \quad \ell(0) = \hat{u}_{\mathrm{BA}}(0) = u^*_{\mathrm{BA}}(0). \tag{5}$$

Moreover, the expectancy condition ensures that, for *any* feasible bidder advisor signaling strategy

$$\mathbb{E}[u^*_{\mathrm{BA}}(\tilde{s})] \overset{1}{\le} \mathbb{E}[\hat{u}_{\mathrm{BA}}(\tilde{s})] \overset{2}{\le} \hat{u}_{\mathrm{BA}}(\mathbb{E}[\tilde{s}]) \overset{3}{=} \hat{u}_{\mathrm{BA}}(w_B) \overset{4}{\le} \ell(w_B). \tag{6}$$

Inequality 1 follows because $\hat{u}_{\mathrm{BA}}$ majorizes u^*_{BA}; 2 follows from Jensen's inequality, 3 follows from the expectancy condition; and 4 follows from ℓ being a 0-support line and hence a majorant for the concave upper envelope.

Next note that if the bidder advisor follows the signaling strategy, $\tilde{s}^*$ of sending the persuasive signal, $s_B = s'$ with probability $\lambda = w_B/s'$ and the unpersuasive signal $s_B = 0$ with probability $1-\lambda$, then, equation (5) and the linearity of ℓ imply that

$$\mathbb{E}[u^*_{\mathrm{BA}}(\tilde{s}^*)] = \lambda\, u^*_{\mathrm{BA}}(s') + (1-\lambda)\, u^*_{\mathrm{BA}}(0) = \lambda \ell(s) + (1-\lambda)\ell(0) = \ell(\lambda\, s') = \ell(w_B). \tag{7}$$

Hence, equations (6) and (7) imply that $\tilde{s}^*$ is an optimal bidder advisor strategy. Thus, the optimal signaling strategy can be identified by finding a 0-support line of the bidder advisor's payoff function, i.e., the persuasive signal that maximizes the reward-to-signal ratio.[20]

Each panel of Figure 2 illustrates one of the three equilibrium configurations of our model using the concave envelope approach. Because, the advisory fee, f, merely scales the bidder advisor's payoff (see equation (3)) and so does not affect the bidder advisor's preferences over signaling strategies, with a slight abuse of notation, we use u^*_{BA} with $f = 1$ to represent advisor preferences in the figure. With this normalization u^*_{BA} also represents the probability of the acquisition. Note also that we have labeled horizontal axes by s rather than s_B. We do this because, while $u^*_{\mathrm{BA}}(s)$ does represent the bidder advisor's payoff from sending

[20]If the maximum bidder advisor's reward-to-signal ratio, $u^*_{\mathrm{BA}}(s)/s$, is attained by more than one s, the parameters of the model could support more than one equilibrium configuration. However, these non-generic edge cases are ruled out in Proposition 2 by the strict inequalities used to define the regions.

signal $s = s_B$, $\hat{u}_{\mathrm{BA}}(s)$ represents the advisor's expected payoff if the expected value under bidder control equaled s and the advisor selected an optimal signaling strategy conditioned on this expected value.

In each panel, the red curves plot u^*_{BA} and the blue dashed curves plots $\hat{u}_{\mathrm{BA}}$. Each panel contains a 0-*support line* for u^*_{BA}. Some also contain other s'-lines for comparison. In each panel, the point $(w_B, \hat{u}_{\mathrm{BA}}(w_B))$, illustrated by a back dot, represents the target's expected value under bidder control, w_B (first component), and the bidder advisor's expected payoff under an optimal signaling strategy (second component). The white dots represent the support points of 0-*support line* associated with the unpersuasive and persuasive signals sent by the bidder advisor.

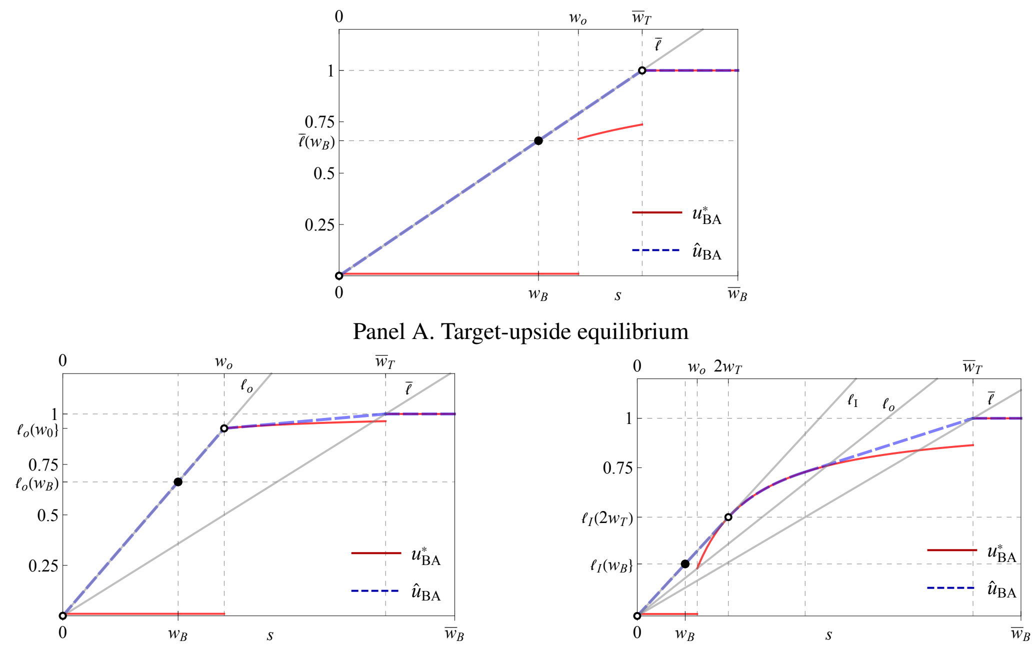

Panel A. Target-upside equilibrium

Panel B. Status-quo equilibrium

Panel C. Intermediate equilibrium

Figure 2: The target advisor, the bidder, and the target are assumed to follow their equilibrium strategies. In Panel A, $\bar{w}_T = 3.8$, $w_T = 1$, $w_B = 2.5$, $\bar{w}_B = 5$, and $w_o = 3$. In Panel B, $\bar{w}_T = 28$, $w_T = 1$, $w_B = 10$, $\bar{w}_B = 34$, and $w_o = 14$. In Panel C, $\bar{w}_T = 14$, $w_T = 1.9$, $\bar{w}_B = 16$, $w_B = 2$, and $w_o = 2.5$. The red (blue) lines depict the bidder advisor payoff function (the bidder advisor payoff function's concave envelope) given the parameters of the panel.

Consider Panel A. In this panel, $\bar{w}_T / w_T = 3.8 < 4$ so hypothesis Hyp-P1 is satisfied and the target advisor has ample capacity to match a toppable signal from the bidder advisor. Thus, the bidder advisor's payoff following a toppable signal is markedly lower than after the target-upside signal $\bar{w}_T$. Visual inception confirms that the 0-support line, labelled $\bar{\ell}$, is the $\bar{w}_T$-line with the support points $(0, \bar{\ell}(0)) = (0,0)$ and $(\bar{w}_T, \bar{\ell}(\bar{w}_T)) = (\bar{w}_T, 1)$. Its slope is $\bar{\ell}$ is $1/\bar{w}_T$. Because $\bar{\ell}$ is the 0-support line, consistent with Proposition 1, the bidder advisor sends the two signals, which are the first components (i.e., the abscissa) of the two support points of $\bar{\ell}$: the persuasive signal $s_B = \bar{w}_T$ and unpersuasive signal $s_B = 0$.

Now consider Panel B. In this panel, $\bar{w}_T / w_T = 28 > 4$ so Hyp-P1 is not satisfied and the variation of the target's value under the alternative is quite large. Moreover, $w_T/w_o = 1/14$ so the target advisor faces long odds of topping even the status quo. Thus, the bidder advisor can ensure that the acquisition occurs with a probability close to 1 by sending the status-quo signal, $s_B = w_o$. Consequently, using the target-upside signaling strategy is not optimal and $\bar{\ell}$ is not a 0-support line. Instead, the 0-support line is ℓ_o, which captures the bidder advisor's payoff from maximizing the probability of sending a persuasive signal. Hence, the bidder advisor ignores ineffective competition from the target advisor and simply aims to send a persuasive signal, i.e., the bidder advisor simply pitches the deal, resulting in the status-quo and an untenable acquisition attempt.

Finally, consider Panel C, in which $\bar{w}_T / w_T \approx 7.7$ so the expected value of the alternative and the status quo are fairly close ($w_T = 1.90$ and $w_o = 2.5$). Hence, the target advisor is quite capable of topping a status-quo signal, and the slope of ℓ_o (the w_o-line) is relatively small. However, Hyp-P1 is not satisfied and substantial variation in the target's value under the alternative ensures that the target advisor faces long odds when trying to match toppable signals well below $\bar{w}_T$. Therefore, the 0-support line is the s'-line, labelled ℓ_I, that is tangent to the graph of u^*_{BA} over $s \in (w_o, \bar{w}_T)$. Over this interval, $u^*_{\mathrm{BA}}(s)/s = (1 - \bar{w}_T/s)/s$ and this expression can be written as $(1/w_T)x(1-x)$, where $x := w_T/s$. Thus, the slope is maximized by $x' = 1/2$, i.e., $s' = 2w_T$ and the slope of ℓ_I is $1/(4w_T)$. In this case, the bidder advisor balances the reward from sending a persuasive signal, captured by ℓ_o, against the reward from sending an untoppable signal, captured by $\bar{\ell}$. This balancing results in an intermediate equilibrium.

5 Target advisor fees

In the baseline analysis, we assumed that the target advisor receives a fee if and only if the acquisition attempt fails. However, target advisor compensation varies considerably in practice in contrast to the almost uniform use of success-fee compensation for bidder advisors. We now consider the effects of alternative fee structures for the target advisor while maintaining hypothesis Hyp-P1.

5.1 Paying for the status quo?

Target advisors often receive contingent fee payments when the target maintains the status quo as we have assumed in the baseline model (e.g., McLaughlin, 1992). However, commentators have argued that boards pay target advisors such fees to obstruct takeovers that would increases the wealth of shareholders but reduce the private benefits of board members. To see if paying contingent fees for status-quo outcomes actually harms target shareholders or reduces takeover market efficiency, we drop the status-quo component of the target advisor's baseline fee and see how the resulting outcomes compare with the baseline outcome.

For this comparison, we refer to the fee in the baseline analysis—a payment of f so long as the acquisition attempt fails—as the *baseline fee*, and denote it by f_T. We refer to the fee that is paid only if the target chooses the alternative, as the *alternative fee*, and denote it by f_a. We use f_B to represent the bidder advisor's fee under both target advisor fee structures.

Ceteris paribus, dropping the status-quo fee component lowers the expected cost of hiring the target advisor. To remove this mechanical effect, we endogenize the target advisor's fee as follows: We assume that the target advisor has a reservation requirement that arises because it must exert a unit of effort which costs the advisor $c > 0$ in order to produce a signal. We describe the target's choice between the two fee structures subject to the restriction that both meet the target advisor's reservation requirement in Lemma 4.

Lemma 4. *The target shareholders will prefer the baseline target advisor fee to the alternative fee when* $\frac{f_T}{w_T} \in (0, B_a)$, $B_a > 0$.

Lemma 4 shows that, when advisor fees are small relative to firm value, as is typically the case in the real world, the baseline fee dominates the alternative fee. To understand the result first note that, under both fee structures, the target advisor must respond to a persuasive signal from the bidder advisor by at least matching the signal in order to have any hope of receiving a fee. So, conditioned on the bidder advisor sending a persuasive signal, under both fee structures the target advisor has the same incentives and follows the same matching strategy. Conditioned on an unpersuasive bidder advisor signal, the acquisition always fails regardless of the target advisor's strategy so the strategy played by the target advisor has no effect on the bidder advisor's welfare. Thus, the bidder advisor faces the same incentives and uses the same strategy under both fee structures: As shown in Proposition 1, the bidder advisor uses the target-upside strategy and the acquisition succeeds if and only if the resulting signal is the target-upside signal, $\bar{w}_T$. Moreover, the target advisor's behavior following the bidder advisor signal $\bar{w}_T$ is the same under both fee structures.

Note however that the target advisor's incentives are fundamentally different under the two fee structures in the event that the bidder advisor sends an unpersuasive signal. Under the baseline fee, the target advisor is indifferent between the status quo and the alternative and thus has no incentive to send a less than a fully informative signal. Under the alternative fee, the target advisor's unique optimal signaling strategy is the status-quo strategy of just "pitching the alternative," i.e., sending a persuasive signal that just matches the status-quo w_o, which eliminates all target gains. Hence, when the bidder advisor fails to send a persuasive signal, from an information production perspective, the alternative fee is inferior.

However, there is a benefit to the target from the alternative fee. Under the baseline fee, whenever the bidder advisor sends a persuasive signal and the target advisor sends an unpersuasive signal, the target's reservation equals $w_o - f_T$. Thus, the status-quo fee payment to the target advisor reduces the bidder's offer dollar for dollar. This bargaining-power effect is proportional to the size of the target advisor's fee and is a pure wealth transfer from target to bidder. The bargaining-power effect is absent under the alternative fee but it has no effect on the efficiency of the takeover market. In contrast, the information production effect reduces market efficiency. Hence, efficiency is always lower under the alternative fee.

Clearly there would be no positive bargaining power effect (for the target) when $f_T = 0$ and, in this case, the baseline fee dominates the alternative fee with respect to target welfare. Lemma 4 provides a range of f_T/w_T ratios between 0 and B_a over which the baseline fee is guaranteed to dominate alternative fee. In order to provide some evidence that the region of baseline fee dominance, $(0, B_a)$, contains plausible parameter choices, consider the following case. Suppose $\bar{w}_T = 1$, $w_o = 0.65$, $w_T = 0.45$, and $w_B = 0.50$. Computing B_a using the formula for B_a provided in the appendix (in the proof of the lemma) reveals that $B_a \simeq 0.233$. Which implies that as long as $f_a/w_T \leq 0.233$ or $f_T < 0.1049$, about 10.5% of $\bar{w}_T$, the expected firm value conditioned on the acquisition occurring, the baseline fee dominates the alternative fee. Thus, status-quo fee payments, while they may appear to be harmful to target shareholders, can promote better outcomes for target shareholders over reasonable levels of target advisor compensation.

5.2 *Paying for performance?*

Target advisors also often receive compensation tied to the increase in target shareholder wealth produced by an acquisition attempt. McLaughlin (1992) documents that 71.4% of the target advisor contracts include fees that are based on value improvement.[21] Cain and Denis (2013) provide further evidence that target advisors are especially incentivized to produce precise information about target firms' values. The obvious argument in favor of such performance-contingent fees is that they better align the interests of target advisors and targets.

To consider the effects of such performance-contingent fees, we add a *performance component* to the target advisor's baseline compensation. The resulting compensation, which we denote by comp_T, has two

[21] An example of such a contingent contract in McLaughlin (1992), is an agreement between Unidynamics and its advisors, Goldman Sachs and Smith Barney, to pay \$125,000 plus 2.5% of the value above \$20 per share for a completed acquisition.

parts: A *baseline component*, b_c, which is the fee $f > 0$ if the acquisition attempt fails (as in the baseline setting), and a performance component, denoted by π_c, equal to some fraction, $\alpha \in (0,1)$ of the target shareholders' payoff in excess of the status-quo value, $w_o - f$.

$$\mathrm{comp}_T = \overbrace{\alpha \max[0, \pi_T - (w_o - f)]}^{\pi_c\text{:performance component}} + \overbrace{\begin{cases} f & r \neq B \\ 0 & r = B. \end{cases}}^{b_c\text{:baseline component}} \tag{8}$$

We refer to compensation schemes like these as *performance compensation*.

We need to address one complication before proceeding: performance compensation depends on the target's payoff but the target's payoff also depends on the payment to the target advisor. To resolve this circular dependency, first consider the case where the acquisition occurs, i.e., $r = B$. In this case the baseline component of target advisor compensation is 0, i.e., $b_c = 0$, so $\mathrm{comp}_T = p_c$. The acquisition price is the target's reservation value, i.e., $P = \max[s_T, w_o] - f$. The target's payoff, π_T, equals the acquisition price less the performance compensation, i.e., $\pi_T = (\max[s_T, wo] - f) - p_c$.

If the acquisition attempt fails, i.e., $r \neq B$, the target's payoff, π_T, is the gross value of the target, $\max[s_T, w_o]$ less target advisor compensation, which includes the baseline component is $b_c = f$. So the target's payoff again satisfies $\pi_T = \max[s_T, w_o] - \mathrm{comp}_T = \max[s_T, wo] - (f + p_c)$.

By definition, $p_c = \alpha \max[\pi_T - (w_o - f), 0]$. Putting these observations together, we see that the performance component, p_c, is defined implicitly by the following functional equation:

$$p_c(s_T) = \alpha \max[s_T - w_o - p_c(s_T), 0], \quad s_T \in [0, \bar{w}_T]. \tag{9}$$

Inspection of equation (9) shows the equation's unique solution is, $p_c(s_T) = \hat{\alpha} \max[s_T - w_o, 0]$, where $\hat{\alpha} = \alpha/(1+\alpha)$ captures the sensitivity of the target advisor's payoff to target value improvement.[22] Thus, under performance compensation the target advisor aims to maximize the expectation of

$$\hat{\alpha} \max[s_T - w_o, 0] + \begin{cases} f & r \neq B \\ 0 & r = B. \end{cases}$$

What does this imply for signaling strategies and informational efficiency? First, suppose the bidder advisor sends the unpersuasive signal, $s_B = 0$. In this case, the acquisition will not occur and the target

[22]The dependency we have just resolved is similar to dependency that must be resolved when determining cash flow sensitivity, i.e., ultimate beneficial ownership, in complex business groups (Almeida et al., 2011).

advisor always receives the baseline component f. Since the baseline component is fixed and not marginal, the target and its advisor's incentives are perfectly aligned. Thus, the target advisor strictly prefers the full information strategy of sending send the target-upside signal $\bar{w}_T$ when the value under the alternative is $\bar{w}_T$ and sending the 0 signal when the value of the alternative is 0. The target advisor uses the same strategy in the baseline setting because, even though it is not the only best response, we have assumed that, when indifferent, the target advisor chooses the information disclosure strategy that maximizes the target's welfare. However, as we have shown, adding an arbitrarily small positive performance component to the target advisor's compensation can ensure that the fully informative strategy is the target advisor's unique best response when the bidder advisor sends the unpersuasive signal.

When the bidder advisor sends the target-upside persuasive signal, $s_B = \bar{w}_T$, which is the persuasive signal sent by the bidder advisor along the equilibrium path in Proposition 1, the effects of performance compensation is similar to the unpersuasive signal case we have just discussed. Under the baseline compensation, the full information signaling strategy is a best response for the target advisor, though not a unique best response it maximizes target welfare. Under performance compensation, the full information strategy is a unique best response. So, the target advisor's signaling strategy is the same.

The situation is more complex when the bidder advisor sends a persuasive toppable signal, $s_B \in [w_o, \bar{w}_T)$. The baseline component promotes the matching strategy of sending $s_T = s_B$. Because of the perfect alignment between the performance component and target welfare, the performance component promotes the full information strategy. Using the concave envelope approach, outlined in Remark 1 in Section 4, the target advisor's optimal strategy, either matching or fully informative, is the one that maximizes the reward-to-signal ratio. The difference between the reward-to-signal ratios for these two strategies is

$$\overbrace{\frac{f+\hat{\alpha}\,(s_B-w_o)}{s_B}}^{\text{matching}}-\overbrace{\frac{f+\hat{\alpha}\,(\bar{w}_T-w_o)}{\bar{w}_T}}^{\text{fully informative}}=\frac{(\bar{w}_T-s_B)\,(f-\hat{\alpha}\,w_o)}{s_B\,\bar{w}_T}.$$

Hence, when $f-\hat{\alpha}\,w_o > 0$ and the baseline fee component dominates the performance component, the target will adopt the same strategy as in the baseline scenario, namely, matching the bidder advisor's signal. Therefore, as Proposition 1 shows, the bidder advisor's equilibrium signaling strategy is the target-upside strategy and performance compensation has no effect on either bidder or target on the equilibrium path signaling strategies.

When $f-\hat{\alpha}\,w_o < 0$ and the performance component dominates, the target advisor will always use the fully informative strategy. Therefore, the bidder advisor faces a different tradeoff than in the baseline setting.

Over all toppable persuasive signals, the probability that the target advisor will send a persuasive signal is the same. Thus, conditioned on sending a toppable persuasive signal, the bidder advisor will send the lowest signal, i.e., the status-quo signal, w_o. Consequently, the only two candidates for an optimal bidder advisor signaling strategy are status-quo signaling and target-upside signaling.

The bidder advisor receives the fee, f, with probability 1 if it sends the target-upside persuasive signal, $\bar{w}_T$. After the bidder advisor sends the status-quo signal, it only receives the fee if the target advisor sends the unpersuasive signal $s_T = 0$. The payment occurs with probability $1 - w_T/\bar{w}_T$ since the target advisor employs the fully informative signaling strategy. Thus, the difference in the reward-to-signal ratios of the two candidate strategies is

$$\overbrace{\frac{f}{\bar{w}_T}}^{\text{target upside}} - \overbrace{\frac{f\left(1 - \frac{w_T}{\bar{w}_T}\right)}{w_o}}^{\text{status quo}} = f\,\frac{(w_o + w_T) - \bar{w}_T}{w_o\,\bar{w}_T}. \tag{10}$$

Consequently, if $f - \hat{\alpha}\,w_o < 0$, the target advisor adopts the fully informative strategy and the bidder advisor strictly prefers the status-quo signaling strategy to the target-upside strategy whenever $(w_o + w_T) - \bar{w}_T < 0$.

In summary, when the target advisor is not very capable, adding a performance component to target advisor compensation can have an adverse effect on the target and efficiency. When the target advisor is highly capable, holding the level of expected target advisor compensation fixed, the mix between baseline compensation and performance component is irrelevant. This result is consistent with the observed cross-sectional variation in target advisor compensation schemes. Such variation could either represent "neutral mutation" or represent targets' attempts to shape advisor behavior along dimensions other than information revelation. For example, stimulating target advisors' efforts to identify profitable alternatives to an acquisition.

6 Extensions

We now consider the effect of modifying our baseline assumptions about the public nature of advisor signals, their sequencing, and bidder bargaining power. We show the robustness of our basic thesis—even though the advisors' compensation appears to generate incentives for inefficient behavior, the dialectic between them promotes efficient information disclosure.

6.1 Private signals

We have assumed that the advisors' signals are publicly revealed before the bidder sets its offer price, and that the target advisor observes the bidder advisor's signal before generating its own signal. It is unclear whether either the bidder or the target would want to publicly reveal their advisors' signals. Therefore, we will now examine a setting in which advisor signals are private.

Suppose each firm privately observes the signal from its advisor, and the target and its advisor act after the bidder makes its offer. This is different from the baseline setting in three ways: (i) the bidder cannot condition its offer on the target advisor's signal; (ii) the target and its advisor cannot condition their responses on the bidder advisor's signal; (iii) the target advisor can condition its response on the bidder's offer.

While the target and its advisor cannot directly condition their responses on the bidder advisor's signal, the target's and its advisor's and strategies are the same in this setting and baseline setting: Accepting (rejecting) the bidder's offer is the target's best response if and only if the offer price at least equals the value of the target net of the advisor fee, conditioned on the target advisor's signal. The target advisor can earn its fee by topping the offer price, and a closure from above condition ensures that its best response is "offer price matching" i.e., when possible, the target advisor's persuasive signal matches the bidder's offer price. Hence, per se, the privacy of the bidder's signal has no relevance to the target or its advisor. The key complication introduced by moving to this setting is that the bidder makes its offer without observing the target advisor's signal. So, we term this setting the "unconditional offer setting".[23]

Despite being fundamentally different from the baseline setting, the unconditional offer setting too produces status-quo, intermediate, or target-upside equilibrium configurations. However, the parameter restrictions that support the configurations in the two settings are substantially different. In Result 1 we describe equilibrium outcomes in the unconditional offer setting when Hypothesis 1 is satisfied. We defer analysis of this setting for all remaining parameter values to the Appendix.

Result 1. Define

$$\bar{s} := \frac{(w_T + \bar{w}_T)^2}{4\,w_T} - \frac{\max\left[w_o - \frac{1}{2}\left(w_T + \bar{w}_T\right), 0\right]^2}{w_T}.$$

If the baseline hypothesis, $4\,w_T > \bar{w}_T$ is satisfied then

(a) In the unconditional offer setting the equilibrium is a status quo equilibrium if

- (i) If $\bar{w}_B \leq \bar{s}$, or
- (ii) If $\bar{w}_B > \bar{s}$, $\frac{1}{2}\left(w_T + \bar{w}_T\right) - w_o > 0$, and $\sqrt{\frac{w_T}{w_o - w_T}} - \frac{\frac{1}{2}\left(w_T + \bar{w}_T\right)}{w_o} < 0$, or
- (iii) $\bar{w}_B > \bar{s}$, $\frac{1}{2}\left(w_T + \bar{w}_T\right) - w_o \leq 0$, and $w_o^2 - (w_o - w_T)\left(w_T + \bar{w}_T\right) < 0$.

[23] That is, the private information setting is equivalent to a symmetric information bidder-commitment setting where the bidder commits to an offer before observing the target's signal—first, the bidder advisor publicly reports its signal; next, the bidder sets the offer price; the target advisor then publicly reports its signal; finally, the target makes its accept/reject decision. Thus, any divergence from our baseline results is entirely produced by the bidder making an offer that is not conditioned on the target advisor's signal, whether this is the result of an inability to observe the signal or because of committing to an offer before observing the signal.

The target and bidder advisors send the status-quo signal w_o, and the takeover attempt is untenable.

(b) Otherwise, the equilibrium is a target-upside equilibrium. The bidder advisor sends persuasive signal $s_B^* = \bar{s} > \bar{w}_T$, and the target advisor sends persuasive signal $s_T^* = \bar{w}_T$.

(c) In the baseline setting, the equilibrium is a always a target-upside equilibrium and both the bidder and target advisor send persuasive signal $\bar{w}_T$.

Result 1 shows that, unlike the baseline setting, hypothesis Hyp-P1 does not guarantee that all equilibria are target-upside equilibria in the unconditional offer setting. There also exist status-quo equilibria.

In the baseline setting, in equilibrium, the offer price matches the target's reservation, which equals the status quo value when the target advisor produces an unpersuasive signal. In this setting, because the bidder sets the offer price without observing the target advisor's signal, the bidder doesn't know whether the status quo value will be the target's reservation. Moreover, because the target advisor picks its signal after the offer price is set, by raising the offer price (so long as it below $\bar{w}_T$), the bidder reduces the probability that the target advisor can send a persuasive matching signal. Hence, the possibility arises that the bidder will make an offer that exceeds the target's reservation value so the target is paid an "information rent". This improves the target's welfare relative to the baseline setting.

The possibility of an information rent payment dramatically changes the bidder advisor's equilibrium signaling strategy. In the baseline setting, in equilibrium, the bidder and target advisors always choose identical signaling strategies. Other than in status-quo equilibria, this is not the case in equilibrium in the unconditional offer setting. To see why, consider, for example, the smallest untoppable signal, $\bar{w}_T$. Conditional on this signal, in the baseline setting, the bidder could earn positive profit if the target advisor produced an unpersuasive signal. In the unconditional offer setting, the bidder cannot condition on the target advisors signal and will not wish to offer a takeover price that ensures success, $\bar{w}_T - f$, because this bid will lead to zero profit. The bidder will offer $\bar{w}_T - f$ only if the bidder advisor's persuasive signal exceeds $\bar{w}_T$.

Thus, in the unconditional offer setting, target-upside equilibria are supported by bidder advisor signals higher than $\bar{w}_T$, which may result in higher bidder welfare relative to the baseline setting in target-upside equilibria. However, the extra capacity required to produce a signal higher than $\bar{w}_T$ encourages the bidder advisor to opt for a status-quo signal. Moreover, for a target-upside equilibrium to exist, the bidder advisor must have the capacity to produce a persuasive signal that exceeds $\bar{w}_T$ by an amount that is fixed by the bidder's best reply correspondence. Such a persuasive signal is feasible only if $\bar{w}_B$ exceeds $\bar{w}_T$ by a sufficient amount. Since the baseline model only assumes that $\bar{w}_B > \bar{w}_T$, this requirement is not necessarily satisfied.

In other words, when $\bar{w}_B$ is not sufficiently high, the target-upside configuration obtained in the baseline setting is not even feasible in the unconditional offer setting.

Figure 3 illustrates Result 1 as well as equilibrium configurations when hypothesis Hyp-P1 does not hold. A comparison with Figure 1 shows how the possibility of an information rent payment and the pressure on the bidder advisor to produce more informative signals change equilibrium configurations relative to the baseline setting. When the takeover is tenable, the bidder advisor produces more information in the unconditional offer setting than in the baseline setting. However, parameters that support intermediate and target-upside equilibria in the baseline setting support status-quo equilibria in the unconditional offer setting.

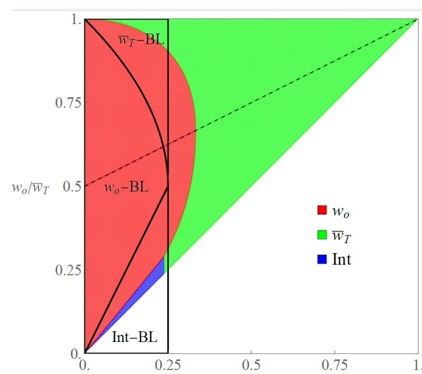

Figure 3: The figure depicts the equilibrium configuration as determined by two factors, $w_T/\bar{w}_T$ (horizontal axis) and $w_o/\bar{w}_T$ (vertical axis). Labels Int-BL, w_o-BL, and w_T-BL refer to the regions enclosed by the black lines. These are the regions in which the intermediate, status-quo equilibrium, and target-upside equilibria obtain in the baseline model, respectively. The regions shaded blue, red, and green represent the intermediate equilibria, status-quo equilibria, and target-upside equilibria in the unconditional offer setting when Hyp-P1 is not satisfied. The dotted line represents the boundary between the $\mathscr{A}(w_T, \bar{w}_T) < w_o$ (above line) and $\mathscr{A}(w_T, \bar{w}_T) > w_o$ (below line) regions.

6.2 Simultaneous advisor signaling

Thus far, we have assumed that the target advisor observes the signal produced by the bidder advisor before selecting its signaling strategy. It is quite possible that a target advisor does not have access to a bidder advisor's signal before choosing its own signaling strategy. So we will now examine a setting, termed the "simultaneous setting," In the simultaneous setting, as in the baseline setting, the bidder and target make their decisions after observing the bidder and target advisor' signals. However, in contrast to the baseline setting, the bidder and target advisors do not observe each other's signals before sending their signal.

Because, bidder and target best replies depend only on the signals' sent by the advisors, Lemma 1 and Part a of Lemma 2 still characterize target and bidder equilibrium behavior. Thus, "equilibria" in this section of the paper refers to equilibrium advisor behavior conditioned on the target and bidder following their equilibrium strategies in the SPNE as defined by these lemmas. As in the baseline analysis, when both advisors submit the same signal, say s, ties are broken in favor of the target advisor when $s \in [w_o, \bar{w}_T)$ and in favor of the bidder advisor when $s = \bar{w}_T$.

6.2.1 *Equilibria in simple strategies*

In the simultaneous setting, the scope for equilibria in simple signal distributions is quite limited. To see this, suppose that, in some equilibrium, both advisors send the same persuasive signal $s_B = s_T = s \in [w_o, \bar{w}_T)$. In this case, the acquisition attempt succeeds only if the target advisor sends the unpersuasive signal, $s_T = 0$. Suppose the bidder defects to sending persuasive signal $s_B = s + \varepsilon$, $\varepsilon \simeq 0$ and $\varepsilon > 0$. The cost of defection, which results from the expectancy condition, is that the probability sending the larger persuasive signal, $s + \varepsilon$, is lees than the probability of sending s. However, this cost can be made arbitrarily small by making ε arbitrarily close to 0. The benefit of the defection strategy is that, for all $\varepsilon > 0$, the takeover succeeds whenever the bidder advisor sends $s_B = s + \varepsilon$, Because the gain from defection is appreciable and its cost can be made arbitrarily small, $s_B = s$ is not a best reply for the bidder advisor. Hence, in any equilibrium in simple strategies, persuasive bidder advisor signals satisfy the following condition, which we term the *no-ties condition*:

$$\text{If } s \in [w_o, \bar{w}_T] \text{ and } \mathbb{P}[\tilde{s}_B = \tilde{s}_T = s] > 0 \Longrightarrow s = \bar{w}_T. \tag{11}$$

At the same time, if the persuasive signals are such that $s_B \neq s_T$, one of the advisors is not choosing a best reply. To see this, Suppose, without loss of generality, that $s_B > s_T$. Then $w_o \leq s_T < s_B$. In this case, in the candidate equilibrium, the takeover succeeds whenever the bidder advisor sends a persuasive signal s_B. Consider a defection by the bidder advisor to $s_B - \varepsilon$, where $0 < \varepsilon < s_B - s_T$. Under the defection strategy, the takeover also succeeds whenever the bidder advisor sends a persuasive signal but because $s_B - \varepsilon < s_B$, the persuasive signal used in the defection strategy can be sent with larger probability. Hence, the defection strategy yields a larger bidder payoff. Consequently, in any equilibrium in simple strategies, both advisors send the same persuasive signal.

The only unpersuasive signal sent by either advisor is 0 and both advisors send this signal with positive probability Thus, the supports of both advisors' signal distributions contain only one signal, $s = 0$. Hence, all equilibria satisfy the *overlapping supports condition*: the supports of the bidder and target advisors' signal distributions are identical.

The overlapping supports condition and the no-ties condition (equation 11) imply that, in all simple strategy equilibria, both the bidder and target advisor send persuasive signal $s = \bar{w}_T$ and unpersuasive signal $s = 0$, i.e., all equilibria in simple strategies are target-upside equilibria.

6.2.2 Target-upside equilibria

When do target-upside equilibria exist? First, note that, when the bidder advisor adopts the target-upside signaling strategy, the acquisition attempt succeeds if and only if the bidder advisor sends the persuasive signal $s_B = \bar{w}_T$. The target advisor's signaling strategy has no effect on the outcome. Thus, any target upside signaling strategy is best reply for the target advisor. If the bidder advisor defects from a target-upside equilibrium to some persuasive signal $s_B \in [w_o, \bar{w}_T)$, the acquisition will occur if and only if $\tilde{s}_T = 0$, which occurs with probability $1 - w_T/\bar{w}_T$. The reward-to-signal ratio for defection is thus

$$f \times \frac{\left(1 - \frac{w_T}{\bar{w}_T}\right)}{s_B} \text{ and } \max_{s_B \in [w_o, \bar{w}_T)} f \times \frac{\left(1 - \frac{w_T}{\bar{w}_T}\right)}{s_B} = f \times \max_{s_B \in [w_o, \bar{w}_T)} \frac{\left(1 - \frac{w_T}{\bar{w}_T}\right)}{s_B} = f \times \frac{\left(1 - \frac{w_T}{\bar{w}_T}\right)}{w_o}.$$

The reward-to-signal ratio for the candidate equilibrium is $f \times 1/\bar{w}_T$. Following the equilibrium strategy is a best response for the bidder advisor if and only if the equilibrium produces at least as high reward-to-signal ratio as defection. So the bidder advisor's target-upside signaling strategy is a best response if and only if

$$\overbrace{\frac{1}{\bar{w}_T}}^{\text{equilibrium}} - \overbrace{\frac{\left(1 - \frac{w_T}{\bar{w}_T}\right)}{w_o}}^{\text{defection}} = \frac{(w_o + w_T) - \bar{w}_T}{w_o \bar{w}_T} \geq 0. \tag{12}$$

These observations are collected in the following result.

Result 2. (a) All equilibria in simple strategies are target-upside equilibria. (b) Target-upside equilibria exist if and only if $\bar{w}_T \leq w_T + w_o$,

Result 2 shows that, in the simultaneous setting, the condition for target-upside equilibria in simple strategies is $\bar{w}_T \leq w_T + w_o$, The baseline sufficient condition for target-upside equilibrium, Hyp-P1, is $\bar{w}_T \leq 4 w_T$. Neither condition implies the other but both conditions are qualitatively similar. Both conditions impose an upper bound on the target-upside signal, $\bar{w}_T$, based on a linear function of the status-quo payoff, w_o, and/or the expected value of the alternative, w_T.

6.2.3 What happens when $\bar{w}_T > w_T + w_o$

As shown by Result 2, if $\bar{w}_T > w_T + w_o$, equilibria in simple strategies do not exist. Thus, as we verify in the appendix, at least one advisor's equilibrium signal distribution is not discrete and the cardinality of that advisor's support is uncountably infinite. These sorts of signal distributions are profoundly different than the signaling strategies considered anywhere else in our analysis. For this reason, we defer detailed analysis of this case to the appendix.

In the appendix, we find that (a) there exist complex target advisor signal distributions that support equilibria where bidder advisor employs the simple target-upside signal distribution, $F_B^{0,\bar{w}_T}$ whenever $\bar{w}_T \leq w_T + \sqrt{w_o^2 + w_T^2}$. (b) When $\bar{w}_T > w_T + \sqrt{w_o^2 + w_T^2}$, in equilibrium, both advisors employ complex signaling strategies; the common upper bound of the supports of the advisor's signal distributions is less than $\bar{w}_T$. Equilibrium signal distribution are fairly complex and determined by conditions that are quite similar to the conditions that identify equilibria in all-pay auction models (Barut and Kovenock, 1998; Siegel, 2009; Fang, Noe, and Strack, 2020).

A rough summary of these results is that, when Hyp-P1 is not satisfied, the simultaneous signaling setting sometimes produces more informationally efficient outcomes than the baseline setting. However, when Hyp-P1 is satisfied. the simultaneous signaling setting never produces more informationally efficient outcomes and sometimes produces less informationally efficient outcomes than the baseline setting. For a large region of the parameter space, the equilibrium outcome is the same in both settings, namely target-upside signaling. Qualitatively, the tradeoffs determining the informational efficiency of advisor signaling strategies are quite similar in the two settings. In both settings, the dialectic between bidder and target advisors frequently result in significant information production even when the compensation contracts offered bidder and target advisors diverges significantly from the contracts that would be optimal in a principle-agent setting.

6.3 Bidder bargaining power

In the baseline setting, the bidder makes a take-it-or-leave-it offer to the target at the price P. This assumption ensures that the bidder captures the entire surplus from the acquisition. How robust are our conclusions to increasing the ability of the target to capture the surplus? In order to avoid overly complicating the analysis by introducing a non-trivial bargaining game played by the bidder and target after advisor signals are received, consider the case where the surplus is allocated via a cooperative bargaining solution.

Specifically, suppose that the acquisition price and the outcome of the acquisition are determined by a mediator who, after observing the advisors' signals, applies a cooperative bargaining solution, e.g., the Nash Bargaining solution. All such cooperative bargaining solutions are Pareto efficient and provide a strictly positive share of the surplus from acquisition, if any, to both the target and bidder. By varying the weights placed on target and bidder payoff, bargaining solutions can provide the target with a fraction of the gains arbitrarily close to one.

Under such solutions, the acquisition attempt will succeed when value under bidder control, s_B, exceeds value under target control, $\max[s_T, w_o]$ and fail when s_B is less than $\max[s_T, w_o]$. When $w_o \le s_B = s_T$, offer failure and offer success produce the same surplus and thus offer success and offer failure are both Pareto efficient outcomes and thus candidate outcomes. However, for the same reasons presented in the analysis of the baseline model, the only tie-breaking rule consistent with an SPNE for in strategic game played by the advisors is for ties to be broken in favor of the target advisor when $w_o \le s_B = s_T < \bar{w}_T$ and broken in favor of the bidder advisor when $s_B = s_T = \bar{w}_T$. Thus, under baseline compensation, which is based only on the outcome (acquisition success or failure), the advisors' strategic incentives are the same as the baseline setting. Hence, equilibrium advisor strategies and therefore equilibrium information disclosure in bargaining solution setting are the same as the baseline setting. Thus, the allocation of bargaining power has mo material effect on our baseline results with respect to information disclosure and thus the total gain from the acquisition.

7 Conclusion

In this paper we examined a paradox: Why do bidding firms compensate their advisors based on acquisition success rather than the economic value generated by the acquisition. We showed that if bidder compensation is viewed as one part of a broader M&A ecology that includes target firms and target advisors, the effect of bidder advisor compensation schemes on M&A outcomes depends not only on the effect bidder compensation has an bidder-advisor actions but also on the effect bidder compensation has, through its effect on the actions of bidders, on target-advisor actions. Thus, mis-alignment between bidder-advisor compensation and the welfare of the bidding firm need not lower the bidding firm's welfare. In fact, we showed that when both bidder and target advisors have compensation contracts that misalign their incentives with those of their principals, the dialectic between these competing biased advisors can reveal the value-add potential of the merger to both the bidder and the target. Moreover, from the perspective of information disclosure, biasing advisors and misaligning their incentives with their principals can be optimal.

We next showed that how aligning target advisor compensation with the target's payoff from the merger can be both unattractive to the target and reduce overall efficiency. Alignment can change target advisor optimal strategies and by so doing change bidder optimal responses in a way that reduces the aggregate informativeness of advisor reports and thus lowers target welfare. Thus, as well as resolving the bidder-advisor compensation paradox we also rationalize some features of target advisor compensation schemes that are viewed as puzzling or inconsistent with boards acting in shareholders' interests such as rewarding advisors simply for thwarting mergers or not conditioning target rewards on the bidder offer.

More generally, we think that the basic insight of this paper—the analysis of the effects of compensation provided to competing experts, e.g. lawyers in civil law suits, advisors to creditors and debtors in insolvency proceedings, depend not only on the direct effect of incentives on a principal's agent but also on the effect the package has on the actions of the rival principal's agent. Hence, compensation designs that are optimal from the perspective of simple principal-agent models may not be optimal when viewed in the larger context of expert competition on behalf of their principals.

References

Almeida, Heitor, Sang Yong Park, Marti G Subrahmanyam, and Daniel Wolfenzon. 2011. "The structure and formation of business groups: Evidence from Korean chaebols." *Journal of Financial Economics* 99 (2):447–475.

Aumann, Robert J and Michael Maschler. 1966. "Game theoretic aspects of gradual disarmament." *Report of the US Arms Control and Disarmament Agency* 80:1–55.

Bao, Jack and Alex Edmans. 2011. "Do investment banks matter for M&A returns?" *Review of Financial Studies* 24 (7):2286–2315.

Barut, Yasar and Dan Kovenock. 1998. "The Symmetric Multiple Prize All-Pay Auction with Complete Information." *European Journal of Political Economy* 14 (4):627–644.

Bergemann, Dirk and Stephen Morris. 2019. "Information design: A unified perspective." *Journal of Economic Literature* 57 (1):44–95.

Blackwell, David. 1953. "Equivalent comparisons of experiments." *The Annals of Mathematical Statistics* :265–272.

Cain, Matthew D. and David J. Denis. 2013. "Information Production by Investment Banks: Evidence from Fairness Opinions." *Journal of Law and Economics* 56:245–280.

Calomiris, Charles W and Donna M Hitscherich. 2007. "Banker fees and acquisition premia for targets in cash tender offers: Challenges to the popular wisdom on banker conflicts." *Journal of Empirical Legal Studies* 4 (4):909–938.

Carlsson, Hans and Eric Van Damme. 1993. "Global games and equilibrium selection." *Econometrica* :989–1018.

Carpenter, Jennifer N. 2000. "Does option compensation increase managerial risk appetite?" *Journal of Finance* 55 (5):2311–2331.

Crawford, Vincent P and Joel Sobel. 1982. "Strategic information transmission." *Econometrica* :1431–1451.

Dewatripont, Mathias and Jean Tirole. 1999. "Advocates." *Journal of Political Economy* 107 (1):1–39.

Duffie, Darrell, Piotr Dworczak, and Haoxiang Zhu. 2017. "Benchmarks in search markets." *Journal of Finance* 72 (5):1983–2044.

Dworczak, Piotr and Giorgio Martini. 2019. "The simple economics of optimal persuasion." *Journal of Political Economy* 127 (5):1993–2048.

Edmans, Alex. 2024. "Conflicts of interest among M&A Advisors." https://alexedmans.com/blog/corporate-finance/conflicts-of-interest-among-ma-advisors.

Elton, J and Theodore P Hill. 1992. "Fusions of a probability distribution." *Annals of Probability* :421–454.

Fang, Dawei, Thomas Noe, and Philipp Strack. 2020. "Turning up the heat: The discouraging effect of competition in contests." *Journal of Political Economy* 128 (5):1940–1975.

Fishman, Michael J. 1989. "Preemptive bidding and the role of the medium of exchange in acquisitions." *Journal of Finance* 44 (1):41–57.

Gentzkow, Matthew and Emir Kamenica. 2017a. "Bayesian persuasion with multiple senders and rich signal spaces." *Games and Economic Behavior* 104:411–429.

———. 2017b. "Competition in persuasion." *Review of Economic Studies* 84 (1):300–322.

Goldstein, Itay and Yaron Leitner. 2018. "Stress tests and information disclosure." *Journal of Economic Theory* 177:34–69.

Golubov, Andrey, Dimitris Petmezas, and Nickolaos G Travlos. 2012. "When it pays to pay your investment banker: New evidence on the role of financial advisors in M&As." *Journal of Finance* 67 (1):271–311.

Grossman, Sanford J and Oliver D Hart. 1980. "Takeover bids, the free-rider problem, and the theory of the corporation." *Bell Journal of Economics* :42–64.

Gul, Faruk and Wolfgang Pesendorfer. 2012. "The war of information." *Review of Economic Studies* 79 (2):707–734.

Hochbaum, Dorit S. 2009. "Dynamic evolution of economically preferred facilities." *European Journal of Operational Research* 193 (3):649–659.

Hunter, William C and Julapa Jagtiani. 2003. "An analysis of advisor choice, fees, and effort in mergers and acquisitions." *Review of Financial Economics* 12 (1):65–81.

Kale, Jayant R, Omesh Kini, and Harley E Ryan. 2003. "Financial advisors and shareholder wealth gains in corporate takeovers." *Journal of Financial and Quantitative Analysis* 38 (3):475–501.

Kamenica, Emir. 2019. "Bayesian Persuasion and Information Design." *Annual Review of Economics* 11:249–272.

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101 (6):2590–2615.

Kleiner, Andreas, Benny Moldovanu, and Philipp Strack. 2021. "Extreme points and majorization: Economic applications." *Econometrica* 89 (4):1557–1593.

Lipnowski, Elliot and Doron Ravid. 2020. "Cheap talk with transparent motives." *Econometrica* 88 (4):1631–1660.

McLaughlin, Robyn M. 1992. "Does the form of compensation matter?" *Journal of Financial Economics* 32:223–260.

Pavan, Alessandro and Nicolas Inostroza. 2021. "Persuasion in global games with application to stress testing." .

Pederson, Kent A. 2024. "Highway to the Revlon zone: Merging Unocal and Revlon into a unified enhanced scrutiny framework." *Corporate and Business Law Journal* 5:62–83.

Phelps, Robert R. 2001. *Lectures on Choquet's theorem.* Springer.

Rau, P Raghavendra. 2000. "Investment bank market share, contingent fee payments, and the performance of acquiring firms." *Journal of Financial Economics* 56 (2):293–324.

Research, Cornerstone. 2018. "Shareholder Litigation Involving Acquisitions of Public Companies." Tech. rep., Cornerstone Research.

Shaked, Moshe and J George Shanthikumar. 2007. *Stochastic Orders*. Springer Science & Business Media.

Siegel, Ron. 2009. "All-Pay Contests." *Econometrica* 77 (1):71–92.

Simon, Barry. 2011. *Convexity: An analytic viewpoint*, *Cambridge Tracts in Mathematics*, vol. 187. Cambridge University Press.

Szydlowski, Martin. 2021. "Optimal financing and disclosure." *Management Science* 67 (1):436–454.