

Who Really Knows What?

Information asymmetry and its structural consequences on M&A deals.

Tommaso Sarli

July 26, 2026

Contents

1	Introduction	1
2	Anatomy of the Gap	2
3	The Toolkit	3
3.1	Due Diligence	3
3.2	Payment Structure	4
3.3	Earnouts	4
3.4	Representations and Warranties	4
3.5	Material Adverse Change Clauses	5
3.6	Fairness Opinions	5
3.7	The Cost of the Toolkit	6
4	Do These Mechanisms Work?	6

1. Introduction

By 2025, roughly two-thirds of US private-target middle market deals referenced representations and warranties insurance, an instrument that exists because nobody has found a way to make the buyer know what the seller knows. [2] An entire insurance market, with its own underwriters, models and claims data, has grown up around a gap that half a century of deal practice has never managed to eliminate.

And it never will. The gap exists because the seller, the buyer and their advisors each sit in a different place relative to the same information, and no amount of deal-making can put them in the same place. It is built into the transaction, and it shapes how every deal gets priced and closed.

A deal is composed of different centres of interest, namely the seller, the buyer and the advisor, each defined by different objectives and levels of knowledge. The seller knows more about itself than any kind of due diligence will ever reveal, from stakeholder relations to

tacit debts. The buyer knows itself, its stock value, synergies and needs, while the advisor has access to profound market intelligence. These parties generate the asymmetry not individually but within the relationships they form with each other. Seller-Buyer, Client-Advisor and Management-Shareholder are the main relations affected by this force.

The force also does not operate uniformly through time. Before the deal is concluded it concerns what each party can discover. The buyer investigates through due diligence, the seller discloses only what it is compelled to disclose, and the advisor reads and gives an opinion on the market. Once the deal is signed, the concern moves to what each party chooses to do once nothing further can be found out.

Throughout the years, the market evolved to find solutions to the problem. An extensive toolkit comprised of contracts, additional clauses and payment structure changes was established. The goal of this article is to analyse each of the major tools available and understand how, and to what degree, it attacks information asymmetry.

2. Anatomy of the Gap

Before we examine the toolkit itself, a more profound overview of the main relations affected by information asymmetry should be given.

The simplest way through which the asymmetric distribution of knowledge impacts M&A deals is the relationship between the Seller and the Buyer. The former knows more about its company than anyone else. This gap does not lie in financial statements or synergies, both easily discoverable and evaluated through due diligence. It lies in board sentiment, management behaviour, stakeholder relationships and undisclosed liabilities.

Undiscoverable knowledge allows information asymmetry to impact the deal in several ways. Firstly, through adverse selection. The gap causes a valuation discount, which pushes good targets out of the equation or requires a much more demanding deal structure. This is a practical application of Akerlof's "The Market for Lemons" (1970). [1] The direction of this effect depends on how the deal is structured. In a negotiated, low-competition sale, the discount dominates, since a single buyer can price in its uncertainty. In a competitive auction, the dynamic reverses. The winning bidder suffers the so-called "winner's curse", having been the party that most overestimated the target. The same asymmetry produces opposite outcomes depending on the number of bidders at the table, which already suggests that these gaps are governed by deal structure rather than removed by it.

Second, further information asymmetry appears on the buyer's side. Especially in stock-financed deals, the seller has no concrete way to accurately determine whether the stock is overpriced.

The gap between Seller and Buyer is only one part of a much bigger picture. A separate relation afflicted by information asymmetry is the one connecting a client and its advisor. The asymmetry here stems from the advisor, usually an investment bank, having greater access to and familiarity with market intelligence than the client. Comparable-deal terms, depth of buyer interest, and knowledge of its own effort, such as how thoroughly due diligence has been conducted, all sit on the advisor's side of the table.

Although this asymmetry is inevitable, whether it is ultimately exploited depends on a second force: what the advisor is incentivized to do with the information it holds. As advisors are mainly compensated if the deal is successfully concluded, these institutions can present numbers in ways that favour a faster, more certain conclusion rather than a favourable client outcome. The client thus finds itself in an unfavourable position compared to its advisor. It cannot verify much of the information or the statements provided, given its lack of market experience, resources and means to audit the advisor's own efforts.

A third relation afflicted by the asymmetric distribution of knowledge is the one connecting a company's management to its shareholders. Management's position inside the company gives

it a persistent informational advantage over shareholders. Stakeholder relations, undisclosed future plans and accounting figures are all known to management well in advance, and are ultimately reported by management itself.

An additional factor which compounds the asymmetry is synergy projections, forecasts with no mandatory methodology, entirely constructed by management. These affect the deal prior to its conclusion but also persist beyond that point, as management reports on synergy integration according to its own timeline and metrics. It knows if the firm is underperforming before the market does.

Whether this informational advantage translates into shareholder harm depends on the structure of management incentives themselves. Gordon (2025) documents how compensation arrangements tied to liquidity events, such as golden parachutes, can invert the incentive to bargain. [4] When a large share of an executive's payoff crystallises on closing rather than on price, the marginal gain from pushing for better terms falls, and the cost of a failed deal rises. The informational advantage becomes harmful once it sits alongside a payoff structure indifferent to the outcome it produces.

Seller and buyer, advisor and client, management and shareholders: three distinct relationships, each shaped by the same structural force. In each case, one side has an informational advantage over the other and, depending on the incentives the superior party holds, it becomes consequential.

At this stage it is worth setting the boundaries of what follows. The vast majority of currently used transaction-specific tools, and most of the ones analysed in this article, deal with the seller-buyer relationship. The other two, management-shareholders and client-advisor, are mainly governed by governance structure, incentives, disclosure obligations and standing agreements, subjects external to the toolkit and beyond the scope of this discussion. The question that follows is not whether the gap can be eliminated, since perfect information is not a realistic standard in any negotiated transaction, but how deal practice has tried to narrow it, or at least govern its consequences.

3. The Toolkit

The instruments that follow fall into three groups. Due diligence stands alone as the only genuine discovery mechanism, and it is treated first because it is the baseline against which everything else is a response. Payment structure, earnouts, representations and warranties, RWI and MAC clauses then operate as risk allocation and pricing tools between buyer and seller. Fairness opinions work differently again. Rather than allocating risk, they produce a certification of the deal's terms, which is why they are evaluated on different grounds than the rest of the toolkit.

3.1. Due Diligence

Due diligence is the market's most direct attempt at shrinking the gap. Unlike the instruments discussed later, it does not change the allocation of risk, nor does it compel disclosure during negotiations. Its sole purpose is the discovery of things the other party already knows, such as financial statements, contracts and undisclosed liabilities.

Although necessary for a deal, this tool is structurally imperfect. It is backward looking, as it cannot verify nor predict the future, and it cannot reach tacit information such as board sentiment or stakeholder relations. Lastly, the quality of due diligence is partly determined by the party that is the object of the research itself. The seller might hide or omit certain information from both public and private records, making it simply impossible to discover.

This imperfection means information asymmetry can be narrowed but never eliminated. It is the residual left by due diligence's structural limits that the following instruments respond to, each in a different way.

3.2. Payment Structure

The first response lies in the payment structure a deal adopts, and specifically in the differences between a cash and a stock-financed deal.

In stock-financed deals, holding equity in the combined firm acts as contingent pricing. The target's shareholders absorb part of the risk associated with the deal, and thus the loss if they prove to be overvalued. Equity therefore becomes the bidder's preferred medium of exchange when uncertainty regarding the target is highest (Hansen 1987). [5] The proportion of consideration paid in equity also acts as a reverse signal. A rational bidder prefers stock when it believes its own currency is overvalued, and targets, understanding this, rationally discount such offers (Myers & Majluf 1984). [8] The signal carries weight because the bidder's choice of currency reveals its own private valuation of that currency.

In cash-based deals the buyer bears the entire "lemon" risk. Cash shows confidence in the purchase, but it might also signal that the buyer's own currency is undervalued, which benefits the target.

No payment structure entirely benefits one of the parties. Cash protects the seller from acquirer-overvaluation risk but loads target-valuation risk onto the buyer, and stock does the exact reverse. The favoured side is whoever has the weaker bargaining position going in.

3.3. Earnouts

A bidder and seller can also resolve disagreements over valuation by changing when the price is fixed, rather than how it is paid. This is the logic behind earnouts.

Through this tool, the seller shares part of the valuation risk by conditioning payment on post-closing performance metrics. This works on two distinct levels. As a signal, it reassures the buyer about the seller's confidence in and commitment to the business. Kohers and Ang (2000) and Datar et al. (2001) find that earnout-based deals tend to reflect this dynamic, with earnouts appearing most frequently where the target is hardest to value from outside. [3, 6] As a risk sharing mechanism, it means that if metrics are reached shareholders participate in the higher earnings, and if they are not, the buyer is spared from having overpaid at closing.

Reflecting these combined benefits, earnouts are usually favourably perceived by the market, with stocks experiencing a 1.3% increase on the day of announcement of an earnout-based deal, rising to 2.2% over a two-day window (Kohers and Ang 2000). [6]

Although beneficial, it is not a perfect solution. Information asymmetry switches from adverse selection to moral hazard, the risk derived from unobservable behaviour by one of the agents. This takes the form of seller-side manipulation, where short-term tactics are adopted to meet targets, potentially compromising long-term results, and buyer-side manipulation, where the new owner can suppress earnings in order to hold down the earnout payment.

As a result, the tool does not directly favour a particular party. In the abstract, it holds both upsides and downsides for both. In practice, it depends on how the parties behave, and the tool seems most beneficial to whichever party is most willing to exploit the deal terms.

3.4. Representations and Warranties

While payment structure and earnouts deal with valuation risk and post-closing behaviour, a different tool is employed to allocate the risk deriving from disclosure uncertainty: representations and warranties.

This tool works on two levels. Its primary function is legal. A breach of a representation allows the buyer to draw a claim against the escrow, thus allocating part of the cost of

undisclosed liabilities to the seller. R&W negotiations can also work as a signal. Sellers with nothing to hide have fewer reasons to decline larger escrows or broader representations, while a seller aware of undisclosed problems will push back and settle for lower escrows.

A further development of this tool transfers the risk away from the seller altogether. Representations and warranties insurance shifts the cost of a breach onto a third-party insurer rather than a counterparty in the transaction.

This shift removes the role of the escrow and only partly resolves the disclosure uncertainty the instrument was originally created for. The signalling effect survives, as the insurer, having to bear the risk, will run a thorough due diligence on the seller, pricing the risk derived from information asymmetry. On the other hand, the seller is now only responsible for deliberate concealment rather than innocent or negligent misstatement, as the insurer can generally only pursue it for fraud. The tool thus appears to be softer.

RWI is the clearest example of how information asymmetry has structurally changed M&A practice. The problem is large enough that an entire market has been built to deal with it, employing sector expertise, models and claims data. Similarly to other tools, RWI does not shrink information asymmetry but builds an efficient market around it, one where insurers specialise in pricing disclosure risk into the deal itself. Its effects are positive enough that by 2025, roughly two-thirds of US private-target middle market deals referenced RWI (American Bar Association 2025). [2]

RWI is framed as a win-win, but its adoption has tracked periods of seller leverage. The clean-exit benefit to the seller is the more durable one, while the buyer's gain is the price of remaining competitive in a seller's market.

3.5. Material Adverse Change Clauses

The tools discussed so far allocate the cost of risk that already exists at signing. An additional tool is needed to address asymmetry that develops during the executory period, when the seller keeps operating and generating information while the buyer, who is usually no longer running due diligence, has to bear the additional risk. This tool is the material adverse change (MAC) clause.

Although the clause gives the buyer a way out in case of significant deterioration during the interim period, it appears weak in practice. The burden of proof is steep, requiring a durationally significant collapse and thorough documentation. A clear example of this stems from the history of Delaware merger litigation, where only one claim has ever resulted in a successful MAC determination, in Fresenius Kabi AG's case against Akorn, Inc.

The clause thus sees major usage as a renegotiation lever. A buyer who suspects deterioration uses the threat of prolonged litigation to negotiate a lower price (Zhou 2016). [9] The information asymmetry that develops during the executory period is therefore not neutralised but simply re-priced, and only when deterioration is severe enough to make the threat credible. Below that threshold, the seller's advantage remains largely intact.

3.6. Fairness Opinions

The tools examined so far directly address the gap between buyer and seller, either through contingency, price or an exit right. Fairness opinions operate differently. Rather than allocating risk, they represent a judgment on the adequacy of the deal's terms themselves. When a board considers a sale, it typically lacks the in-house expertise to judge independently whether the price offered is adequate, so it commissions a third-party financial advisor's assessment that the transaction is fair to shareholders from a financial point of view.

The market reacts differently depending on who provides that assessment. If the deal advisor also writes the fairness opinion, it is seen as conflicted and the market discounts the price. Makhija and Narayanan (2007) find that deals with conflicted fairness opinions carry merger

premiums 7.7% lower, rising to 13% where both sides have conflicted opinions. [7] The same research shows that board structure affects the shareholder reaction: a captured board usually draws a negative one, as its members are understood to be pursuing personal benefits, whereas an independent board produces the opposite.

Fairness opinions are therefore a viable instrument only when provided by non-advisors and to an independent board. Outside those conditions, the tool benefits other parties in the deal rather than its intended beneficiary. Advisors gain twice over, since certification obtained at lower merger premiums brings more certainty around closing, and therefore payment, while also limiting legal liability by evidencing a rigorous valuation. A captured board gains as well, as shareholders are presented with an expert certification they have no means to scrutinise, and board members may exploit it while collecting contingent fees.

3.7. The Cost of the Toolkit

One dimension runs through all of these instruments and is easily overlooked. None of them is free. Escrows immobilise a portion of the seller's consideration for the duration of the survival period. RWI carries a premium and a retention, paid regardless of whether a claim ever arises. Earnouts require post-closing accounting infrastructure, and generate a meaningful share of post-deal litigation. Fairness opinions carry advisory fees. MAC-based renegotiations consume time and legal resources at precisely the moment when both parties are most exposed.

These costs are the price of transacting across the gap, and they explain adoption patterns better than risk logic alone. RWI spread through the middle market partly because it allocates risk efficiently, and partly because its premium is small relative to the delay and negotiation it removes. Earnouts persist despite their moral hazard problem because abandoning a deal over an unbridgeable valuation gap costs more. The toolkit works as a set of trades in which parties accept a known cost in exchange for a reduction in unknown risk.

4. Do These Mechanisms Work?

Not every tool resolves the problem it addresses in the same way. Some, such as R&W and RWI, are self-enforcing. They do not require trust between the parties, nor further intervention after closing. Others are recursive, introducing a new version of the problem elsewhere. Earnouts shift the problem from adverse selection to moral hazard, while fairness opinions allow certification to be abused for personal gain.

To summarise, none of the tools examined here eliminates information asymmetry outright. Only due diligence and the signalling function narrow it. Most of the toolkit reallocates or prices the residual risk rather than closing the gap. Earnouts transform adverse selection into moral hazard, RWI allocates residual risk rather than reducing it, R&W insurance prices that risk and transfers it to a third party, MAC clauses reserve a mostly theoretical exit and are used in practice as negotiating leverage, payment structure shifts which side bears the risk, and fairness opinions work only as a market signal.

The toolkit also fails to protect the intended beneficiary with any consistency, and at times protects the party it was meant to discipline. Fairness opinions are often read as conflicted by the market, causing a discount in both merger premiums and stock value.

This does not mean the tools are useless. They exist to render deals transactable despite the gap rather than to close it, and three implications follow from that.

The first concerns the direction of financial innovation. The sophistication of these instruments measures the gap's persistence rather than its retreat. Each new tool has moved further from discovery and closer to pricing. Due diligence tried to find the information. Representations and warranties allocated the cost of not finding it. RWI outsourced that cost to a party whose entire business is estimating it. The trajectory runs from knowing,

to allocating, to underwriting, and there is little reason to expect the next instrument to reverse it.

The second concerns distribution. Because these tools are priced, they are available in proportion to the ability to pay for them, and to the sophistication required to negotiate them. A seller advised by a bank that runs dozens of processes a year will extract a cleaner exit, a narrower survival period and a better-structured earnout than one advised by a firm running its first. The toolkit converts informational advantage into a negotiable commodity, and the party better able to buy and structure it captures most of the benefit. This is the mechanism behind the observation, made throughout this article, that tools frequently favour the party they were designed to discipline.

The third, and the most consequential, is that the residual risk has been concentrated rather than dispersed. When disclosure risk was borne by counterparties, it sat across thousands of individual transactions and escrows. Under RWI, it accumulates on the balance sheets of a small number of specialised underwriters, priced on models built from a claims history compiled during a benign period for private-target M&A. That concentration is efficient while the pricing holds. It is worth asking what happens to deal flow in the middle market, where two-thirds of transactions now depend on this instrument, if a claims cycle proves those models optimistic and capacity withdraws.

References

- [1] George A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- [2] American Bar Association. Private target mergers & acquisitions deal points study, 2025. Mergers & Acquisitions Committee, Business Law Section.
- [3] Srikant Datar, Richard Frankel, and Mark Wolfson. Earnouts: The effects of adverse selection and agency costs on acquisition techniques. *The Journal of Law, Economics, and Organization*, 17(1):201–238, 2001.
- [4] Jeffrey N. Gordon. Too many mergers? the golden parachute as a driver of M&A activity in the twenty-first century. *Stanford Journal of Law, Economics & Business*, 30(2), 2025.
- [5] Robert G. Hansen. A theory for the choice of exchange medium in mergers and acquisitions. *The Journal of Business*, 60(1):75–95, 1987.
- [6] Ninon Kohers and James Ang. Earnouts in mergers: Agreeing to disagree and agreeing to stay. *The Journal of Business*, 73(3):445–476, 2000.
- [7] Anil K. Makhija and Rajesh P. Narayanan. Fairness opinions in mergers and acquisitions. *Working Paper, Fisher College of Business, The Ohio State University*, 2007.
- [8] Stewart C. Myers and Nicholas S. Majluf. Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, 13(2):187–221, 1984.
- [9] Y. Carson Zhou. Material adverse effects as buyer-friendly standard. *New York University Law Review Online*, 91:171–185, 2016.