TRACK B · MISALIGNMENT TAX SERIES

The Deal Certainty Discount

Close-probability and post-LOI value compression in lower-middle-market transactions

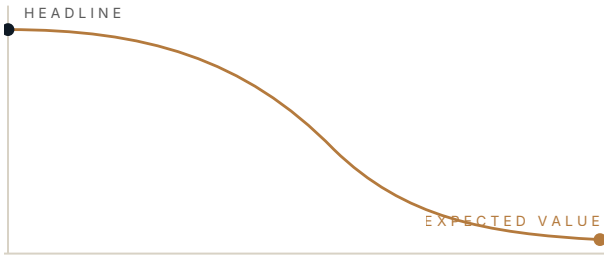

PUBLISHED JUNE 2026
Cordis Institute for Lower-Middle-Market Research
cordisinstitute.org

BY CORDIS INSTITUTE

THE FINDING

A signed letter of intent is a probability-weighted number, not a price. In a thin process the highest headline is frequently the least certain, and a seller who ranks by headline rather than by expected realized value can choose wrongly.

1 in 3	90%+	~21c	45-90d
ENGAGEMENTS · NO DEAL	PPA PREVALENCE	EARNOUT REALIZED	EXCLUSIVITY
Sell-side engagements ending without a transaction. Pepperdine PCMR 2025.	Private-target deals with working-capital adjustments	Per dollar of face, across private-target deals	No-shop window, lengthening since 2021

ABSTRACT

The argument in brief.

A sell-side process is usually decided on the highest letter of intent, on the assumption that the headline number is the number the seller will receive. It is rarely the number the seller receives. A signed letter of intent is non-binding, a material share of signed letters never reach close, and conditional on closing the price is routinely revised downward through purchase-price adjustments, escrows, and earnouts that now appear in the large majority of private-target deals. This paper defines the gap between a headline offer and its expected realized value as the Deal Certainty Discount and develops it as an organizing framework rather than a measured result. The empirical inputs are drawn from the cross-market M&A deal-terms and bidding-strategy literature, so the lower-middle-market claims are advanced as testable propositions rather than segment-specific measurements. Its central proposition is that offer aggressiveness raises both the risk of failure and the expected post-signing compression while process competition lowers both, so that the highest headline bid is frequently the least certain, and ranking offers by headline rather than by expected realized value can invert a seller's choice. The discount is the execution-stage component of the founder-to-close gap documented in prior Cordis Institute work, and the paper proposes a two-stage framework for where lower-middle-market value leaks.

KEYWORDS mergers and acquisitions, lower middle market, deal certainty, letter of intent, purchase price adjustment, deal completion

JEL G34, G32, D82, L26

KEY TAKEAWAYS

Four takeaways.

01 The letter of intent is not the deal.

A signed LOI is non-binding, and a material share of signed letters never reach close. The no-shop period it imposes, commonly 45 to 90 days, removes the seller's alternative during exactly the window in which price is most exposed to revision.

02 A headline is a probability-weighted number.

Expected realized value equals the headline times the probability of closing times one minus the expected post-signing compression. The gap between the headline and that value is the Deal Certainty Discount.

03 Aggressiveness and certainty pull apart.

Higher, more aggressive offers tend to win the bid and exclusivity, but winning the letter is not closing it. In a thin process the bid written to win is disproportionately the one that later fails or re-trades.

04 The tax is levied twice.

Buyer-lane divergence sets the number before the letter is signed (WP-002). Close-probability and post-signing revision set how much of that number survives. Two stages, two different fixes.

SECTION 01

Introduction.

A headline is a forecast, not a price.

Ask a founder which of two offers is better and the answer is almost always the larger one, which is a reasonable instinct. The headline number is the one part of a transaction that is legible without a lawyer, and after a working life spent building the business it reads as a verdict on that work. But a headline is a forecast rather than a settled fact, and what it forecasts is that the buyer who wrote it will still be there, at that price, several months later.

A letter of intent is a non-binding expression of price subject to diligence and documentation, and the period that follows its signature is the one in which the seller holds the least leverage. In that window a deal can fail outright, or it can survive at a revised and lower price, and the headline number anticipates neither outcome.

The Deal Certainty Discount is defined here as the difference between a headline offer and its expected realized value, once the probability of closing and the expected post-signing revision are taken into account. From that definition the paper advances a behavioral proposition that runs against the seller's instinct: the features that make an offer aggressive tend to lower its probability of closing and to raise the revision it later suffers, so that the highest headline is often the least certain. Two things should be clear about the claim. The expected-value arithmetic involved is elementary and is not presented as novel, and the empirical figures are cross-market, so the lower-middle-market reading is offered as a set of testable propositions, with the measurement that would confirm them identified as the next step.

The idea is not unfamiliar to practitioners, but the framing here differs. Transactional-law commentary and deal advisers discuss deal certainty mostly from the buyer's side, as a certainty premium that a seller may concede for a cleaner, faster close (Goodwin, 2025). This paper takes the seller's side of the same coin: the discount that a rational seller should apply to an uncertain headline before comparing it with a more certain one.

The argument is meant to complete a sequence of Cordis Institute work, and the relationship to that work should be stated honestly rather than leaned on. WP-001 measured the total gap between a founder's pre-marketing expectation and the realized close, the finding named the Misalignment Tax (SSRN 6515478). WP-002 located one source of that gap before the letter is signed, in how different buyer lanes underwrite the same business to different numbers (SSRN 6735844). These are prior papers by the same institute, not independent corroboration, and they are used here only to motivate the framework, not to prove it. What they leave open is the portion of the gap that forms after signing, which is the gap this paper organizes. Section 2 states the data and its limits. Sections 3 through 5 build the discount and the proposition about aggressiveness and certainty. Section 6 sets out the two-stage framework. Sections 7 and 8 turn to implications and to what the evidence cannot yet settle.

SECTION 02

Data and Methodology.

An analytical model built on public secondary data.

This is principally an analytical paper. Its central object, the Deal Certainty Discount, is a definition and a small reduced-form model, and its empirical content is secondary, drawn from published studies rather than from a proprietary transaction sample. That choice bounds the claims, and the paper tries to stay inside those bounds.

Four bodies of work supply the inputs. The M&A deal-terms studies published by SRS Acquiom document the mechanisms through which a signed price is revised after signing, principally working-capital purchase-price adjustments, escrows, and earnouts. The Pepperdine Private Capital Markets Report surveys intermediaries on why engagements end without a transaction. The academic literature on bidding

strategy and deal outcomes, including Boone et al. (2024) on the private phase of the deal process, relates the structure of an offer to whether it wins and to how it is subsequently revised. Prior Cordis Institute work, WP-001 and WP-002, is cited to locate this paper in a research program, with the caveat in Section 1 that it is not independent evidence.

Two features of these sources require care. First, they are cross-market. The SRS Acquiom samples span private-target deals well beyond the lower middle market, the Pepperdine respondents are intermediaries across size bands, and the bidding-strategy studies use broad transaction databases. The prevalence of the mechanisms is therefore well established, while their precise magnitude within the lower middle market is not, and this paper does not assign one. Second, the probability of close, which the model treats as the load-bearing parameter, is not established for the lower middle market in any published large-sample study known to the author. A figure circulates in the advisory community that roughly half of signed letters fail to close; it is treated throughout as a practitioner estimate, not a measurement, and Section 8 makes the resulting limitation explicit. No Cordis re-trade frequency or average compression is presented, because none has been measured to the standard this series requires.

SECTION 03

The Letter of Intent Is Not the Deal.

~50%
LOIS THAT FAIL

Exclusivity removes the seller's only alternative.

The legal character of a letter of intent is well understood and routinely underappreciated by first-time sellers. With the usual exception of the exclusivity and confidentiality provisions, it binds neither party to the transaction. What it does bind is the seller's optionality, through the no-shop clause that accompanies most letters. That clause, commonly 45 to 90 days, takes the seller off the market for the length of diligence, and exclusivity periods have tended to lengthen over the past several years (Goodwin, 2023). For the stretch in which price is most exposed to revision, the seller has agreed not to build an alternative.

How often the process then fails is harder to state than practitioners sometimes suggest. The often-repeated claim that about half of signed letters do not close is plausible and consistent with what intermediaries report, but it rests on consensus rather than on a published study of lower-middle-market completion, so it is used here only as an order-of-magnitude indication. The nearest survey evidence measures a different object: the Pepperdine Private Capital Markets Report (2025) finds that close to a third of sell-side engagements end without a transaction, with a buyer-seller valuation gap the most frequently cited reason and most such gaps falling in an eleven-to-thirty-percent band. Engagement failure and letter failure are not the same measurement, but they run in the same direction, and a non-trivial fraction of priced deals do not survive the path from intent to close.

Exhibit 2 Practitioner estimates put letter-to-close survival near one in two.

Source · Practitioner consensus (advisory community); corroborated in direction by Pepperdine PCMR 2025 (~one-third of engagements end without a transaction). Treated as an order-of-magnitude estimate.

SECTION 04

The Deal Certainty Discount.

Rank offers by expected value, not headline.

The quantity that ultimately matters to a seller is the cash that clears at close, not the number printed on the letter. Let H be the headline offer. Model the outcome of a signed letter as a simple branch. With probability q the transaction fails and the seller realizes nothing from it; with probability 1 - q it closes, and conditional on closing the seller realizes a fraction 1 - c of the headline, where c is the expected proportional reduction between signing and close. Expected realized value is then, in reduced form,

$$E[V] = (1 - q) \cdot H \cdot (1 - c)$$

The Deal Certainty Discount is the gap between the headline and this expected value, H - E[V]. The arithmetic is elementary, and the contribution is not the formula but what it forces into view: that an offer can fail outright, and that a surviving offer is usually not paid in full.

A caution about the model belongs here rather than in a footnote. The two adjustments, q and c, are not independent, and treating them as if they were would misstate the discount. Re-trading and outright failure are partial substitutes: a buyer who extracts a price reduction late in diligence is, in many cases, choosing to close at a lower number instead of walking away, so heavier expected compression can coincide with a higher probability of closing rather than a lower one. The reduced form above should therefore be read as an accounting identity for a given offer, not as two independent levers, and the interesting variation is in how q and c move together across offers, which is the subject of Section 5.

The components are worth grounding in what is actually known, and the evidence on the failure side has to be read with care. Research on the private bidding phase finds that higher and more aggressive initial offers tend to win the deal and to draw fewer subsequent revisions (Boone et al., 2024); that is how an aggressive bid secures exclusivity in the first place. Winning the bid is not the same as closing it. A separate failure channel operates after signing, when a deal is abandoned because its economics deteriorate during diligence, and this is the outcome q is meant to capture. The conditions that win a letter and the conditions that carry it to close are not identical: a high number paired with thin financing commitments and questions deferred into diligence can win the first while raising q on the second. On the compression side, the machinery is nearly universal even where its magnitude is not measured for this segment. Working-capital purchase-price adjustments appear in more than ninety percent of private-target transactions, up from roughly half a decade earlier (SRS Acquiom, 2025a), so an explicit channel for revising price after signing is

present in almost every deal. Earnouts, where used, realize far less than face value, on the order of twenty cents on the dollar across private-target deals outside life sciences once non-paying deals are counted, and the lower middle market fares worse, with smaller deals carrying proportionally larger earnouts (SRS Acquiom, 2025b). None of these mechanisms is inherently abusive, and each allocates real risk. Each is also a lever, available once the seller is off the market.

A short illustration makes the structure visible, and it is only an illustration; the figures are chosen to expose the arithmetic, not to report a frequency. Take two offers for the same business. The first is twenty million from a single buyer that won on price, financing not yet committed, diligence conditions long. The second is eighteen million from a buyer with committed capital, a short and specific diligence list, and a record of closing what it signs. A seller ranking by headline takes the twenty. But suppose the first fails three times in ten and, when it closes, gives back a tenth of the price; its expected realized value is twenty times 0.7 times 0.9, about 12.6 million. Suppose the second fails one time in ten and gives back almost nothing; its expected realized value is near 16 million. The lower headline is worth more. The specific numbers carry no weight here. What the illustration shows is that once q and c are admitted, ranking by headline and ranking by expected realized value can disagree, and the disagreement tends to be largest in the cases least examined.

Exhibit 1 A $20M headline is worth about $12.6M once failure risk and compression are priced.

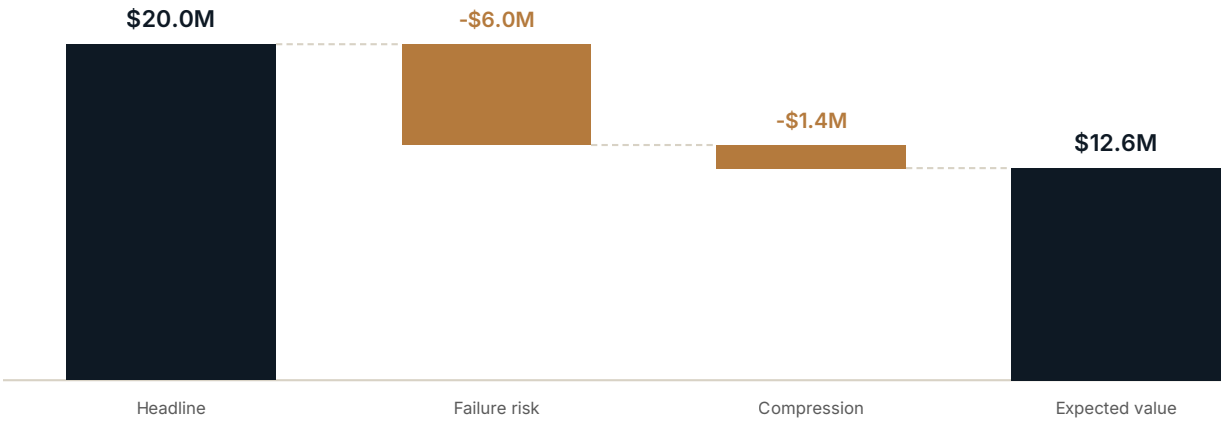

Source · Cordis Institute illustration (figures chosen to expose the arithmetic, not a measured frequency).

Exhibit 3 Earnouts realize roughly 21 cents on the dollar across private-target deals.

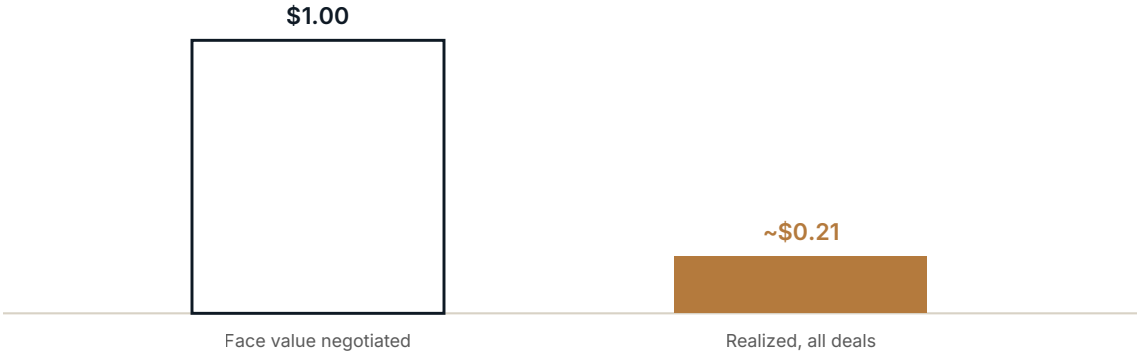

Source · SRS Acquiom 2025, M&A Claims Insights (outside life sciences; lower middle market fares worse).

SECTION 05

Aggressiveness and Certainty.

90%+

DEALS CARRY A PPA

Thin processes invite re-trades; competition suppresses them.

Section 4 establishes only that certainty matters and that q and c move together. The sharper proposition, and the one this paper puts forward for testing, is that offer aggressiveness raises both q and c while process competition lowers both, so that in a thin process the highest headline and the lowest certainty tend to be the same offer. This is argued from the bidding-strategy evidence and from practitioner observation of re-trading; it is not established here on a matched sample of lower-middle-market bids, and it should be read as a hypothesis with a clear test rather than a measured result.

Practitioners describe the mechanism plainly. Re-trading, the revision of price after exclusivity on pretexts that were visible before it, is widely reported to thrive in single-buyer processes and to struggle in competitive ones, and the logic is mechanical: a buyer who knows the seller has no alternative bears little cost in testing a lower number late, while a buyer who knows a credible underbidder waits does not. Competition, in this account, is the variable that compresses both q and c at once, which is why thinness of process, rather than the character of any individual buyer, is the better predictor of trouble.

There is a selection effect beneath the proposition. The bid that clears the room, high enough to win exclusivity and to send the other bidders home, is disproportionately the one written to win rather than to close. This is where the framework connects to WP-002 without depending on it: buyer lanes underwrite the same business to different numbers, and the lane that can justify the highest standalone figure is not always the lane with the most disciplined path to close. For a seller, financing and investment-committee specificity are the observable proxies for a low q that can be inspected before exclusivity is granted. An offer that cannot say where its money comes from, or that reserves the central questions for after the no-shop begins, is in effect disclosing a high q, whether or not the seller registers it.

SECTION 06

A Two-Stage Framework for the Misalignment Tax.

Two leaks, two different fixes.

The certainty discount is most useful read as half of a larger structure, and the structure is offered here as an organizing framework rather than as a demonstrated identity. WP-001 measured the total founder-to-close gap. WP-002 attributed part of it to buyer-lane divergence settled before the letter is signed. The Deal Certainty Discount describes a second part that forms after signing, in the probability of close and the post-signing revision. Framing the total as a pricing-stage component plus an execution-stage component is a proposition about where to look, not a measured decomposition, and the paper does not claim the two components have been shown to sum to the WP-001 total; establishing that would require the joint measurement the series does not yet have.

The framework earns its place by implying different responses at each stage. The pricing-stage leak is worked months before any letter exists, by shaping which lane underwrites the business and on what terms. The execution-stage leak is worked at the moment exclusivity is granted, by preserving competition and demanding specificity. Managing one stage well does not protect a seller who neglects the other, and strong preparation can still end in a weak close if the offer chosen at exclusivity is the uncertain one.

Exhibit 4 The Misalignment Tax forms in two stages, before the letter and after it.

Source · Cordis Institute framework. A proposition about where value leaks, not a measured additive decomposition; the stages interact.

SECTION 07

Implications.

Preserve competition; price certainty as a feature.

Three implications follow, and none requires the seller to accept less. The most important is to keep competitive tension alive into the grant of exclusivity and, where the structure permits, through it, since a credible alternative, even a quiet one, is the strongest discipline on a buyer's incentive to re-trade. Specificity should then be treated as a priced feature of an offer rather than as paperwork: committed financing, a stated investment-committee posture, and a short and concrete diligence list are evidence of a low q, and they deserve weight against a higher but vaguer headline. Finally, offers should be ranked by expected realized value, which means asking of each bid not how high it is but how likely it is to be paid in full. Sellers are rarely encouraged to ask that question, and the reason is structural: an intermediary paid on close does not carry the cost of a deal that fails.

SECTION 08

Limitations and Further Research.

Close-probability is unmeasured for this segment.

The principal limitation is the one named in Section 2. The probability of close is the load-bearing parameter of the model, and no published large-sample estimate of it exists for the lower middle market to the author's knowledge. The familiar "half of letters fail" figure is a practitioner estimate, and the Pepperdine engagement data is suggestive but measures a different object. A defensible estimate of completion for this segment would materially strengthen the framework, and producing one is the obvious next project.

Two further limits follow from the cross-market nature of the evidence. The compression term is documented through the prevalence of its mechanisms, not through a measured average reduction specific to the lower middle market, and the studies establishing prevalence draw on broader samples. And the central proposition of Section 5, that aggressiveness and certainty pull apart, is argued from the bidding-strategy evidence and from practitioner accounts of re-trading rather than from a matched sample of lower-middle-market offers tracked from letter to close. Each of these is testable. A disclosed Cordis engagement panel, with stated size and methodology, would let the failure rate, the compression magnitude, and the aggressiveness-certainty relationship be estimated directly, and is the empirical complement this paper is built to invite rather than to pre-empt.

SECTION 09

Conclusion.

The highest headline is often the least certain.

The headline offer is the most visible number in a sale and among the least reliable. A signed letter is a probability-weighted claim whose expected realized value depends on whether the deal closes and on how much of the price survives the close. The Deal Certainty Discount names that gap, and the proposition advanced here is that in the lower middle market aggressiveness and certainty tend to move against each other, so that the highest headline is often the least certain and a seller ranking by expected realized value will sometimes, rightly, prefer a lower number. Set beside the prior work in this series, the result suggests a two-stage way to read the Misalignment Tax. Value can be lost when the number is set and again when it is tested, and because those are different problems they call for different fixes.

DISCLOSURE

The author of record is affiliated with Cordis, an advisory firm active in lower-middle-market mergers and acquisitions. This paper presents an analytical framework built on publicly available secondary sources and does not draw on confidential engagement data.

DATA AVAILABILITY

All data referenced in this paper are drawn from publicly available secondary sources cited in the references. No proprietary or confidential data were used.

REFERENCES

References.

Boone, A., De Maeseneire, W., Dereeper, S., Luypaert, M., & Nguyen, M. T. (2024). *Unravelling Bidding Strategies in M&A Transactions: Evidence from the Private Phase of the Deal Process.* European Corporate Governance Institute, Finance Working Paper No. 1015/2024. SSRN 4964800.

Everett, C. R. (2025). *Private Capital Markets Report.* Pepperdine Graziadio Business School, Pepperdine Private Capital Markets Project.

Goodwin Procter LLP. (2023). *Durations in M&A Exclusivity Periods Increased Significantly Since 2021.* Goodwin Insights.

Goodwin Procter LLP. (2025). *How Deal and Financing Certainty Wins M&A Competitions.* Goodwin Insights.

SRS Acquiom. (2025a). *Working Capital Purchase Price Adjustment Study.*

SRS Acquiom. (2025b). *M&A Claims Insights: Undisclosed Liabilities and Earnouts.*

Cordis Institute. (2026). *The Preparation Gap (WP-001).* SSRN 6515478.

Cordis Institute. (2026). *The Buyer Lane Preparation Map (WP-002).* SSRN 6735844.