AEIOU ACADEMY

STRUCTURAL LITERACY SERIES: Mergers & Acquisitions

M&A: What Determines the Number

Where the price of a private technology company is actually decided

Christopher J. Velis, NACD.DC

Founder, Auris Health · M&A Advisor, Spine Solutions / Synthes Transaction

August 2026 · White Paper No. 1

aeiouacademy.org · aeioumedical.org

Written by Christopher J. Velis for Velis & Company · info@velisandco.com · 617.528.9482

Disclosure of interests

I represent technology companies seeking exits, and I have done so for years. I have also sat on boards and served as chief executive at companies that were for sale or hoped to be. I expect to hold those roles again.

So this paper is not disinterested, and I am not going to present it as though it were. What it argues about preparation, about how buyers are identified, and about the number a board is handed at the outset describes work I am paid to do. A reader should weigh it accordingly.

Two things follow from that, and I have tried to hold myself to both.

The first is sourcing. Every empirical claim carries a citation to peer-reviewed research, and each is graded. Where I am reasoning rather than reporting, the text says so and labels the inference. A reader who distrusts my motives can check the argument without taking my word for any part of it.

The second is that I have reported the evidence that cuts against me. Section IV sets out a formal result that values one additional bidder above the negotiating skill I sell, and Section VII argues against hiring on the strength of an opening number, which is the instrument my own industry most often uses to win mandates.

Most boards and management teams will run this process once. A very few will run it two or three times. That asymmetry opens this paper and it is the reason I wrote it. The intent here is to step outside all of our roles, mine included, and look plainly at how the process and the outcome can be improved.

Executive summary

A board sells a company once. The people across the table do it forty times a year. Boards hear that asymmetry described as a negotiating problem and hire accordingly. They are hiring against the wrong problem.

Four things set the price of a private technology company. Three are fixed before a buyer is ever contacted: the condition of the records, the number of credible buyers live at the same moment, and whether the consideration is cash or structured. The fourth is the number the board was given when it hired its advisor.

Two simpler explanations fail on the evidence. If the market set the price, two companies with the same asset and the same timing would not sell for different amounts because one kept better records. The accounting research finds that they do. If the table set the price, bargaining power would matter more than the number of bidders. Auction theory proves otherwise. What survives is a short list of conditions established before anyone sits down.

A board's leverage over the first three expires before the process starts. Its leverage over the fourth is exercised once, on the day it picks an advisor. Preparation, buyer identification and consideration structure are decisions, not administration. The opening number is a claim, not a forecast, and it should be interrogated as one.

Every empirical claim below is sourced and graded. Where the evidence runs out, the paper says so.

I. Ready and prepared are different words

Management tells the board the company is ready. Management is not lying. The company has a data room, a folder structure, and most of the documents are in it. What the company lacks is a record built to answer the questions of someone whose job is to find a reason to pay less. That difference does not surface until exclusivity has been signed.

In practice it looks like this.

Not whether you own intellectual property. Whether the assignment documents are proper and signed, including by the contractor who wrote the firmware in 2019 and the postdoc who left for a competitor.

Not whether the cap table is in order. Whether every shareholder was notified, on the record, on time, and in the manner the operating agreement requires.

Not whether management understands the regulatory pathway. How certain they are, and whether that certainty survives a reviewer who has run the pathway before.

Not whether the company keeps records. Whether the books and the design history files are in a condition a stranger can audit.

And the list goes on.....

A director who cannot answer the questions does not know what the company is worth. Neither does the chief executive. Neither, at that point, does the buyer, whose ignorance is the only one of the three that is temporary. The buyer is about to find out, on a schedule it controls, at a moment when the seller has run out of alternatives.

A buyer cannot tell the difference between a company with problems and a company that cannot show it has none. Both look the same from across the table, and both get priced the same way. This part of an exit sits entirely within the seller's control and is almost never exercised. It appears on no process timeline, because it happens before the timeline starts.

II. What the research shows

The claim that records determine price is not advisor folklore. It has been tested.

Skaife and Wangerin built a measure of financial reporting quality from indicators the prior literature associates with less reliable, less relevant and less precise reporting, then applied it to acquisitions where the outcome could be observed [1]. They found three things, and the order matters more than any of them separately.

Acquirers offered higher premiums for targets whose reporting was poor. Those same transactions were then significantly more likely to be renegotiated. They were also significantly more likely to collapse. Targets whose deals failed were more likely to restate their financial statements after announcement.

The generous number arrives first, before deep diligence and verification of facts. The renegotiation arrives months later, after diligence, when the seller has been off the market, management has been off the business, and the market knows a process is running.

Two adjacent studies point to the same thing. Neither is a replication of the other; each looks from a different perspective, one the buyer, the other the seller. Marquardt and Zur found target accounting quality positively related to the likelihood a transaction completes at all [2]. Wangerin, looking from the buyer's side, found that acquirers who cut diligence short under competitive pressure paid for it later in goodwill impairments and weaker performance [3]. That last finding explains something sellers often misread as obstruction. A serious acquirer will not shorten diligence to suit a seller's timetable, because they remember what happened the last time they did.

The obvious objection

A director can reasonably say this is what an advisor would claim, and that a good company sells well whatever its filing discipline. The objection is the right one.

The answer is that nobody who produced this finding has a stake in it. Accounting researchers do not care whether medical technology companies retain advisors. Their result was not that sellers should prepare. Their result was that reporting quality predicts renegotiation and failure. The advice is something a reader adds after reading the finding. The finding is also unflattering to buyers, since it shows acquirers systematically overpaying at announcement for companies they had not yet checked. Nobody in this industry commissioned that.

What the research does not show

These studies look at public companies, because that is where the data is. A privately held, venture-backed company at early revenue is not the study population. However, the mechanism carries. It does not depend on being public. It depends on whether a buyer can verify what it was told, and how long it takes to find out that it cannot. In a private transaction, there is no prior public disclosure to check against, and diligence does all of the verifying, so the effect should be larger.

III. The zone of negotiation

Preparation determines whether a transaction survives. It does not determine what the company receives. For that it helps to be precise about where the value in an acquisition comes from.

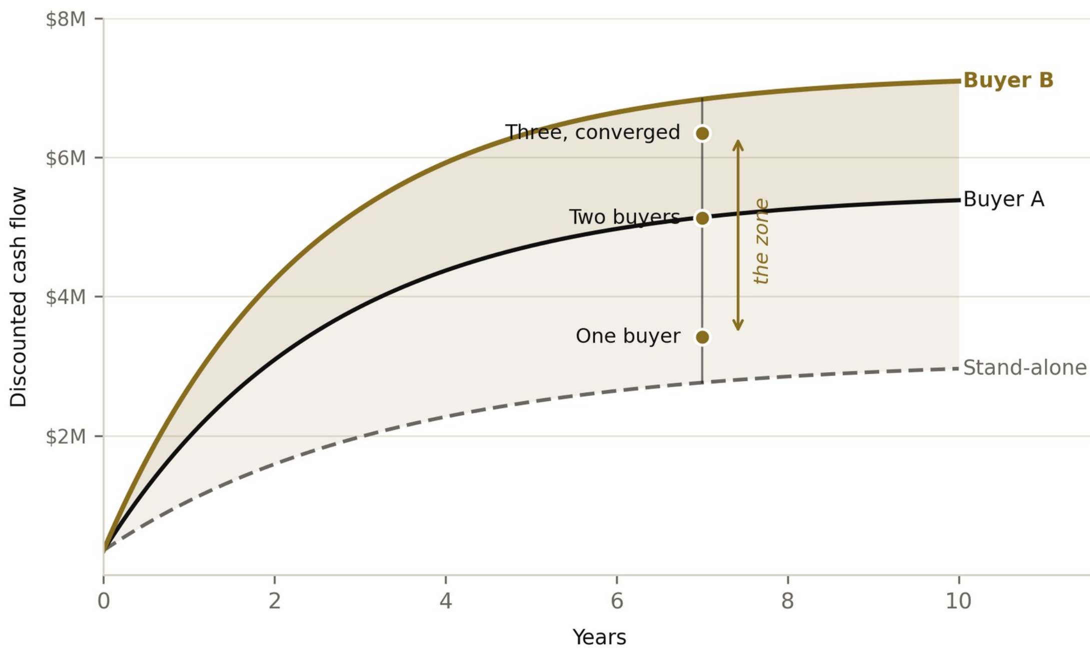

Figure 1. The zone of negotiation. The stand-alone curve is the floor. Each buyer's curve sits higher because the same technology reaches an apparatus already built. Discounted cash flow measures the width of the zone. It is silent on where inside it the price lands, and the number of live parties decides that. Illustrative.

Project the company's free cash flow over time, on its own plan, with its own capital and its own sales channel. Discount it. Discounted cash flow gives you a stand-alone value. A seller can calculate it without a counterparty, and a seller should never take it. One caution on the arithmetic. A private company at early revenue and a strategic acquirer running the same technology through an existing channel do not share a cost of capital. The width of everything that follows depends on which rate is used.

Now put the same technology inside a particular acquirer. It leverages a sales force that already exists, customers that are already in place, a regulatory and reimbursement apparatus already built. The cash flow arrives sooner and it arrives larger. Discount that curve and you have a second number, higher than the first.

The gap between the two curves is value that exists only because the transaction happened. The seller could not produce it alone, and neither could the buyer. Discounted cash flow will tell you how wide the gap is. It tells you nothing about where inside it the price lands, and that is the only question that decides the outcome.

One mechanic decides where it lands. In a process with genuine competition the price is bounded not by the highest bidder's valuation but by the second highest. The winner pays roughly what it takes to beat the runner-up, and everything above that stays with the winner. So the identity of the second buyer matters more than the count. A weak second party raises the price to a weak ceiling, and a board told it ran a competitive process may have run one against the wrong opponent.

That also corrects a convenience I allowed myself above. The stand-alone curve is drawn as the seller's floor, but a board's real alternative is rarely to execute that plan unchanged. The alternative is usually another financing round, a partnership, a licensing structure, or remaining independent, each carrying its own risk and its own probability. The reserve a board should defend is a comparison against those, not against a single curve.

Why the buyer's curve is not a valuation exercise

For transformational technology, there are no comparable transactions, because nothing is comparable. The curve is set by what the technology changes for a specific acquirer, and it runs in three directions at once.

First, what they cannot do now. A category entered without building it internally. A capability they have failed to develop internally, sometimes more than once. A strategic position that becomes defensible. A competitive advantage.

Second, what they stand to lose. This is often the larger number. An incumbent watching a technology shift threaten an installed base and its consumables stream is not evaluating an acquisition. That buyer is working out whether its existing business survives at all, and certainly at what rate it will deteriorate.

Third, what the asset gives them access to. Intellectual property, a regulatory position, a data asset, a platform that unlocks something adjacent. The buyer's technical and legal people usually value this; corporate development, management, and the board factor it into perceived value, and it is often worth more than the revenue attached to it.

These carry different weight with different acquirers. A seller who leads with the wrong one negotiates against the wrong variable and never finds out what was available.

Expect one answer to all of this, because it arrives in almost every process: we will not pay for value we create with our own capital and our own channel. It sounds reasonable. It is not an argument. The acquirer cannot create that value without the asset, and the seller cannot create it without the acquirer.

Neither party has a principled claim to the surplus. Logic does not allocate the surplus. A second party who wants it does.

Value that never appears on the curve

Discounted cash flow prices cash flows. Some assets are not cash flows. They are permissions.

A patent estate that clears freedom to operate in a category is worth the difference between entering that category and not entering it, not the royalties it might earn, and for a particular acquirer that difference is often binary. There is no revenue line for being allowed to compete. A regulatory clearance, a clinical dataset, a manufacturing process others cannot design around, all behave the same way. Finance has a name for this class of asset and a method for it. Luehrman put it plainly: a business strategy is much more like a series of options than like a single projected cash flow, and an investment made now creates the option to invest later [7]. Conventional discounting understates anything whose worth lies in flexibility under uncertainty rather than in a forecastable stream.

The understatement happens inside the buyer's own process. That is what makes it expensive. Christensen and colleagues studied acquisitions meant to change a company's growth trajectory by using the target as a platform. Those are the hardest for an acquirer to evaluate. Chief executives routinely conclude such deals are overpriced. They are also, on the evidence, the acquisitions that pay off spectacularly [8]. The buyer is not being difficult when it says the number is too high. It is applying a method that cannot see what it is buying.

There is a harder version of this, and I have watched it decide transactions. The acquisition that opens a category carries value no later acquisition in that category can carry, because the later ones become possible only after the first has happened. The entrant buys position, and position is priced once. An incumbent watching a competitor take that position compares two futures, not two assets: one in which it competes and one in which it does not. The closest empirical test I have found sits in cross-border transactions. A firm's first international acquisition earns a higher announcement return than its subsequent ones, and the effect is specific to cross-border deals rather than a generic first-deal artifact. The authors read it as a premium attached to the initial move into a new arena [9].

Before anyone accepts this, here is the evidence against a proposition standing next to it. The capitalization hypothesis held that the market prices an entire acquisition program at the first deal, so later deals should show little incremental gain. Asquith, Bruner and Mullins tested it directly and found bidders gained significantly on each of their first four merger bids. Their conclusion was that the results fail to support the capitalization hypothesis [10]. Later work has rejected market anticipation as the explanation for declining serial returns as well.

That finding does not reach the argument here, and I would rather state the difference than gloss over it. Asquith asks whether an acquirer's own shareholders capture a whole program's value at deal one. This section asks whether a particular asset is worth more to a particular buyer because of what it permits. Different question, different unit of analyzis. The entrant argument is an inference built on the option-value literature and the cross-border evidence. It is not a tested finding, and a board should weigh it that way.

None of this is discoverable by auction. An auction finds the highest bid among parties applying their own frames, and the frame is where the error lives. Only someone who knows which driver is live for which acquirer can supply the frame. Supplying it is negotiation. It is not process management.

That is the reconciliation Section IV leaves open. The auction result assumes bidders drawn from the same distribution. Where that holds, competition dominates and negotiating skill is worth little. Where acquirers differ sharply in what the asset permits them to do, the assumption weakens, and what a negotiator knows about a specific buyer's position starts to matter more than the number of buyers. Section VI reaches the same place from the payment structure. This is the same finding arriving from the asset side.

The buyer who has not decided not to be a buyer

Corporate development teams work from a mandate, and the mandate reflects a strategy set some time ago by people who were not looking at this technology when they set it. A company whose stated priorities do not name the category is usually unaware rather than uninterested. No process built on matching an asset to a published strategy will ever reach it.

Reaching it means building the case: what the technology makes possible for that specific acquirer, why the position holds, how it fits what they already run, what it does to their standing in five years, and what it costs them if a competitor owns it instead. The executive who receives that argument has to carry it into his own investment committee and defend the integration plan and the price. Whether he can do that decides whether there is a transaction. The seller's task is therefore to arm one executive, not to persuade him.

IV. One extra bidder

Where the price lands inside the zone is largely a question of how many parties are live at once, and the economics here is unusually clear.

Bulow and Klemperer proved that a seller running a simple auction with one more genuine bidder does better than a seller running the most skillfully optimized negotiation with one fewer, even where the auction gives up every element of the seller's bargaining power including the right to withdraw the asset [4]. The result holds under both independent private values and public values. Their own conclusion was that the value of negotiating skill is small next to the value of additional competition.

Two words in that sentence carry the weight. The bidder must be genuine, meaning it has capacity, a mandate and a committee date, not an expression of interest. And the theorem is silent on what an additional bidder costs. In a live process the cost is not zero: every approach is a disclosure, and disclosure reaches customers, employees and competitors. Cast the net too visibly wide and the serious acquirer concludes the company is being shopped, and leaves. Add one more bidder is free in the model and expensive in the room. A board is entitled to hear how that tension is being managed.

That deserves a moment, given who is writing this. A negotiator with thirty-five closings behind him is telling you the published research says his negotiating skill is worth less than one more real buyer at the table. It says exactly that, and I will not argue with it.

Why a board should prefer the result to the instinct

The natural objection is that a seasoned negotiator has sat in hundreds of rooms and a model has sat in none. In general that objection is sound. Here it fails, for a specific reason. Experience is the wrong instrument for this particular question.

A negotiator sees what happened in the deals he ran. He does not see what would have happened if a third party had been live that month, because that outcome was never produced. Every practitioner's estimate of his own value is built from a sample that leaves out the comparison the question turns on, and confidence accumulates from the visible half of the evidence. A formal result has no such problem. Bulow and Klemperer compare the two arrangements directly.

The conclusion is not that negotiating skill is worthless. The ranking is simply stable. Competition beats skill, and a practitioner is the last person in a position to notice. So the result is not a reason to distrust negotiators. Use it to ask a negotiator what he does about the thing the research says matters more.

And the second objection

Every advisory firm says its process is better, so a board is right to treat process claims as unfalsifiable and discount them. Three things separate this argument from that one.

It is structural rather than a claim about quality. The assertion is not that one firm runs a tidier process. A process producing buyers in sequence produces the outcome of a single buyer, however well each conversation is handled. The next section sets out the mechanism, and a board can check it against its own process by asking which parties were live in the week terms were discussed.

It concedes something a marketing claim would not. The auction result says a board should prefer a mediocre negotiator with three live bidders to an excellent one with two. That is an argument against hiring for negotiating reputation, which is what most firms sell and what I could sell most easily.

It is falsifiable at the board's own table. Section VIII sets out six questions. Advisor rhetoric survives questions about philosophy. It does not survive being asked how three named parties will be brought to a decision point in the same fortnight without a date the market respects. Either there is a specific answer or there is no answer.

V. Two buyers three months apart are one buyer, twice

Proper prior planning prevents piss poor performance. Seven words, no known author, and every officer who has ever run an operation knows it [11]. Sun Tzu said the same thing in a form a board might find easier to repeat: the general who wins makes his calculations before the battle is fought [12].

Both make the claim this paper makes. The negotiation does not begin when the parties sit down. By then most of what will decide the outcome has already been decided.

This is the part I have watched people miss for thirty years. They are impulsive. They are eager to be in a conversation, eager for a close, eager to have something to report. So they take the first live buyer and negotiate against it alone, and they call it a process.

A competing party is worth something only if it is live at the moment the leading party is deciding. Its interest, its fit and its capacity to pay all count for nothing if it arrives after terms are set. Bringing several parties to that moment together, without a date the market respects, is the whole of the art.

And it is done by telling each of them the truth, carefully, at a pace. Not a fabricated deadline. Not a competing bid that does not exist. You do it by reading committee calendars, fiscal years and the personal timelines of the executives who carry the argument inside their own companies. Then you tell each party honestly where the process stands, so that each moves at the rate its own governance allows and all of them arrive together.

The irony is total, and I want to be clear about it. The thing that produces genuine competitive tension is accurate information delivered well. The counterfeit destroys the very tension it was built to create, because an acquirer that discovers it was played does not merely leave the table. It returns, if it returns at all, knowing exactly what the seller's position is, and it prices that knowledge into every term that follows. That last observation is my own testimony rather than a sourced finding, and I mark it as such.

There is no shortcut to maximizing value. There is only the work, done by someone who has done it before, and the discipline to wait when the situation calls for waiting.

Take a process that looks healthy on a status report. Buyer A engages in week one and reaches a decision point in week seven, on a schedule set by its own investment committee. Buyer B surfaces in week ten. Buyer C is not reached until week eighteen, because the buyer list was built from comparable transactions and C did not look like an acquirer until somebody went and asked. All three are real. All three are interested. The board is told, accurately, that three parties are in the process.

Price is set in the window where the leading party demands an answer, not across the process as a whole. Look at what is live inside that window and the answer is one party, every time. The seller negotiates alone, exactly as though the other two did not exist, and three interested buyers produce the outcome of one.

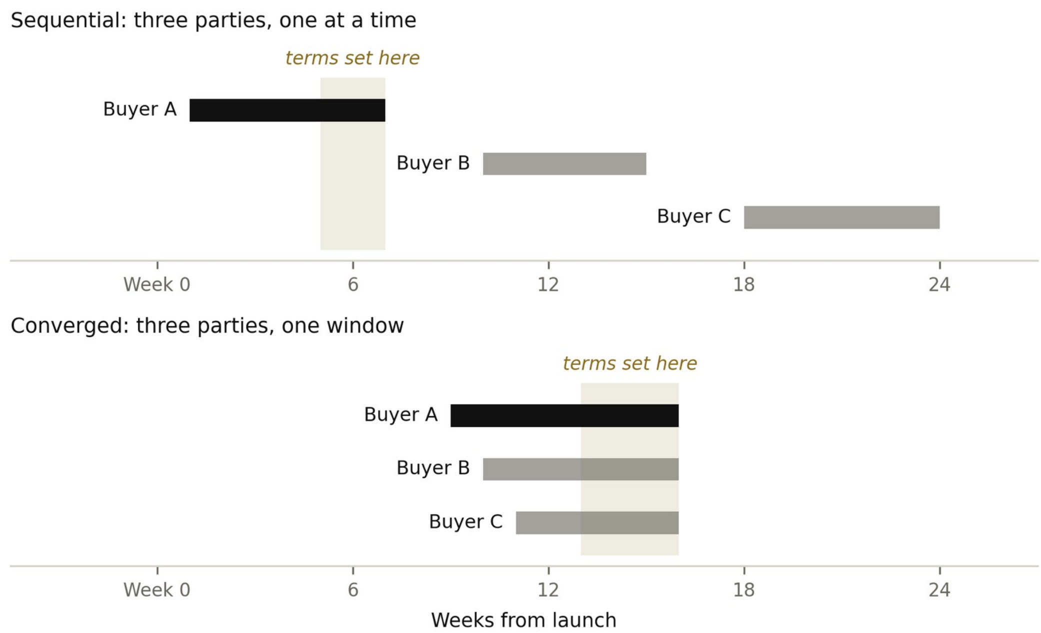

Figure 2. Same three buyers, same company, same numbers. Above, they arrive in sequence and the seller negotiates alone inside the only window that sets terms. Below, they arrive together. The difference is not more buyers. It is the same buyers, converged. Illustrative.

The remedy is not more buyers. It is the same buyers, converged.

Why dates do not solve it

Where an asset is genuinely in high demand and the risk of transacting is low, a process can be run to a date, and the date will hold. That case is the exception, and it appears most often where revenues are large and consistent.

For early-revenue technology companies, deadlines are frequently ignored on purpose. A buyer that misses the date is not disorganized. That buyer is waiting for the auction to break so that it can negotiate against a seller with no alternatives left and a board that has already been told a process was under way. From the buy side, the broken auction is the objective.

Convergence without an enforceable date is produced by managing several counterparties' internal clocks at once: committee calendars, fiscal year ends, budget cycles, conference schedules, and the personal timelines of the executives who have to carry the argument inside their own companies. The work is invisible, it is an art, and it stays invisible until it is missing.

VI. Where competition stops being the answer

The auction result in the previous section carries an assumption that is easy to miss and that changes everything in technology M&A. It assumes the price is cash.

Once consideration can be contingent, through earnouts, equity, royalties or milestone payments, the result no longer holds cleanly. Hoffmann and Vladimirov show that a seller who can negotiate contingent terms often does better against fewer bidders than in a wider auction, because the ability to set the mix of cash and contingent pay is worth more than the extra competition, with reserve prices a secondary matter [5]. They find negotiations more likely to dominate where synergies rise with the buyer's own capability, which is the acquirer and target complementarity case that describes most technology M&A. Boone and Mulherin found the empirical counterpart. Negotiated transactions are at least as common as auctions, and sellers do not appear to do worse in them [6].

This also closes a loop opened in Section II. Marquardt and Zur found target accounting quality positively related not only to completion but to a transaction being structured as a negotiation rather than an auction [2]. Read alongside the result above, preparation stops being hygiene. It becomes leverage over the sale mechanism itself. The prepared seller can choose the arrangement in which structure is worth more than competition. The unprepared seller is left with the arrangement its records permit.

Competition sets the price when the price is a number. When the consideration has structure, what the seller can design is worth more than what the seller can auction.

This is the formal case for deal architecture, and the people it favors barely use it. Most sell-side processes are built to produce competition and stop there. That serves until the buyer proposes an earnout. The whole outcome then moves onto ground the process was never built for, and the seller discovers that the headline number and the realized number are two different things.

Structure as diagnosis

There is a second use for structure that has nothing to do with price. A proposed mechanism is a test.

A buyer that intends to fund and develop an acquired program will accept obligations that bind it to the program: funding commitments with forfeiture consequences, reversion of rights on failure to perform, seller-retained development control, independent subsidiary status, completion incentives. A buyer that intends to shelve the technology, or that is buying defensively, will resist every one of them, and it will resist on grounds that sound reasonable.

Willingness to accept a constraint is evidence about intent that no representation in a purchase agreement will give you. Propose the mechanism and watch what happens to it. You get the answer months before it would otherwise appear, at a point where the seller can still do something about it.

VII. The number as a sales tool

The fourth determinant is the one a board meets first, before any of the work above has been done.

Every firm a board interviews will offer a number. It will be higher than the last number the board heard, because the firm knows it is being compared and knows how the decision gets made. The number costs nothing to say. It wins the mandate. The seller pays for it later.

A high opening number is not a neutral error. It sets a reserve the market will not meet, aims the process at the acquirers who would pay it rather than the ones who will, and produces a process in which nobody bids. Then the number comes down. Then it comes down again. Every reduction is visible to the parties watching, and the buyers who were waiting for exactly that arrive one at a time, each negotiating alone, against a seller whose position is now measurably weaker.

One clarification before going further, because the whole section turns on it. The number under discussion is enterprise value. What reaches common shareholders is what remains after the preference stack, indebtedness, the working capital adjustment, escrow and transaction expenses. The gap between the two is usually larger than the gain anyone is arguing over. A board that fixes on the headline without modeling the waterfall may win the argument it is having and lose the one that matters.

One consequence of a mandate-driven number runs longer than the others and stays invisible until late. Sell-side fees now commonly rise as a percentage as the price rises, and there is nothing wrong with that structure. The problem is where the tiers get set. Tiers built off an opening number that was itself built to win the mandate put the meaningful thresholds out of reach before anyone has spoken to a buyer. The team then works for months against hurdles it cannot clear, and the arrangement meant to align everyone stops doing anything at all. Realistic tiers cost nothing to set and are worth more than they cost. That is the whole of my point on fees, and this paper is not the place for the rest of it.

The mechanism is easier to see in residential property, where asking prices are published. The figures below are illustrative rather than empirical.

> Priced to win the listing: asked $5.9 million, reduced to $5.4 million, reduced again, sold at $4.75. One buyer, no competition, five months of a stale listing behind it.
>
> Priced to win the sale: asked $4.495 million, twenty parties in the first week, cleared at $5.15. Same house, same market, same month.

The asking price was the variable, and the party who set it was the party who wanted the listing.

Private companies are not listed at a price, but the pitch is the listing. The number a board is given at the start shapes which buyers get approached, when a bid gets called disappointing, and how long the board waits before it takes one.

The anchoring tension

None of which makes a low number the safe answer.

Anchoring is real, and it works. The first number stated pulls on every number after it, and a seller who opens conservatively has given away as much as one who opened impossibly. He has simply done it quietly, and he will never find out. Processes that clear well are run by sellers who understood their value, set the reserve deliberately, and had a plan for defending both.

The tension is genuine, and it does not resolve into a rule. Anchor too high and the market never engages, so you never learn what it would have done. Anchor too low and you learn, at a price. No formula finds the point between them. There is only judgment about a particular company, in a particular window, in front of particular buyers.

Height does not separate a working anchor from a failed one. Purpose does. Was it set to win the mandate, or set to win the negotiation? A number chosen because a competing pitch was higher has no argument behind it, and it collapses the first time a buyer pushes. A number built from what a specific acquirer's position can bear survives contact, because when it is challenged there is something to say.

How a board should act under that uncertainty

The absence of a formula does not leave a board without a procedure. Two tests are available, and neither requires knowing the right number.

The first is the derivation test, applied before signing. Ask the advisor to derive the number in front of the board: from which acquirer's position, on what assumptions about that acquirer's alternatives, and with which variables driving it up or down. A number that can be derived can be defended in the room where it matters. A number that cannot be derived is a bid for the mandate, and the board should treat the difference between it and the derivable number as a marketing expense it is being asked to fund with its own outcome.

The second is the escalation test, applied during the process. Decide in advance what a reduction means. One revision after new information is ordinary. A second reduction that arrives without new information is evidence, not a pricing event. It shows the original number had no derivation, and it should trigger a governance response rather than another price discussion. The board reviews who is live, what each party has actually said, and whether the process still has the structure it was sold. Boards routinely spend that moment renegotiating with themselves. The reduction is the signal, and it is worth deciding what it means before it arrives.

VIII. What a board can ask before it signs anything

This paper does not put a director in a position to run a process. It should put one in a position to judge the person who will. Each question has a right answer, and each is hard to answer well without having done the work.

Who will run this day to day, how many transactions has that person negotiated to close, and what else is he carrying? The people who present are frequently not the people who execute, and a board is entitled to the answer in writing before it signs.

What will you find in diligence that we have not told you about? Anyone who has read a data room properly has a list. Anyone without a list has not looked.

How was the buyer list built? A screen of who has bought companies like this one is a different exercise from an analyzis of who stands to gain and lose most by owning it.

How will you get three parties to a decision point in the same window without a date the market respects? The answer is specific or it does not exist.

What is your number based on, what moves it, and what would make you revise it? Then ask the calibration question, which is harder: on your last three completed processes, what was the opening number and what was the realized price? An advisor can derive a number fluently and still be wrong, and the record is the only evidence a board can check.

What do we owe our current advisor if we change, and for how long? Tail provisions in sell-side engagement letters commonly oblige a fee to the prior advisor on buyers it contacted, sometimes for a year or more after termination. A board that switches can owe two fees on one transaction. The tail is negotiable at signing and fixed forever after. Signing day is the only day a board has any leverage over it.

A firm that answers all six well may still be the wrong choice. A firm that cannot answer them is the wrong choice whatever number it offered.

What this leaves unresolved

Three of the four determinants can be worked on directly. Records can be brought into condition. Buyers can be found by strategic analysis rather than by precedent. Consideration can be structured rather than accepted.

One determinant sits outside this paper altogether. When to run a process at all, judged against a company's own trajectory and the state of its market, may matter more than three of the four discussed here. That is a different question and this paper does not answer it.

The fourth cannot be reduced to method. Where to set the number is a judgment made under uncertainty by somebody, on a particular Tuesday, about a particular company. This paper has not resolved that and does not claim to. Anyone who tells a board he has resolved it is describing a sales tool in a bid for engagement.

What can be said is that the judgment improves with the number of times it has been exercised, and that a board is entitled to ask whose judgment is being applied and how often he has applied it.

Velis & Company advises boards of privately held, venture- and private equity-backed technology companies on exit preparation, process design and transaction structure. Every engagement begins with a written evaluation of what a buyer will find. Boards that want only the evaluation can stop there.

Notes

Sources are graded. Tier A means peer-reviewed research or primary regulatory and statutory material. Claims resting on inference rather than evidence are identified as such in the text.

[1] Skaife, H. A. and Wangerin, D. D. (2013). Target Financial Reporting Quality and M&A Deals that Go Bust. Contemporary Accounting Research, 30(2), 719–749. Tier A. Summarized from the published abstract and from descriptions in the subsequent literature.

[2] Marquardt, C. and Zur, E. (2015). The Role of Accounting Quality in the M&A Market. Management Science, 61(3), 604–623. Tier A.

[3] Wangerin, D. D. (2019). M&A Due Diligence, Post-Acquisition Performance, and Financial Reporting for Business Combinations. Contemporary Accounting Research. Tier A. Summarized from the published abstract.

[4] Bulow, J. and Klemperer, P. (1996). Auctions versus Negotiations. American Economic Review, 86(1), 180–194. Earlier circulated as NBER Working Paper 4608 (1994). Tier A.

[5] Hoffmann, F. and Vladimirov, V. (2025). Auctions versus Negotiations: The Role of the Payment Structure. The Journal of Finance, 80(3), 1769–1813. DOI 10.1111/jofi.13446. Tier A.

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

[7] Luehrman, T. A. (1998). Strategy as a Portfolio of Real Options. Harvard Business Review, 76(5), 89–99. Tier B.

[8] Christensen, C. M., Alton, R., Rising, C. and Waldeck, A. (2011). The Big Idea: The New M&A Playbook. Harvard Business Review, 89(3). Tier B.

[9] Dandapani, K., Hibbert, A. M. and Lawrence, E. R. (2020). The Shareholder's Response to a Firm's First International Acquisition. Journal of Banking & Finance, 118, article 105852. Tier A. Finding and interpretation verified at the publisher; sample size and period not verified, and none is asserted here.

[10] Asquith, P., Bruner, R. F. and Mullins, D. W. Jr. (1983). The Gains to Bidding Firms from Merger. Journal of Financial Economics, 11(1–4), 121–139. Tier A. Cited against the interest of the argument in Section III. See also Ahern, K. R. (2008), The Returns to Repeat Acquirers, working paper, which likewise rejects market anticipation as the explanation for declining serial acquisition returns.

[11] "Proper prior planning prevents piss poor performance," an adage of uncertain authorship long associated with British and United States military planning. Variants substitute "planning and preparation" for "prior planning." Not attributable to a datable source; cited here as a maxim, not as evidence.

[12] Sun Tzu, The Art of War, ch. I (Giles translation, 1910): the general who wins a battle makes many calculations before it is fought. See also ch. IV, that the victorious strategist seeks battle only after victory is already assured. Note: the widely circulated line "every battle is won before it is ever fought" does not appear in the text. It reached popular use through the 1987 film Wall Street and is not Sun Tzu.