

# The **C**ompetitor

Issue 1 | 2026



## IN THIS ISSUE

**The real world is full of games: When AI replaces people, prices them out, or preserves human work**

Page 3

**The three levers that determine every negotiation – and the mistakes that neutralize them**

Page 8

**Why AI changes the logic of negotiation – but does not replace it**

Page 12

## Welcome to the latest issue of The Competitor.

AI has arrived in procurement – but at most companies it stops at the copilot stage. Summarizing documents, pre-structuring bids, generating reports. Useful, but far removed from what is actually possible. The reason is almost always the same: risk awareness outweighs the courage to implement. More capable models are not being tested, autonomous agents are not being deployed, structural advantages are not being built – because the organization would rather hedge than experiment.

At the same time, research is already showing where this is headed: in our current research project, Agentic Economics, we are studying how AI agents can be guided through economic incentives – compensation, reputation, skin in the game. The early results are promising: a cheaper AI model with economic incentives can come close to the performance of a much more expensive model without incentives. The principle behind this is not new – it is classical principal-agent theory applied to autonomous systems. Whoever understands these mechanisms can do more than use AI; they can steer it.

This issue brings together three perspectives on the interplay of AI, negotiation, and strategic procurement.

In the first article, we develop a framework that moves the debate about AI and employment to the task level – where it becomes operational. Two variables produce four clearly distinguishable development paths: a diagnostic tool for anyone who wants to assess their own position – or that of their organization – realistically.

In the second article, we show why the negotiation process – not the conversation itself – is the decisive playing field, and which concrete levers procurement professionals have to shape that field. We also address common misconceptions that systematically neutralize those levers.

The third article addresses the question of what AI should take over in negotiation processes – and what it structurally cannot. Drawing the line between tail-spend automation and strategic negotiation leadership is one of the most important decisions procurement organizations are making right now.

We hope this issue gives you useful perspectives for your day-to-day work, and as always, we look forward to your feedback.

*Christoph Pfeiffer*

# The real world is full of games: When AI replaces people, prices them out, or preserves human work

---

*The debate about AI and work is stuck in a false dichotomy: does AI replace people – or make them more productive? The answer is neither one nor the other. It is predictable – task by task – based on two structural variables.*

---

## 1. From jobs to tasks

The question “Will AI replace my job?” is too coarse to guide action. Jobs are bundles of tasks. Technology will fully take over some tasks. Others it will not. The question becomes answerable only when we ask it task by task.

Two structural variables determine the trajectory of every task: first, how open is the problem – can it be fully specified in formal terms, or does it resist a complete description? Second, how steep are the payoffs – does the market reward quality differences proportionally or disproportionately?

## 2. Is the problem open or closed?

A problem is closed if a fully specified mapping from state to optimal action exists in principle – even if it is computationally expensive. A problem is open if such a mapping cannot, in principle, be fully specified because boundary conditions, evaluation criteria, or solution paths remain indeterminate or endogenous.

What matters is this: the issue is not difficulty. Chess is enormously complex, but fully closed – 64 squares, defined pieces, complete information. Zermelo’s theorem proves that a perfect strategy exists. Medical diagnosis, by contrast, is often less complex than chess but structurally more open: symptoms are ambiguous, patient histories are incomplete, and identical symptoms can have different causes depending on context.

The structural consequence is this: closed problems tend toward substitution – if technology can in principle solve the problem completely, the structural rationale for human involvement erodes. Open problems tend to preserve a human role – because machines and humans are fundamentally different information-processing systems that make different errors and recognize different patterns.

Openness is not a binary attribute but a spectrum. The same profession can appear at different points on it: a lawyer reviewing contracts (more closed) and a lawyer designing a negotiation strategy (highly open) occupy structurally different positions. The same is true for a buyer verifying invoices (closed) and a buyer designing an award mechanism (open). The unit of analysis is the task, not the job title.

A third variable – complexity – affects the speed, not the direction. Simple closed problems are substituted immediately. Complex closed problems take longer. Complexity is the time constant, not the vector.

### 3. How steep are the payoffs?

Here lies the insight most of the AI debate misses. In open problems, humans often remain meaningful complements – but complementarity alone is not enough protection. People are expensive. And as technology advances, the marginal contribution of the human shrinks. The decisive question is: does the market reward that shrinking marginal improvement enough to justify its cost?

In markets with linear payoffs, slightly better performance yields slightly higher rewards. As the human’s marginal contribution shrinks, the pay attached to it shrinks as well. At some point, the costs exceed the benefits. The human is priced out – not because no value is being created, but because the value created no longer justifies the cost.

In markets with steep payoffs – winner-takes-most structures in which small quality differences produce disproportionately large differences in reward – the economics works differently. The gap between the best and second-best result can be tenfold or hundredfold. If both sides have access to the same technology, the human contribution – small as it may be in absolute terms – becomes one of the remaining sources of differentiation. The steep payoff structure ensures that even small human advantages continue to pay off.

For procurement organizations, sourcing standardized C-parts has linear payoffs – a slightly better price produces a proportional advantage. Negotiating a strategic framework agreement with a key supplier has steep payoffs – one percentage point more or less can amount to millions over the life of the contract.

### 4. The 2×2 framework: Four development paths

Combining these two variables yields four qualitatively different trajectories.

	Linear Payoffs	Steep Payoffs
Closed Problem	Replaced	Closing Window
Open Problem	Priced Out	Decisive

## Quadrant I – Replaced (closed + linear)

The problem can be fully specified, and the market rewards meeting a standard, not exceeding it. Technology solves the core task; the human becomes economically redundant. The textbook example is the ATM: the task was fully specifiable, and the market did not reward human precision disproportionately. Substitution of the core task was immediate and complete – and did not even require AI. Interestingly, banks then opened significantly more branches, because lower costs made expansion profitable. Total employment initially rose – but in different, higher-value tasks. That did not help the individual teller, who was not automatically a good advisor.

## Quadrant II – Closing Window (closed + steep)

The problem is closed, but steep payoffs create a transitional phase in which human-machine hybrids make economic sense. In 1997, Deep Blue defeated Kasparov. In 1998, Kasparov launched Advanced Chess: human-machine teams competing against one another. In 2005, at a freestyle tournament, two amateurs with three ordinary computers defeated both the strongest engine and grandmasters with better hardware – through strategic orchestration. For nearly two decades, the steep payoff structure rewarded the human who could steer technology better than the opposing human. Then the window closed. Today, at the top level, human intervention worsens the outcome. Steep payoffs protected the human only as long as there was a human contribution for them to amplify. In closed problems, that contribution eventually disappears.

## Quadrant III – Priced Out (open + linear)

The problem resists full specification – the human creates real value. But the market pays proportionally: a slightly better result brings slightly more. If the human’s marginal contribution shrinks to 2 percent, then 1 percent, then 0.5 percent, the cost no longer justifies the difference. Professional translation is the textbook example: machine translation now handles a large share of global translation volume. Human translators are increasingly concentrating in subsegments with steeper payoffs: literary translation, luxury brand localization, diplomatic communication. The same profession contains both Quadrant III and Quadrant IV segments.

## Quadrant IV – Decisive (open + steep)

The problem resists full specification, and the market rewards the best performer disproportionately. Two mutually reinforcing mechanisms are at work here: first, the human remains a complement because the task cannot be fully formalized. Second, the steep payoff structure turns even a small marginal advantage into a large difference in rewards. When all players have access to the same technology, the technological advantages cancel out – and human judgment becomes one of the primary sources of competitive advantage.

That is why this article’s title is no coincidence: the real world is full of games – not optimization problems. In an optimization problem, the best solver wins and the human is irrelevant. In strategic environments – where outcomes depend on interaction, adaptation, and relative positioning – the remaining competitive advantage shifts from access to the tool to judgment in its use.

## 5. The core-job principle

The key is to apply the framework to the core function – not to peripheral activities. Clayton Christensen’s Jobs-to-be-Done thinking provides the bridge: what is the customer actually paying for?

A taxi driver who holds the door and makes small talk provides genuine human complementary value – but in peripheral tasks. The core job is transportation: closed, linear payoff. Quadrant I. No amount of warmth changes the trajectory of the core task. An accountant who builds deep client relationships provides important relationship value – but the core job is accurate bookkeeping: structurally closed, with largely linear payoffs. Quadrant I.

Many professionals define their value unconsciously through peripheral contributions – precisely because they sense that the core task is under pressure. The framework makes that self-assessment visible.

## 6. Within a profession: Medicine and procurement

The explanatory power of the framework is most evident within a single profession.

Routine diagnostics in an employed setting: the problem is open (every patient is different), but the payoff structure is linear. The hospital is paid roughly the same whether the human adds marginal value or not. If AI independently reaches 95 percent accuracy, the human’s marginal correction shrinks to a few percentage points that may no longer justify the cost. Quadrant III.

Elite oncology at a top institution: the problem is equally open, but the payoff structure is steep. The best diagnostician attracts the most difficult referrals, and reputation compounds over time. The gap between the best and second-best outcome is enormous. Quadrant IV.

Same medical degree. Same open problem. Different trajectory – because the payoff structure is different.

The same applies in procurement: operational invoice verification sits in Quadrant I – fully automatable. Should-cost analyses for standard products move toward Quadrant III – AI delivers “good enough” results. Designing complex award mechanisms for strategic categories remains in Quadrant IV – here, human judgment determines disproportionately different outcomes. CPOs should neither automate across the board nor rely on human capability across the board; they should place each task individually within the framework.

## 7. The social problem

This is where the analysis becomes uncomfortable. Humans remain economically strong primarily in Quadrant IV – precisely where steep payoffs already tend to concentrate income. AI amplifies that mechanism.

An entire generation was advised to choose stable, predictable careers: accounting, law, medicine, administration. Many of these are linear-payoff environments. The places where human work remains economically defensible are inherently riskier: entrepreneurship, dealmaking, strategy, competitive markets. And education systems are still largely optimized for closed problems with linear payoffs: follow rules, pass tests, produce correct answers to defined questions.

---

Earlier waves of automation replaced physical tasks – and the new tasks were cognitive. That worked because much of the population could move into cognitive work. AI is increasingly substituting for cognitive tasks. Many of the new, higher-value tasks will concentrate in more open, steeper environments. Not everyone can move there. That does not mean collapse – but it does mean the transition requires conscious design, not faith that the market will sort it out on its own.

## 8. Conclusion

The real world is full of games – not optimization problems. In optimization problems, the best solver wins and the human becomes irrelevant. In strategic environments – open, competitive, with steep payoffs – human judgment remains a decisive differentiator. Whoever understands the structural position of their own work can act strategically. Whoever does not remains the object of a development whose logic they can feel but do not understand.

# The three levers that determine every negotiation – and the mistakes that neutralize them

---

*Negotiation power has nothing to do with negotiating skill. Confuse the two and you lose – structurally, systematically, and repeatedly. This article shows the three levers that actually determine negotiations, and the most common mistakes that render them ineffective.*

---

## 1. Levers, not arguments

You have the cost data. You have the margin analysis. You can show that the supplier is 30 percent above the market. You present all of it – and the supplier holds its position.

This is not an isolated case. It is the standard outcome when negotiators do not understand the decisive difference: the difference between an argument and a lever.

An argument appeals to reason. A lever changes the other side's business case. As long as your analysis has no consequence for the supplier's calculation – no real threat, no measurable loss – it is economically irrelevant. Saying "Your margin is too high" is not a lever. It is an appeal. And appeals move nothing in negotiations because they do not affect the other side's incentive structure.

The test is simple: does this measure change the other side's incentive structure? If not, it is not a lever – no matter how well supported it is. What are real levers? Competition, bundling, specification changes, a credible make-or-buy option. Cost data becomes a lever – but only when it provides the basis for one of these actions.

## 2. Competition: The strongest lever – and the one most often simulated

Competition is the most powerful instrument in procurement. The news is not that it exists – it is that it is systematically underdeveloped.

Everyone knows the standard pattern: the contract expires. The incumbent supplier is invited. Price is negotiated. An agreement is reached – or one side gives in. The result is not a market price. It is the result of a bilateral conversation without a real alternative.

Competition does not mean collecting three bids. It means creating credible alternatives – other suppliers, in-house options, the realistic possibility of stopping a project or fundamentally restructuring it. Credible is the decisive word. A supplier who knows the alternative is not being seriously pursued prices exactly that into its position. Faux competition is often more dangerous than no competition at all – it creates effort without creating pressure.

What most people miss: even weak bidders create pressure. Intuition says otherwise – a new bidder with a higher bid than the incumbent supplier seems useless. The opposite is true. The weak bidder changes the incumbent supplier's probability calculation. As long as a competitor remains in the race, the incumbent supplier cannot be certain of winning. It has to consider a concession to protect its probability of winning. The weak bidder does not win the business – but it forces concessions that would not arise without it. The cost: more preparation. The benefit: a structurally stronger negotiating position. That is not a bad trade-off.

### **3. Comparability: Whoever defines the rules of the game wins**

Competition without comparability is blind. Three offers based on different logics are not competition – they are structured opacity.

Comparability means: all options are translated into a common monetary model. Not just the quoted price – all relevant cost differences. Switching costs, integration effort, quality differences, logistics, risk premiums, payment terms. Every advantage and disadvantage gets a price. Only then is comparison possible. Before that, it is an estimate.

The Delta Cost of Ownership model is the right instrument for this. It attacks the incumbent supplier's real advantage – not its lower price, but the incomparability of its position. Long-standing suppliers know the internal processes, the switching costs, and the buyer's real willingness to pay. New bidders have to price in risk premiums – and therefore appear more expensive without truly being so. The fact that new suppliers bid higher is not a quality signal; it is a sign of missing comparability.

Creating comparability is therefore not an administrative step. It is a strategic attack on the incumbent supplier's information asymmetry. The operational consequence is clear: suppliers respond in the buyer's pre-structured pricing model, not their own. Anyone who fails to enforce that hands definitional power to the other side.

### **4. Commitment: The precondition without which everything else is worthless**

Competition and comparability create impact only under one condition: suppliers have to believe that the process holds. Commitment is therefore not the third lever – it is the precondition that makes the first two work.

Commitment means: the defined mechanism is binding. The best offer in the final round wins. No renegotiation. No bilateral exceptions. No extra round for the incumbent supplier when it applies pressure.

That sounds obvious. In practice, it is the procurement principle most often broken. Well-intentioned flexibility. Escalation to senior management. The desire “not to jeopardize the relationship.” The outcome is always the same: suppliers learn that the process is not binding. That there is a second round. That the real negotiation begins after the formal end.

Anyone who breaks commitment destroys more than the current process. They retroactively destroy the value of every negotiating position that has been built. Because the effect of competition and comparability rests on a single expectation on the other side: that the mechanism counts. Once that expectation collapses, the lever collapses with it.

Commitment is not a question of individual discipline. It is a governance issue. It requires clear escalation rules, predefined decision points, and management backing that is not negotiable. The CPO does not belong in the first negotiation round. The CPO is the final escalation signal – and loses its effect as soon as it is used too early.

## **5. Frontloading: Where the three levers have their maximum effect**

Competition, comparability, and commitment are not tools for negotiation day. They must take effect earlier – in the RFP phase, long before the first commercial conversation.

Frontloading means: all relevant information, pricing models, and evaluation criteria are introduced as early as possible. Suppliers receive the buyer’s pricing model. They are asked for package prices, bundling options, and alternative payment terms – not to confuse them, but to make their commercial degrees of freedom visible before bilateral dynamics distort the process.

The key discipline here is that technical questions are closed before the commercial round. Every open technical issue in the commercial conversation is a legitimate tool for repricing. Suppliers use it. The rule should be clear and consistently enforced: technical clarity before the commercial round. After that, the playing field is considered defined.

## **6. Three traps that neutralize every lever**

Even a cleanly designed process can be devalued by specific execution errors. Three patterns are especially common.

### **Trap 1: Early offers under time pressure**

An early, unusually high discount tied to time pressure or special terms feels like a win. That is exactly the problem. This pattern creates reciprocity pressure – whoever receives something feels obliged to give something back. It creates time pressure that crowds out rational analysis. And it prevents the buyer from systematically building its own negotiating position.

An offer is only good if it has been compared with your reservation price and your realistic alternatives. Without that comparison, “good” is not an economic judgment – it is a psychological feeling. And feelings are manipulable. The rule is simple: no deal before all advantages and disadvantages have been monetized.

## Trap 2: The indexation error

Indexation negotiations are one of the most overlooked sources of structural overpayment. A price from past years is rolled forward each year based on inflation or commodities and is treated internally as fair because it has been “adjusted to the market.”

That is exactly the wrong conclusion. An inflated base price is not corrected by indexation – it is carried forward. If you are 15 percent above the market today, you were 15 percent above the market then. The index merely became more efficient at cementing the gap.

The right question in an indexation negotiation is not “Which index applies?” It is “Is the base price still right?” Productivity gains on the supplier side, market price developments, and technological shifts have to flow into the base. Anyone who does not insist on that is financing the supplier’s efficiency gains.

## Trap 3: Designing after rather than before the negotiation

The evaluation matrix belongs before the negotiation, not after it. If suppliers know which criteria and weightings will determine the decision, they change how they bid and how they negotiate. Transparency about the scoring model is not an information leak – it is behavior shaping through structure.

Likewise, lot design is a strategic instrument, not an administrative step. Lots that are too large exclude mid-sized suppliers and reduce competition. The optimal lot structure creates maximum, credible competition – and therefore the strongest lever.

And finally: management backing is not a symbol; it is a credibility signal. When suppliers know the process is binding and that no backchannel conversations will take place, their behavior changes fundamentally. Credibility is not a personality trait – it is a structural property of the negotiation design.

## Conclusion

Competition, comparability, commitment. Three levers. One system. Whoever builds it consistently wins structurally – before the conversation, not in it. Whoever knows the three traps protects the advantage they have built from being eroded by process errors. The best negotiators do not win through better arguments. They win because they designed the game better than the other side. Competition, comparability, commitment – these are not prerequisites for a good negotiation. They are the negotiation.

# Why AI changes the logic of negotiation – but does not replace it

---

*AI does not replace negotiators. It exposes the wrong negotiators – the ones who believe the conversation decides. It is preparation that decides. And that is exactly where AI is changing the rules fundamentally.*

---

## 1. The wrong question costs millions

“Will AI replace our buyers?” That is the question being asked in executive suites. It is the wrong one.

The right question is: which parts of the negotiation process should AI take over – and which parts must it under no circumstances take over? Anyone who fails to make that distinction commits one of the costliest mistakes in modern procurement: automating the wrong things and leaving the right things to chance.

The reason many AI projects miss their expected results is rarely the technology. It is the approach. Most companies ask, “Which problem can we solve faster with AI?” That is optimization thinking – and optimization thinking, at best, scales what already exists.

The more effective approach reverses the logic: what would the best possible outcome be if technological constraints played no role? Then work backward to the present and build the path there. Not “How do we speed up the existing process?” but “What would the process look like if we reinvented it today?” That difference is not semantic. It determines whether AI remains a tool or becomes a structural advantage.

## 2. Two worlds, one fatal confusion

In procurement, there are two fundamentally different negotiation realities. Confusing them is the most common reason AI investments in procurement disappoint.

The first world: tail spend. Supplier known, product defined, price negotiable, process repeatable. Here, the task is fully specifiable. AI acts autonomously, without human intervention, with consistent results. In this world, the human usually adds no further value in the process – and often even slows execution.

The second world: strategic negotiations. Objectives are not fully defined, the other side acts strategically, and information is distributed asymmetrically. The outcome depends on trust, credibility, timing, and power dynamics – factors that cannot be fully modeled. Here, full AI autonomy is not a robust option. It increases risk.

Anyone who uses AI in strategic negotiations the same way they use it in tail spend is not making an implementation mistake. They are making a conceptual mistake – and handing strategic information advantages to the other side.

### **3. Where AI has the greatest leverage: before the conversation**

The most powerful use of AI in negotiations does not take place in the negotiation room. It takes place weeks earlier – in preparation.

Research and market intelligence form the foundation of every strong negotiating position. And this is exactly the area that most often collapses under time pressure in practice. An experienced buyer needs days to analyze ownership structures, identify antitrust-relevant interlocks, and reconstruct market prices with confidence – and even then remains systematically incomplete. AI does the same in hours, with greater depth and without fatigue effects.

The same applies to analyzing the negotiating counterpart: decision logic, historical behavior patterns, public positioning, known preferences. If you understand the person across the table before the first word is spoken, you shape offers, arguments, and timing not by intuition – but by evidence.

Cost transparency is the third lever. In award processes, knowledge of real market prices and cost structures determines whether a negotiating position is credible – or merely sounds that way. AI delivers that foundation faster and more robustly than any manual approach.

### **4. AI as a sparring partner – not as the decision-maker**

Beyond research lies the second application area: strategic preparation of the negotiation itself. Here, too, the distinction between support and autonomy applies – and it is absolute.

A coherent negotiation strategy includes objectives, reservation prices, BATNA, and concession planning. AI can serve as a structured sparring partner through this process: it asks uncomfortable questions, tests the consistency of positions, and uncovers blind spots. The final decision stays with the human – but its quality rises measurably.

The same applies to the simulation of bilateral negotiations: anticipating counterarguments, playing through concession sequences, thinking through roles and timing. Experienced buyers do this intuitively. AI makes it systematic – and therefore scalable.

The most demanding field is game-theoretic and auction-theoretic design. In complex tenders with multiple suppliers, interdependent lots, and strategic interactions, the design of the tender process itself becomes a negotiation instrument. Which auction format maximizes competition? How should lots be split? Which information should be disclosed when? These questions follow a formal logic. AI systems with a game-theoretic foundation can support this more systematically than pure human intuition – especially at high volume and with many variables.

## 5. The structural reason AI does not negotiate

Here lies the central insight that carries the entire argument: there is a fundamental difference between designing the game and playing the game.

AI can design the game. It can build mechanisms, calculate scenarios, and optimize information structures. What it cannot do is play.

Strategic negotiations are structurally open and have a steep payoff function. Small quality differences in negotiation execution produce disproportionately large differences in outcome. That makes them a domain in which human capability does not merely remain relevant – it is the decisive differentiator.

The reason is precise: negotiations are not an information problem. They are a credibility problem. The effectiveness of a threat depends on whether the other side believes you will carry it out. Credibility arises through nonverbal signals, reputation, and the perceived willingness to escalate. AI has no reputation. It cannot assume responsibility. It does not build trust.

There is also the adaptive complexity of the real conversation: dynamics shift, unexpected information appears, the atmosphere turns. A negotiator reads those signals and responds in real time. That is not delegable – neither to a system nor to an algorithm.

## 6. What follows from this – in concrete terms

Three conclusions for procurement decision-makers.

Segmentation is the prerequisite for everything else. Without a clear separation between fully automatable processes (tail spend), AI-supported preparation (strategy, research), and inherently human execution (negotiation leadership), resources will be misallocated. This segmentation is not an IT decision. It is a strategic one.

Preparation is the greatest untapped lever. Companies that deploy AI consistently in the preparation phase have a structural advantage over those that do not. That advantage is not marginal. It is the difference between grounded strategy and well-intentioned improvisation.

Accelerate capability building. Experiential knowledge – the ability to anticipate counter-strategies, understand game-theoretic dynamics, and structure arguments consistently – normally takes years of practice to acquire. AI-supported systems make that knowledge accessible earlier. That not only accelerates the onboarding of new buyers. It changes the capability profile of an entire procurement organization – and makes investment in structured negotiation preparation scalable.

---

## Conclusion

AI changes where negotiations are won. Not in the conversation – that was always the wrong arena. In the process. In the preparation. In the design. Whoever uses AI there wins structurally. Whoever looks for it in the negotiation room is looking in the wrong place.

The question is not whether to use AI. The question is whether you understand where it matters – and where the human remains irreplaceable.

# The mpetitor

Questions, feedback, criticism, topic ideas or new signups?

Please send them to  
[newsletter@competitio.de](mailto:newsletter@competitio.de)