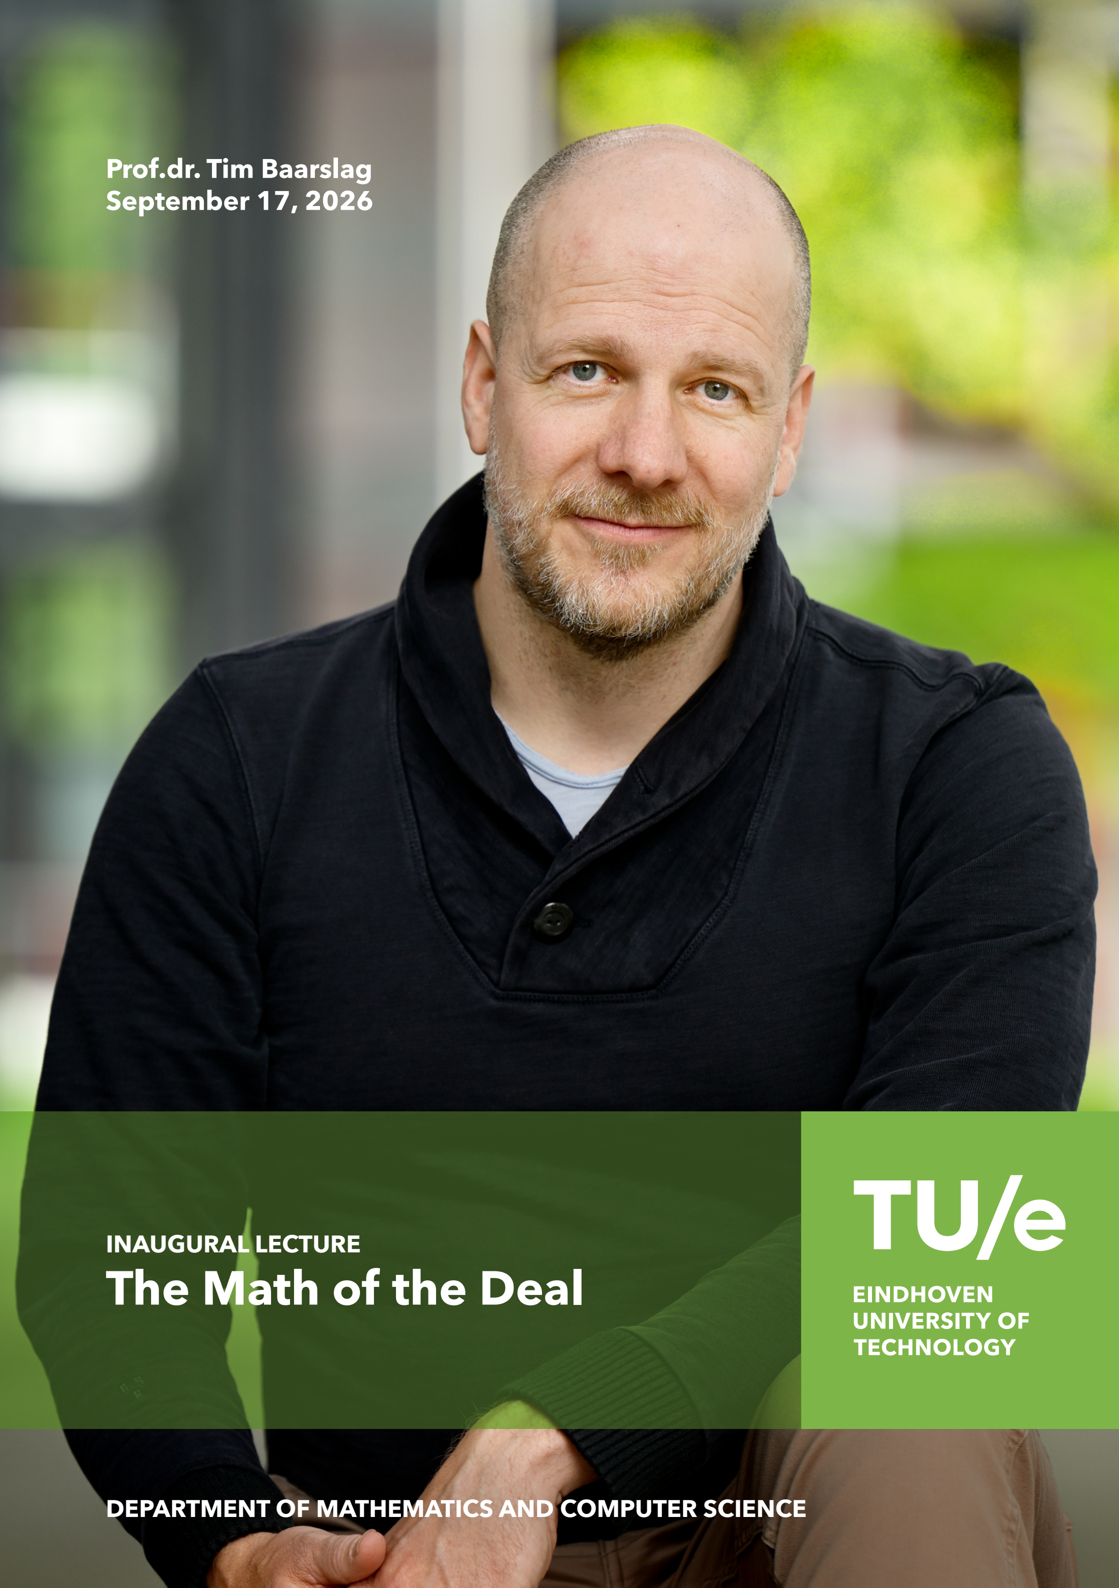

Prof.dr. Tim Baarslag

September 17, 2026

INAUGURAL LECTURE

The Math of the Deal

EINDHOVEN UNIVERSITY OF TECHNOLOGY

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE

INAUGURAL LECTURE PROF.DR. TIM BAARSLAG

The Math of the Deal

September 17, 2026
Eindhoven University of Technology

Welcome

Madam rector magnificus, the dean and members of the Executive Board, members of the Department Board of Mathematics and Computer Science, dear colleagues, students, family and friends.

How the negotiation seed was planted

Every year, during Queen's Day – the national holiday on which the whole country turns into a street market – I would go out as a little kid, with a guilder burning in my pocket, to score as many Donald Duck comics as possible. I'd pass the stalls all day, haggling about the price. The asking price would vary wildly: some would charge a guilder apiece, but there was also this little kid that sold them for just a few cents.

It was nearing the end of day and I still hadn't bought anything. And much worse, the little kid with the cheap Donald Ducks had gone! With a heavy heart, I decided to make one more round and found two old guys (I now realize they were probably teenagers) who were clearly done for the day and were packing up crates full of unsold Donald Ducks. With wide eyes I asked them what I could possibly get for my guilder. One of the guys looked at my guilder and exchanged looks with his friend, who nodded affirmatively, and he spoke the magic words: "You know what? For that guilder, you can have them all."

I can still recall the fear of losing my way back as I sprinted to my parents to load up these treasures in our car. I was blissfully happy, and my fascination with negotiation had begun.

The Math of the Deal

Fast forward a few decades and I am very honored to stand before you as a professor in this very area. In these 45 minutes, I'd like to talk about the scientific, mathematical insights in getting a good deal. And so, the title of this talk is 'The *Math* of the Deal'.

I'd like to give you a taste of the beauty of the mathematical problems arising in negotiation. I'll give an overview of the field and my work in particular, brightened, hopefully, by anecdotes. And I'll dive deep at carefully selected moments, hoping that you'll bravely jump with me.

Of course, the title is also a topical nod to *The Art of the Deal*. Its author sees the world as a zero-sum game, where you lose when the other wins. This might conjure up thoughts about taking what you can get, manipulation, clever maneuvering, and even outright lying.

But in my research, I want to stay away from the self-help books and focus on the analytical side of negotiation: the *Math* of the Deal. The idea is as follows. Suppose that almost everything would be clear (mathematicians would say: well-defined). You know what you want and you have certain beliefs about everything you do not know. Even then, it is still not always clear what to do. But let's solve that first and understand how the result depends on our beliefs. This is the so-called decision theoretic viewpoint and it allows us, at the cost of certain assumptions, to investigate the underlying principles of negotiation.

Think, for example, of buying a house. For anyone who has managed to get onto the housing market – and lately that is fewer and fewer of us – it is a rude awakening to the importance of negotiation. It's scary, you have to be careful what you say, you can't be sure of what others are doing, and a mistake can be extremely costly.

We try to tame that unrest by trusting an intermediary (a broker) or by reaching for the few things we can control, such as writing the seller a personal letter.

Instead, I ask: how can we learn what the other finds acceptable? How do you measure that? And can we then calculate the winning offer? Why exactly am I making a 15,000-euro concession and not less? In principle, this line of reasoning could save a good deal of money.

These questions induce all kinds of pocket-sized mathematical machinery. They are highly stylized and idealized, mind you, but wherever the assumptions hold, so do the insights – always and everywhere.

Negotiation as a scientific discipline

The first to seriously consider this subject and the mathematical way of thinking about negotiation were Nash and von Neumann [1], [2]. This kind of research is traditionally at the interface between math and economics. We owe some beautiful theorems to this approach, but they require very strong assumptions on the rationality of the negotiators, to the point where the whole negotiation collapses into one offer, accepted on the spot [3]. This has always been gnawing at us: what we see every day actually bears little resemblance to these results.

- The agent community thinks about the problem a little differently. They view a negotiator, even an automated one, as a small autonomous and intelligent entity: an agent [4], [5]. The agent perceives the environment (often, a bid coming in), does some deliberation, and sends a bid back [6].
- The decision-theoretical view on negotiation owes a lot to Harvard Business School and flag-bearers Keeney and Raiffa [7]. Their formalization also opened the door to computational solution concepts for negotiation, such as reinforcement learning, and recent work by Yasser Mohammad and Amy Greenwald shows us the way here [8], [9].
- You may wonder how we can ever make true progress with an indifferent, mathematical view on a topic as emotionally charged as negotiation. Well, there is a lot of important research on this aspect too, supplementing what I present here [10]. The machine is still sorely lacking in this area, so for the foreseeable future, we can safely assume that the human still tends to the relationship, in synergy with the machine, which is better at calculating.
- And finally, I can sense the question in the room, so let me rip off the band-aid: can't we throw this all in the trash? Surely, we can now just ask an LLM to do the negotiation for us?

Let me answer that with a test case we ran a while ago, with two AIs negotiating where to go for dinner. There were three valid outcomes: pizza, sushi, or fries. One preferred pizza; the other wanted sushi. And we hoped they would agree on fries, which was the reasonable middle ground. Instead of that, however, after a long-winded discussion, the bots happily agreed to opt for sushi…on pizza.

And this summarizes the state of LLM negotiation nicely: it is the first natural-language negotiation capable of inventing its own outcomes. Earlier this year, Michael Wooldridge – one of the fathers of negotiation agents – rightly said in his AAMAS 2026 keynote that we've underestimated the power of language in the agent community. We do not yet know how far their capabilities will go, but right now, it is very challenging to extract actionable contracts or meaningful insight about negotiation from these tools. On top of that, and I will come back to this later, LLMs need many guardrails to keep them in line, and not just gastronomically.

My work is at the intersection of these disciplines: I adopt the agent vision, combined with the mathematical focus of game theory, but viewed through the decision-theoretical lens. Let me show you how we can go about this.

Automated negotiation

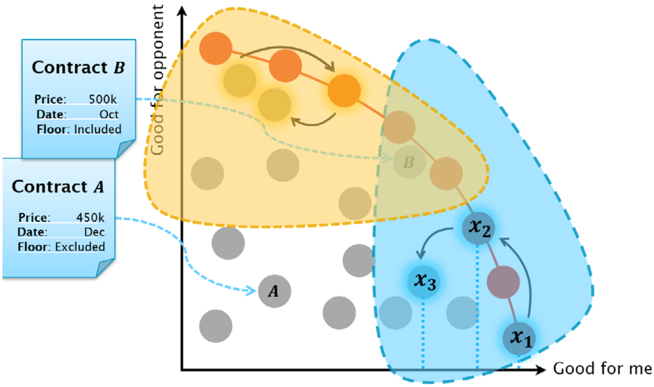

So, this is how we normally represent a negotiation, with our utility and the opponent utility on the axes. And every dot is an outcome that specifies values for multiple issues (for example, a real estate contract specifying price, delivery date, and so on). The desired, or Pareto-optimal, outcomes are in red in the top right; the Nash point is a special, fair Pareto-optimal outcome that maximizes the product of the utility of both players.

Note that it is no longer a line like in the housing example. Because we need to settle multiple issues, it has now become a game with private information and involves concession moves, exploring the space of agreements together.

The tricky part is: we have a private set of acceptable outcomes and so does the opponent. The intersection defines the *zone of agreement*, and one of these will be the outcome. Which one will it be? Well, that is determined by who is the best negotiator…or who has the best negotiation algorithm!

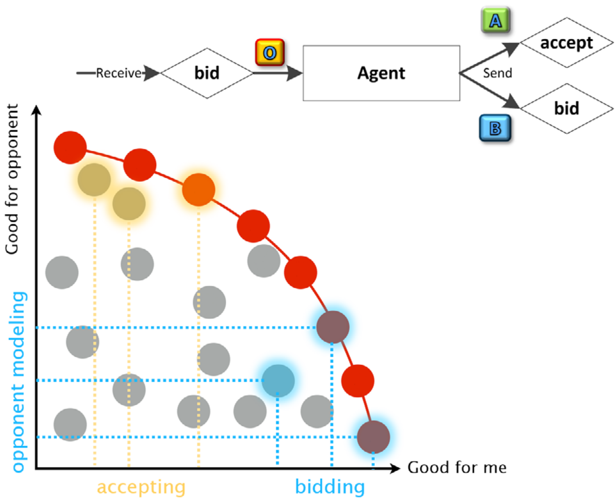

One of the pillars of the mathematical approach for a negotiation algorithm is a framework I have been working on since my PhD: the BOA framework [11]. With BOA, a negotiation agent decomposes into three small mathematical modules which can be studied in isolation. Little engines that can be put back together with a soft hum to form a larger whole.

- The B is for the Bidding strategy. Given the current state, which bids should be proposed? What concession is appropriate?
- The O is for Opponent model, which takes place on the vertical axis: how can we guess what the opponent prefers?
- And finally, A for Acceptance strategy. Someone has to accept to get an agreement, but when should we do that?

Over the years, BOA has become a textbook standard when designing negotiation agents, and has been extended with reinforcement learning by Yasser Mohammad [12], deep learning by Katsuhide Fujita [13], and LLMs by Reyhan Aydoğan [14] and Takayuki Ito [15]. This is a good time to stress that my work, like everything I present here, is joint work with my amazing students and collaborators. For the next part in particular, I thank Tamara Florijn, Reyhan Aydoğan, Yasser Mohammad, Enrico Gerding, Rafik Hadfi, Catholijn Jonker, Koen Hindriks, Takayuki Ito, and Pınar Yolum, many of whom I am happy to have here in the room with me. Thank you so much; I owe this podium to you.

But in its purest form, the 'B, O and A' induce three challenges, and if we zoom in on them, we discover elegant mathematical puzzles hiding in each one of them. So, let's take a deep dive, with your permission, into the bidding strategy.

Bidding

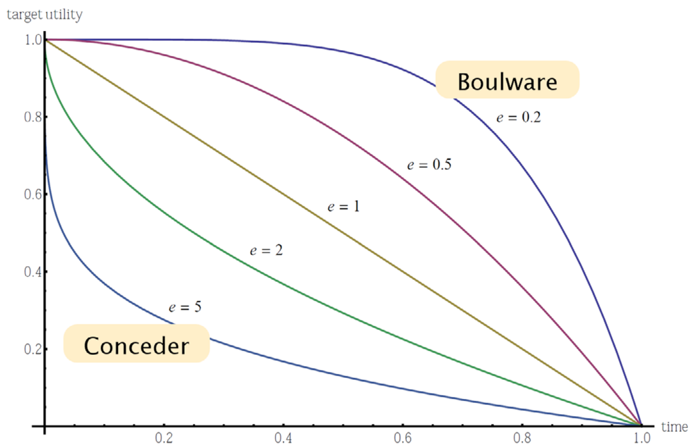

Let's start with the bidding. Many agents adjust their aspirations based on a concession curve to slowly concede over time. But what curve should we pick?

Some of the first simple, time-based negotiating agents were the time-dependent tactics proposed by Faratin and Sierra [16]. They concede over time, dependent on parameter *e*. For high values of *e*, we concede a lot in the beginning and tighten up toward the end. For low values of *e*, this leads to the Boulware strategy, which hardly concedes over time until the very end. We'll see more of Boulwarism later; it is a negotiation tactic named after General Electric's former vice president Lemuel Boulware [17], who pioneered the strategy in his negotiations with the workers' union. He would simply present his offer and stick to that position without making virtually any concessions. The shortcoming here is that whatever value we pick for *e*, this concession curve is selected rather arbitrarily and is not informed by any other fundamental insights.

So, how can we think about concessions more fundamentally? We have found the first 'optimal' curve – meaning the best way to concede, given that we have to fix, ahead of time, the number *k* of bids we will make [18].

Optimal bidding

Suppose through opponent modeling, we have some idea of the chance that the opponent will accept a certain bid x_i. Note: this is the decision-theoretical viewpoint I was talking about!

Now we need to formulate a sequence of bids x^* of length k. If we only need to make *one* bid, so k = 1, this is relatively easy: this is exactly the Nash point. But for k = 2, it already gets more complicated and, in principle, we'd have to consider every pair of offers.

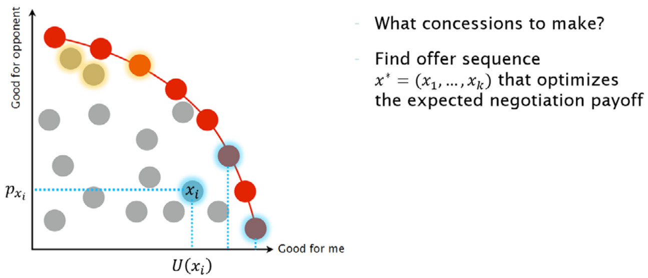

Acceptance probability of x_1

Utility of agreement x_1

Subtracted by the costs of making k bids

$$p_{x_1}U(x_1)+(1-p_{x_1})\underbrace{\left[p_{x_2}U(x_2)+(1-p_{x_2})\left[\cdots\right]\right]}_{\text{Otherwise, we try } x_2}-c(k).$$

More generally, we need to optimize the expected utility of sending the first bid or, when this is rejected, the expected utility of the second bid and so on, minus the costs. These can be constant or suddenly go up to infinity after the deadline of k bids.

Now I am going to claim that finding this sequence is not as hard as it seems. Using theory from simultaneous search [19], we can prove that we can build up our sequence incrementally.

	A	*B*	*C*
			
Utility	1	0.5	0.48
Acceptance probability	10%	90%	100%

- Applying for admission costs 0.05 per college.
- Question: Which colleges should we apply for?

Note
The answer is a subsequence of $\{A, B, C\}$ with the highest expected utility

I am going to illustrate this with something a little easier to visualize: suppose a student must make a costly and simultaneous application to several colleges. College A is a very prestigious college with the highest utility, but extremely hard to get into, while B and C are safer bets.

The question is: which colleges would you apply to?

The student cannot simply apply to all colleges because there are application costs. For instance, it might not even be worth applying to A! How can we decide what to do?

A logical start would be to just consider the immediate payoff first. And I'm going to claim that, in fact, if we keep doing that, this is good enough! That is: we can construct it in a greedy way, adding bids while there is a marginal benefit in doing so.

The first step is pretty straightforward, looking at the numbers and multiplying: if we could pick only one college, we should clearly pick College C.

So, we'll add it to our portfolio and we will not remove it again. We have now incurred 0.05 costs.

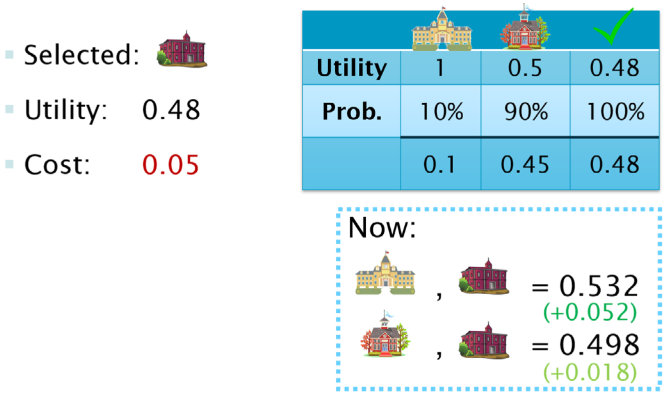

Now we can consider adding A or B to the mix. Note that we don't want to be admitted into C if we can get into A or B instead, so we should always apply to the best college first! So, if we consider adding A or B, we should put them in front of College C.

A simple calculation verifies that adding option A is far superior to adding B and also beats the cost of applying.

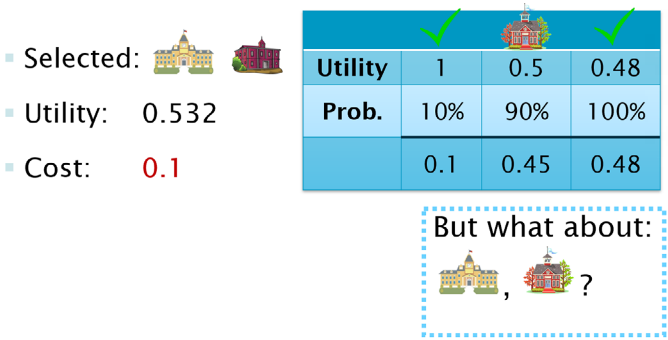

Therefore, we select A as well, creating a portfolio of two colleges at cost 0.1.

But wait! Shouldn't we also consider the last pair, A and B? This is a reasonable question, since maybe we need to take our initial choice of C out of our portfolio again.

Here is the kicker: we don't. Compare the two cases: both contain College A first, so there is no difference there. And if we don't get into College A, we are back in the case where we have only one application left. And we already determined that C is the best choice in that case!

And therefore, A and C must be the best two choices.

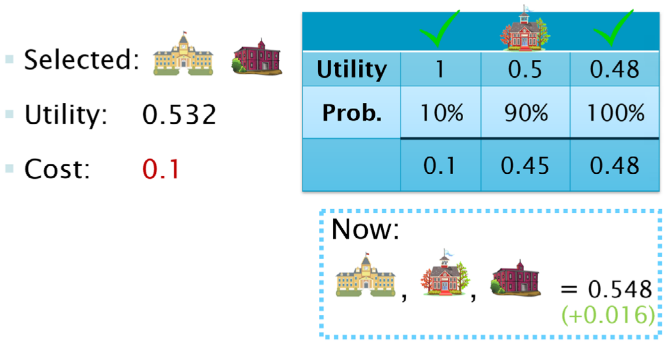

And we can verify easily that it is not worth it to add B to the mix because the costs are higher than the marginal gains.

So, we are done, and the solution to the puzzle is {*A*,*C*}. The general idea is to start with some safe options in our portfolio and then add riskier ones.

We can prove that this always works under certain conditions on the cost function that typically hold in negotiation.

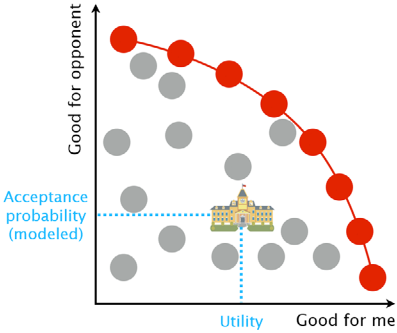

Now, bidding in negotiation is simply an enormous college problem. We can view all of our outcomes as colleges, where the costs can be constant or deadline-induced. The number of possible bids is huge, of course, but because we never need to backtrack, this is actually feasible.

Then we can apply the greedy concession algorithm [20] to find the curve we were looking for.

I am aware of the irony that I have omitted the most beautiful part, namely the actual proof – which is in the paper [20] – while at the same time trying to convince you of the beauty and universal use of math in negotiation. But I believe you might be more interested in the intuition and what the solution looks like.

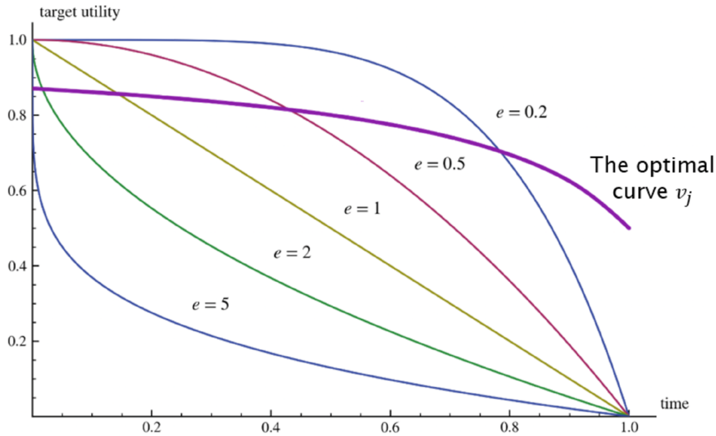

When plotted against time, the results show that the optimal curve is a highly risky Boulware-like strategy.

And anyone working with Boulware strategies in their simulations or the Automated Negotiating Agents Competition (ANAC) [21] knows this is actually a pretty good strategy. But unlike Boulware, this strategy follows straight from theory; as far as we know, this was the first concession curve we got this way.

We now understand why Lemuel Boulware had it roughly right, and this is my take-home bidding tip: it is a mathematical necessity to only concede quickly toward the end. This is how we should generally approach any negotiation: with only a few bids left, we need to work hard to get to an agreement. And conversely, with much time, we are allowed to take risks for a while, much higher than you might expect, by insisting on a very good outcome.

Bidding on a house

And if you're not convinced by the math, maybe it is worth knowing that I trusted this algorithm so much that I bought my own house using a variant of this bidding strategy.

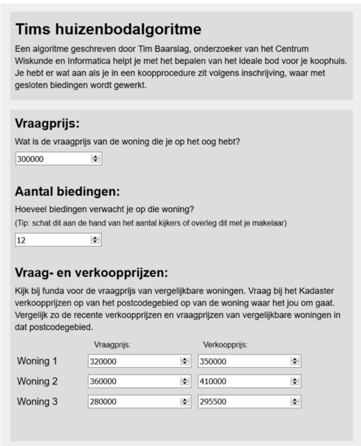

http://go.ntr.nl/tims-algoritme

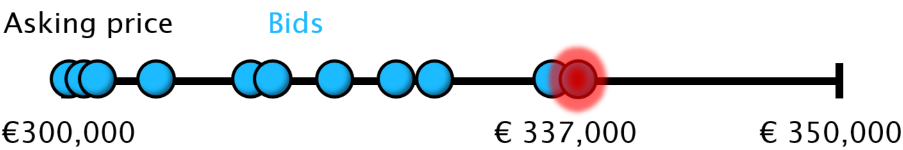

Some of the calculations I did to make a successful bid turned into an app, which I made together with science TV show Focus, where you can calculate the optimal bid on a house, given some input like the asking price and how sure you want to be to get it.

The idea is that it tries to model bids that will be made by others. It then calculates a bid that is as low as possible, but still higher than all others!

It's quite a simple calculator, but it gave us some insights. For instance, security costs an extraordinary amount of money: the last percentages to be sure to get a house are by far the costliest.

Looking to the future, I think entrants to the housing market who rely on data and mathematics rather than on 'feel' and intuition will have a huge advantage in the coming years of digital transition. Right now, it's already possible to bid online on some houses on Funda, so computerized support will not be far away. We are also looking at the inverse problem: how should you set the asking price as the seller?

We're working together with real estate agents to learn from their data, and we will soon be able to find out whether an AI can do as well as, or perhaps even better than, their human counterpart.

So, how did the AI do for my house? It did make the highest bid – though for all I knew, I had overbid enormously.

This is extremely anecdotal evidence, of course, but it is very rare you get the full picture. My broker was not really supposed to tell me, but they were nice enough to inform me anyway – it turned out the bid was only 1,500 euros above the number two.

Opponent modeling

Leo Hendrik Baekeland, born in 1863, was always at the head of his class. He graduated at just sixteen and he received his doctor's degree maxima cum laude when he was still only twenty-one. By 1891, he had opened an office in the US as an independent consultant and invented a special type of photographic paper, Velox, that could be developed under artificial light, but he was still struggling with his manufacturing business.

One day he received an invitation letter from George Eastman, who had established the Eastman Kodak Company (a name you might recognize) in Rochester, New York. Eastman suggested that if Baekeland was willing to sell his Velox manufacturing company, he was welcome to visit him for a talk. During the long carriage ride up to Rochester, Baekeland planned to ask for $50,000 but kept wondering if he would be able to get even $25,000 for his manufacturing process. As the story is usually told, Eastman invited Baekeland into his office and, fortunately for Baekeland, Eastman spoke first! He right away offered him one million dollars. Baekeland immediately took the offer. He could now afford to do his research in a well-equipped laboratory and went on to invent the first plastic, Bakelite.

What Eastman failed to do is to engage in proper opponent modeling. What does the opponent want, and can we learn from the opponent's bids? It's the second challenge I want to talk to you about, taking place on the vertical axis.

For this, we can use methods from machine learning and pay attention to the bids and trade-offs. For instance, if the opponent really wants something, they will probably stick to that issue.

But I will be brief about this one because I have a message that might lead to some frowning in the front row here: we need way less research here. And the reason is simply that doing a reasonable job here is good enough.

In 2016, we surveyed the opponent-modeling methods proposed up to that point [22] and we found something interesting: you can simply count. The rule is so simple, it is routinely rediscovered by master's students and can be used in everyday life. When a particular value is offered often, it is probably preferred. And if it does not change very often, then it is probably important. We have shown that this simple rule does surprisingly well, and even better than more sophisticated methods such as Bayesian learning.

In fact, the situation is even more extreme. Quite a few years ago, we conducted a sort of meta-analysis on the influence that each BOA component has on the final outcome [23]. It turns out the influence of the opponent model is only 4% – way less than the bidding strategy. Don't get me wrong: I'm not saying opponent modeling isn't important, and the interaction effects with the other components are quite large, but this particular metric (called the effect size) tells us that the benefit of doing something sophisticated is entirely marginal. In fact, even momentarily cheating and knowing the opponent's preferences does not do much better [23]. In other words: while it may be fun to build, we really do not need another machine learning technique to do opponent modeling reasonably well, and we should expend our effort on the actual bidding and accepting instead.

However, in keeping with the theme of this talk, I do see a valuable way forward in this area. And that is to understand mathematically why certain methods work well. And this is something we have recently started working on here in Eindhoven. For instance, under certain circumstances, which are still unknown, we can show that this simple counting method is, in fact, guaranteed to converge to the right preferences! Another interesting avenue is the converse problem: how do we stop the opponent from inferring too much about our own preferences? The trick we found there is paltering [24], which means telling the truth but very selectively, to the point where you say very little or even nothing at all. We found, using linear programming techniques, that we can construct long sequences of bids where we can concede, but with upper bounds on how much information we reveal. I think this will be very important in the future, with long, drawn out automated negotiations where bots can exchange thousands of bids.

When to accept

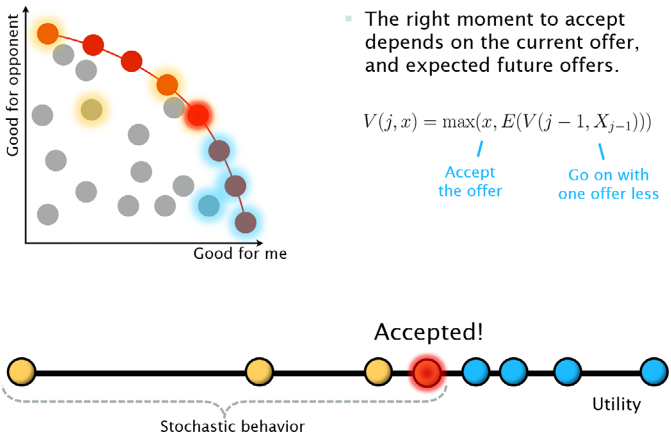

If the opponent is also doing proper opponent modeling of us, then we might be presented with some enticing offers as well. How do we decide whether it is time to accept?

This is a hard problem because the decision to accept depends not just on the current offer but also on the expectation of what offers might present themselves in the future.

For instance, suppose we open with a good bid for ourselves and the opponent starts in their corner. We make a small concession; the opponent makes a big one. So far, so good! We continue and then something funny happens: we get a worse bid! We can see what happened: the opponent gave up utility but failed to give us more because the opponent does not know our preferences perfectly! Therefore, we can model their behavior as stochastic in our utility.

This can be seen as a sequential decision-making problem of when to stop and say yes, I want it. This leads to a recurrence relation in which the value of having j bids left and current offer x is then simply either x when we accept it, or the expected value of doing essentially the same thing but with one round less. So, how do you choose the right time to stop?

For the famous astronomer Johannes Kepler, this was much more than an abstract mathematical problem. In 1613, Kepler wrote a long letter to his friend Strahlendorf in which he describes a great problem that he faced [25], [26]. Kepler had lost his wife and set about finding a new wife through a series of interviews among eleven candidates. In order not to hurt the feelings of his potential wives, he had to see them one at a time, and decide whether to propose before he was allowed to move on to the next.

He spent two years trying to find his ideal woman and kept somewhat troubling notes of his small disappointments as he wooed: the first candidate, he wrote, had “stinking breath”, so she was out. The second one had been “brought up in luxury that was above her station.” It got worse for him – the third one turned out to be engaged already! Luckily, No. 4 and No. 5 were good choices, but could he really stop so soon already, with six candidates unseen? In fact, he deliberated for so long that both got impatient and took themselves out of the running.

His problem was essentially that of any eligible bachelor: how do you decide when to settle and stop looking for other candidates?

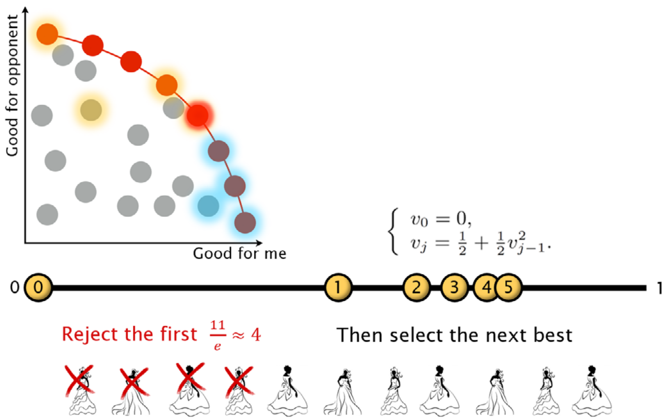

Thankfully, Kepler's problem is a bit antiquated in modern times and is now better known as an instance of an optimal stopping problem [26]: a stochastic decision problem of determining whether to accept among offers appearing sequentially and randomly. You can show that in the simplest case, you have to blindly decline the first 37% (or, more precisely, 1/*e*) of all candidates and then accept the first one who is better than everyone you have already seen. That is the optimal thing to do.

And if you look at these thresholds with varying number of candidates, we see something surprising: we find exactly the same acceptance threshold graph v_j as the optimal concession curve we have seen before! There is still some mystery regarding this curve and how it pops up everywhere. We found it in optimal bidding too [18], and de Jonge has shown it even pops up as an equilibrium strategy [27]. This underlines another great strength of the mathematical approach: we don't need to exchange our Python code or synchronize our LLM version; we speak the same universal language and seem to be uncovering parts of something bigger.

There are many avenues for moving forward from here. In practice, of course, you don't always know the number of options in advance, but it is also possible to extend the theory to non-fixed numbers, population learning, or recall of previous options.

And trade of goods and procurement is another example that we are working on currently, where you have to choose wisely between offers from different suppliers. In all those cases, you get multiple offers from different parties at the same time.

In terms of optimal stopping, that creates multiple lines of marriage candidates, and this is not equivalent to one larger queue because the current candidate in one queue starts acting as the backup option of the other one! How to solve this is still an open problem I am working on myself right now.

So, that concludes our third and final challenge, the agent's acceptance decision! I will end by zooming out a bit to sketch what it all means and how we can think more broadly about the implications for Cooperative AI.

But before I do: you might be wondering, what happened to Kepler in the end?

Well, Kepler was so crippled by indecision that after two years, he had gone through all his candidates, totally wooed-out. He decided that maybe he'd done this all wrong. "Was it Divine Providence or my own moral guilt," he wrote, "which, for two years or longer, tore me in so many different directions and made me consider the possibility of such different unions?" After a period of reflection, Johannes Kepler came crawling back to the fifth woman, who he described as "modest, thrifty, and diligent", and proposed to her, and she said…yes!

In fact, he could have saved himself all that heartache had he followed the cold, mathematical rules of optimal stopping. With eleven candidates, the rule says to look at the first four ($11/e \approx 4$) and then take the first who beats them all. She was the best candidate at that point, so choosing the fifth candidate is exactly what you would expect from optimal stopping theory.

And it turned out very well for Kepler: he became extremely happy in his second marriage – much happier, in fact, than in his first.

Mathematics of Cooperative AI

Alright, now let me zoom out to argue that these kinds of topics will soon not just be matters of the human heart. We're fast approaching a future where AI will be doing most of the heavy lifting, especially in domains where people are just not fast or efficient enough. So, AI systems won't just help us; they will represent us while interacting and cooperating with each other. And to become truly Cooperative AI, they will need negotiation skills in their toolbox to make joint decisions, just like we do. This is the broader scope of what I have actively worked on with the Intelligent and Autonomous Systems Group that I lead at CWI:

- For instance, the smart grid is fast approaching, in which computers have to negotiate complex energy contracts in milliseconds.
- Or think of autonomous vehicles. Traffic lights will slowly disappear and instead, vehicles will need to decide jointly on who gets priority on an intersection.
- Another example is the rise of the Internet of Things. Here, our smartphone needs to decide autonomously on the usage of our sensitive data with our apps and digital services, making countless trade-offs between price, privacy concerns, and convenience.
- Or the trading of goods and services, which is the topic of my current Vidi research: this could be supply chain management to real estate negotiations to kidney exchange.

The overarching question here is how we can teach computers to represent us properly in these negotiations [28].

To do so effectively, the agent needs to know more about us by asking questions, but not too many and only at the right time.

And therefore, we have to accept that Cooperative AI will need to perform these tasks with very limited preference information, where the negotiation space becomes one big blur!

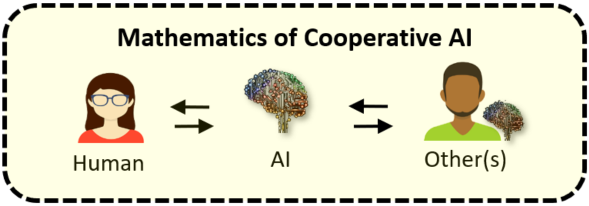

So, this is the challenge I am taking on with my chair here in Eindhoven: the *Mathematics of Cooperative AI*.

The key challenge is to formulate new mathematical theories to represent humans properly while dealing with AIs or humans on the other side.

It is my long-term ambition to crack this, and my team and I will be working on formulating mathematical models of elicitation and the fundaments of AI interaction under uncertainty in the upcoming years.

Future of Cooperative AI

- One product: $\arg\max_{i} p_{x_i} \cdot U(x_i)$
- Two products: $\arg\max_{i,j} p_{x_i} \cdot U(x_i) + (1 - p_{x_i})\, p_{x_j} U(x_j)$

Finally, I ask your permission to speculate a little bit from my brand-new professorial armchair. I'd like to convince you that the type of Cooperative AI I have just outlined will take great flight in the coming years.

To do so, I ask you to consider the following: why does one shampoo cost three times as much as the other? Surely both get the basics right, and Shampoo A cannot possibly be three times better.

As any economist will tell you: this is a simple form of second-degree price discrimination. When we can put only one shampoo on sale, we are confronted with a difficult trade-off: the luxury option A will scare off the budget-conscious, while option B does not extract enough value from buyers with disposable income. If only we could have a dynamic price tag per customer! Offering multiple – yet marginally different – products achieves just this: it lets customers self-select their willingness to pay.

So far, so good, but here is the twist: these formulas should by now seem familiar; this is just a simple form of optimal bidding by the seller, with only two colleges! In this light, the luxury option can be seen as the first, high-risk/high-reward attempt, and the budget option as the safe fallback offer.

And where the old trick of price discrimination is limited by real-life constraints such as customer fatigue, inventory capacity, and marketing budgets, negotiated products can vary endlessly and, with automation, effortlessly.

With the advent of LLMs, true agentic AI is at our doorstep. Agents can now browse the human web, and I predict that websites will cater more and more to agentic visitors. Rather than painstakingly shopping for products and services ourselves, we will have agents procure them on our behalf. They will ask for the best deal from multiple providers and fine-tune their offerings based on some of the negotiation principles I set out in this talk.

Back in 2000, Amazon's early brush with dynamic pricing provoked a lot of backlash, because buying the same product for a different price is widely perceived as unfair. However, I predict extremely personalized products in the future, where second-degree price discrimination collapses, at last, into first-degree: near-perfect and personal. You get exactly the color, warranty, and so on, that you prefer, as negotiated by your agent. This will be an almost unique offer, with no envy-inducing reference point to others, and it will only be possible because of automated negotiation in the background.

You might wonder what will happen then if everyone starts using these negotiation bots. We will have come full circle in that case, with entirely rational players, where game theory provides the right framework for studying their outcomes. But until that day, fairness will be a big sticking point: how do we prevent it simply being that whoever has the best bidding bot prevails?

A more subtle danger is implicit collusion among agents [29], where they may simply learn to rig prices as an emergent feature, without ever having been told to explicitly. We will need regulations and ethical guidelines to retain a fair and transparent market in this new time of agentic negotiation.

Finally, looking even further forward, it is worth noting that people negotiate very differently through intermediaries (for instance, lawyers) than they would face-to-face: they show less regard for fairness and ethical behavior [28]. We need far more research to understand when and how such effects arise when we are being represented by AI, which I consider a fundamental part of my chair.

Thanks

It is time to close off. I am incredibly grateful for the honor of becoming a professor and would like to thank Frits and Edwin for their faith in me. Thank you, Combinatorial Optimisation and the SPOR members, for welcoming me so warmly. Thank you to everyone in my group at CWI, and my direct colleagues in particular: Tamara, Loes, Brian, Valentin, Hendrik, Jurjen. Thank you Eric, for your trust in me to take over the reins as group leader. And my thanks also go to: Peter, Sander, Bert, Daan, Lynda, Jos, Ton, and Vanessa for your endless support.

I'm also indebted to the lovely people at Utrecht University, my alma mater, for teaching me so much about academia! Thank you, Pınar, Mehdi, Marc, and the Intelligent Systems group.

And it is during my PhD in Delft I finally became that scientist, and don't think I would've made it without you, Catholijn. And thank you Koen, Pradeep, Mathijs, Joost, Reyhan, for making Delft such a fond memory.

I salute my dear Jonge Akademie members for educating me, for coming all the way here, and for giving me a way to call myself Jong and be technically correct.

And thank you to all my other students and collaborators, without whom this simply could not have been done. I am lucky to have worked with so many of you. Ik heb gezegd.

References

[1] J. F. Nash, "The bargaining problem," Econometrica, vol. 18, no. 2, pp. 155–162, 1950. DOI: 10.2307/1907266.
[2] J. von Neumann and O. Morgenstern, Theory of Games and Economic Behavior. Princeton, NJ: Princeton University Press, 1944.
[3] A. Rubinstein, "Perfect equilibrium in a bargaining model," Econometrica, vol. 50, no. 1, pp. 97–109, 1982. DOI: 10.2307/1912531.
[4] J. S. Rosenschein and G. Zlotkin, Rules of Encounter: Designing Conventions for Automated Negotiation among Computers. Cambridge, MA: MIT Press, 1994.
[5] M. Wooldridge, An Introduction to MultiAgent Systems, 2nd ed. Chichester: John Wiley & Sons, 2009.
[6] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. J. Wooldridge, and C. Sierra, "Automated negotiation: prospects, methods and challenges," Group Decision and Negotiation, vol. 10, no. 2, pp. 199–215, 2001. DOI: 10.1023/A:1008746126376.
[7] R. L. Keeney and H. Raiffa, Decisions with Multiple Objectives: Preferences and Value Tradeoffs. Cambridge: Cambridge University Press, 1993. (Originally New York: John Wiley & Sons, 1976.)
[8] Y. Mohammad, S. Nakadai, and A. Greenwald, "NegMAS: a platform for situated negotiations," in Recent Advances in Agent-based Negotiation, Studies in Computational Intelligence vol. 958, Singapore: Springer, 2021, pp. 57–75. DOI: 10.1007/978-981-16-0471-3_4.
[9] Y. Mohammad, S. Nakadai, and A. Greenwald, "Automated negotiation in supply chains: a generalist environment for RL/MARL research," in PRIMA 2024: Principles and Practice of Multi-Agent Systems, LNAI vol. 15395, Cham: Springer, 2025, pp. 19–24. DOI: 10.1007/978-3-031-77367-9_2.
[10] J. Gratch, D. DeVault, G. M. Lucas, and S. Marsella, "Negotiation as a challenge problem for virtual humans," in Intelligent Virtual Agents (IVA 2015), LNCS vol. 9238, Cham: Springer, 2015, pp. 201–215. DOI: 10.1007/978-3-319-21996-7_21.

[11] T. Baarslag, K. Hindriks, M. Hendrikx, A. Dirkzwager, and C. Jonker, "Decoupling negotiating agents to explore the space of negotiation strategies," in Novel Insights in Agent-based Complex Automated Negotiation, Studies in Computational Intelligence vol. 535, Tokyo: Springer, 2014, pp. 61–83. DOI: 10.1007/978-4-431-54758-7_4.

[12] A. Sengupta, Y. Mohammad, and S. Nakadai, "An autonomous negotiating agent framework with reinforcement learning based strategies and adaptive strategy switching mechanism," in Proc. 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2021, pp. 1163–1172.

[13] R. Higa, K. Fujita, T. Takahashi, T. Shimizu, and S. Nakadai, "Reward-based negotiating agent strategies," in Proc. 37th AAAI Conference on Artificial Intelligence, vol. 37, no. 10, 2023, pp. 11569–11577. DOI: 10.1609/aaai.v37i10.26367.

[14] E. Kuru, A. Doğru, M. Doğan, and R. Aydoğan, "Evaluating theory-of-mind in large language models through opponent modeling," in Proc. 25th ACM International Conference on Intelligent Virtual Agents (IVA), Berlin, 2025, pp. 1–9. DOI: 10.1145/3717511.3747081.

[15] Y. Konishi, K. Yamamoto, E. Sonomoto, R. Takeda, R. Furukawa, Y. Muraki, T. Shimizu, K. Fukumura, Y. Kanemoto, T. Ito, and S. Ding, "Preference estimation via opponent modeling in multi-agent negotiation," in Findings of the Association for Computational Linguistics: ACL 2026, pp. 28049–28058.

[16] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," Robotics and Autonomous Systems, vol. 24, no. 3–4, pp. 159–182, 1998. DOI: 10.1016/S0921-8890(98)00029-3.

[17] C. J. Gitomer, "Boulwarism: the philosophy and method, from 1949 to 1960," B.S. thesis, Dept. of Economics and Social Science, Massachusetts Institute of Technology, Cambridge, MA, 1965. [Online]. Available: https://hdl.handle.net/1721.1/70892

[18] T. Baarslag, R. Hadfi, K. Hindriks, T. Ito, and C. Jonker, "Optimal non-adaptive concession strategies with incomplete information," in Recent Advances in Agent-based Complex Automated Negotiation, Studies in Computational Intelligence vol. 638, Cham: Springer, 2016, pp. 39–54. DOI: 10.1007/978-3-319-30307-9_3.

[19] H. Chade and L. Smith, "Simultaneous search," Econometrica, vol. 74, no. 5, pp. 1293–1307, 2006. DOI: 10.1111/j.1468-0262.2006.00705.x.

[20] T. Baarslag, E. H. Gerding, R. Aydoğan, and m.c. schraefel, "Optimal negotiation decision functions in time-sensitive domains," in Proc. 2015 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), vol. 2, Singapore, 2015, pp. 190–197. DOI: 10.1109/WI-IAT.2015.161.
[21] T. Baarslag, R. Aydoğan, K. V. Hindriks, K. Fujita, T. Ito, and C. M. Jonker, "The Automated Negotiating Agents Competition, 2010–2015," AI Magazine, vol. 36, no. 4, pp. 115–118, 2015. DOI: 10.1609/aimag.v36i4.2609.
[22] T. Baarslag, M. J. C. Hendrikx, K. V. Hindriks, and C. M. Jonker, "Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques," Autonomous Agents and Multi-Agent Systems, vol. 30, no. 5, pp. 849–898, 2016. DOI: 10.1007/s10458-015-9309-1.
[23] T. Baarslag, A. Dirkzwager, K. V. Hindriks, and C. M. Jonker, "The significance of bidding, accepting and opponent modeling in automated negotiation," in Proc. 21st European Conference on Artificial Intelligence (ECAI 2014), Frontiers in Artificial Intelligence and Applications vol. 263, Amsterdam: IOS Press, 2014, pp. 27–32. DOI: 10.3233/978-1-61499-419-0-27.
[24] T. Rogers, R. Zeckhauser, F. Gino, M. I. Norton, and M. E. Schweitzer, "Artful paltering: the risks and rewards of using truthful statements to mislead others," Journal of Personality and Social Psychology, vol. 112, no. 3, pp. 456–473, 2017. DOI: 10.1037/pspi0000081.
[25] P. R. Freeman, "The secretary problem and its extensions: a review," International Statistical Review / Revue Internationale de Statistique, vol. 51, no. 2, pp. 189–206, 1983. DOI: 10.2307/1402748.
[26] T. S. Ferguson, "Who solved the secretary problem?," Statistical Science, vol. 4, no. 3, pp. 282–289, 1989. DOI: 10.1214/ss/1177012493.
[27] D. de Jonge, "Theoretical properties of the MiCRO negotiation strategy," Autonomous Agents and Multi-Agent Systems, vol. 38, art. 46, 2024. DOI: 10.1007/s10458-024-09678-1.
[28] T. Baarslag, M. Kaisers, E. H. Gerding, C. M. Jonker, and J. Gratch, "When will negotiation agents be able to represent us? The challenges and opportunities for autonomous negotiators," in Proc. 26th International Joint Conference on Artificial Intelligence (IJCAI), Melbourne, 2017, pp. 4684–4690. DOI: 10.24963/ijcai.2017/653.
[29] E. Calvano, G. Calzolari, V. Denicolò, and S. Pastorello, "Artificial intelligence, algorithmic pricing, and collusion," American Economic Review, vol. 110, no. 10, pp. 3267–3297, 2020. DOI: 10.1257/aer.20190623.

Curriculum Vitae

Prof.dr. Tim Baarslag was appointed part-time professor of Mathematics of Cooperative AI at the Department of Mathematics and Computer Science at Eindhoven University of Technology (TU/e) on June 1, 2024.

Tim Baarslag is a Professor of Mathematics of Cooperative AI at Eindhoven University of Technology and a Senior Researcher and leader of the Intelligent and Autonomous Systems group at Centrum Wiskunde & Informatica.
He graduated from Utrecht University with an MSc in Mathematics and obtained his PhD from Delft University of Technology in 2014. He then became a Research Fellow at the University of Southampton. In 2016, he joined CWI as a Researcher, while working as an Assistant Professor (2018, NWO Veni) and then Associate Professor (2022, NWO Vidi) at Utrecht University.
Prof. Baarslag has served as a member of De Jonge Akademie of the Royal Netherlands Academy of Arts and Sciences (KNAW). He is also a founding member of the Netherlands Academy of Engineering and the ACM Future of Computing Academy. His research is featured in Science, Artificial Intelligence, BBC Technology of Business, AI Magazine, MIT Technology Review, and New Scientist. He is a recipient of the ERCIM Cor Baayen Early Career Researcher Award and has been named an Academic Pioneer by Elsevier, Young Talent by Het Financieele Dagblad, and Science Talent by New Scientist.

Colophon

Production
Communication Expertise Center TU/e

Cover photography
Bart van Overbeeke Photography, Eindhoven

Design
Grefo Prepress, Eindhoven

Digital version:
research.tue.nl/en/

Visiting address
Building 1, Auditorium
Groene Loper, Eindhoven
The Netherlands

Navigation
De Zaale, Eindhoven

Postal address
PO Box 513
5600 MB Eindhoven
The Netherlands
Tel. +31 (0)40 247 9111
www.tue.nl/map

EINDHOVEN
UNIVERSITY OF
TECHNOLOGY