

AI In Sales: Where It Works, Where It Fails, And What's Next

Introduction

[0:00] Alec Pritzos: Seller resistance is a big part of this. Like, you can put all the tools out in the world, but if your sellers don't use those tools, it's not going to help. So initially we tried to roll out these agent use cases. It didn't go well.

Once we got adoption over 80%, opportunity conversion improved by about 15%. Revenue per seller, when adoption went up, was like 9-something percent. For us, it was pretty dramatic.

[0:43] Anthony E Byrne: Welcome to Growth Decoded, MarketStar's podcast for B2B sales leaders who want to cut through the noise and get into how sales organizations actually execute across different growth stages, shifting markets, and increasing complexity.

At MarketStar, we run sales programs for some of the world's leading technology companies. About 4,000 sales professionals are deployed across all our client programs around the world today.

Over the past couple of years, AI has gone from something we were testing on the edges to something that's now very central to how we operate. We've built it into our workflows, our enablement, and all our tools.

Our philosophy is what we like to call AI + AI, which is Artificial Intelligence plus Actual Intelligence. So, it's human steering, technology amplifying—not replacing the seller, but multiplying what the seller can do every day.

Which is exactly what today's podcast and conversation is all about. Because if you're a CRO or a sales leader in 2026,

you're past the "what should we do with AI?" question. In fact, your board has probably already answered that.

The real question now is: how do you actually make it work?

How do you get from a pilot all the way to full-scale implementation?

And where does the human still matter? More importantly, where does that human matter most?

There's actually no one better placed to answer that than our guest today.

Alec Pritzos is the Global Sales Director for AI Business Process at Microsoft, working across telecom, media, and gaming—some of the most complex, high-velocity B2B sales environments in the world.

He's not just selling AI. He's embedding it into how enterprise organizations try to turn it into real sales outcomes. He sees what works. He sees what stalls. And he sees what leaders get wrong. So Alec, it's great to have you today. Thanks for joining us.

[2:40] Alec: That's perfect. Thank you so much.

Section 1: Context & Clarity: Why AI in Sales Is at an Inflection Point

[2:48] Anthony: Alec, from where you sit at Microsoft, selling AI into complex enterprise environments, what's actually fueling urgency right now?

Is it board pressure? Is it the C-suite? Is it genuine operational need? Or is it something else? What are you seeing?

[3:01] Alec: Yeah. So when I think about urgency, there are a few different things that come to mind for me.

I think there's a lot of urgency from IT organizations today that are really wanting to funnel in AI use cases.

There was an event I was at in New York, and somebody quoted that business cases are actually outpacing IT's ability to deploy solutions, which kind of feels relevant here.

What I mean by that is there's a lot of competitive fear around not releasing agentic AI solutions, especially in the sales space.

We see a lot of organizations being fueled by either having a competitive advantage or not having a competitive advantage.

For what I do here at Microsoft, I support a lot of our telecom, media, and gaming customers, and they all have that same feeling of wanting to gain a competitive advantage.

One of the things we see as an ROI metric that sometimes fuels this urgency is that there have been studies showing that for every dollar invested in generative AI in the sales space, organizations can get up to a \$3.50 return.



Whenever we share stats like that, or customers come to us with stats like that, it resonates.

Not everyone is seeing that return. So, if they're not seeing it, the question becomes: why aren't they seeing it, and what do they need to do differently? If they see competitors getting those kinds of results, that really fuels urgency.

[4:48] Anthony: It's interesting.

And when you talk about returns, one of the most common words in everyone's meetings and conversations around AI is "pilot".

Section 2: Reality Check: Scalable Leverage vs. Systemic Risk

[4:55] Anthony: Everyone's running pilots. We hear about pilots everywhere. Where we don't seem to be seeing as much success is the pull-through into scale.

There's often a gap between running a pilot and identifying the ROI case early enough to justify broader implementation.

What's your perspective on where that "one dollar in, three-fifty out" actually comes from?

Where does it align from a sales program perspective or a sales workflow perspective?

[5:28] Alec: I think of a production-facing solution that you're actually testing.

Whenever I see organizations looking at a pilot, they're almost always starting small.

Let's say they have a population of 200 sellers. Their pilot might focus on 10 of those people to start. The thing that's important to remember is that the technical deployment cost is often the same.

The integration work, the implementation effort—whatever work you're doing to deploy those solutions—you're paying those costs whether it's for 10 people or 200 people.

Often, you also have a very similar runway to get that solution into production. So, when I see organizations spending significant time and effort on very small-scale pilots, that's usually a red flag.

Now, sometimes you can't avoid it. You don't want to put the business at risk. You want your sales teams to be set up well. If something is highly experimental, it makes sense to start small.

Where I've actually seen a lot of success is putting guardrails around what you call a pilot before you even get started.

Instead of asking, "How many people should we deploy this to?" ask, "What can we build in 10 days, and what value can we get from it?"

When I see IT teams operating that way, they're rapidly putting pilots and tests in

front of the business, getting immediate feedback, and learning quickly.

You're not keeping teams tied up on a pilot for four months only to find out it failed after a lengthy deployment process.

That's where I think we've seen our own business become much more disciplined over the last few quarters. We've become better at determining the appropriate pilot size.

Sometimes that decision has to be made based on the motion itself and the time-to-value expectations.

Any time we've deployed something successfully, we've had a very clear hypothesis around the impact we expected to see.

Then the next phase becomes determining the right pilot size and balancing that against how quickly you'll be able to gather meaningful data.

[7:52] Anthony: Right. And you make a great point. From a technology perspective, you're often investing the same resources to get a pilot up and running.

What you really want is a control group where you can isolate as many variables as possible while still having a large enough sample size to generate meaningful insights and make decisions quickly.

Alec, budgets are clearly moving very fast into AI.

Looking at the pilots and projects you've seen over the last one or two years, do you think that money has generally been well spent?

Or are we still in a phase where some failed investment is acceptable because we're still experimenting?

[8:38] Alec: You know, it's funny. If I rewind three years, I'd probably say no. There was a real sense of urgency to just get started. I'm sure you saw it too.

Generative AI arrived and the narrative became, "Everybody should be using it. Let's figure out where and how." The mindset was basically, "Just go do it. Turn it on and see what happens."

That's why you've seen reports from research firms showing limited return on investment and significant amounts of lost capital. When I look at how organizations were deploying AI early on, that doesn't surprise me.

But I think companies are getting much smarter now. I'm not seeing organizations rush to deploy something simply because they can.

I'm seeing organizations that know exactly what they want to accomplish and are looking for help on how to do it as effectively and as quickly as possible. More than anything, what I hear now is time to value.

That's what comes up in every conversation I have—whether it's sales, service, or any other function.

If I deploy a solution, when do I start seeing value?

That's different from an ROI conversation. An ROI discussion might span four years, depending on what you're deploying. Time to value is much more immediate.

Organizations are saying, "We deployed generative AI, but we didn't see value." Maybe it was in place for a year, and they still didn't see value relative to what they were paying for it.

So overall, I think AI investments are being made much more intelligently today. There's another topic that's coming up now that I think is helping organizations spend more effectively.

[10:31] Alec: Another thing that's interesting when you think about AI budgets in general is that you have a salaried employee, a sales employee or someone else in the organization, and then you have the token spend associated with that person.

Last week, Nvidia CEO Jensen Huang made a comment that really caught my attention. He said that he expects engineers earning \$500,000 a year to consume around \$250,000 in AI token costs.

In other words, he wants highly paid employees using AI at such a rate that the

organization is investing significantly in AI alongside that individual.

When I think about that, Nvidia is probably ahead of a lot of organizations in this space, but if you extrapolate that idea into sales, I can't help but think some version of that becomes our future.

There's going to be a value equation tied to how much you're spending on AI tooling once you've dialed it in and understand exactly what value you're getting from every token you purchase.

[11:50] Anthony: Alec, if we cut through the noise, which AI applications in sales are consistently delivering what you would consider real, measurable value right now?

[12:02] Alec: That's a big question. There are a few low-hanging-fruit use cases that I'll mention because I think they're often overlooked, but they're important.

The first is meeting intelligence. If you think about a B2B seller's day, they might start work at 8:00 a.m. and finish at 5:00 p.m., spending most of that time on calls.

One of the biggest early value propositions was simply being able to record and transcribe meetings.

That was huge. But initially, it only got you so far because you still had to figure out what to do with those transcripts.

You had all these notes, but the question became: now what?



What Microsoft started focusing on was how to leverage all that meeting and call information inside our system of record.

Let's say you have an opportunity with twelve people working on it, but only one person is responsible for updating the opportunity record.

The rest of the deal team still needs visibility into what's happening, but historically that created a lot of friction. People just accepted that friction as normal.

What we've done is layer AI into that process so opportunities can be updated automatically.

The system identifies exactly what information matters, who needs to see it, and how it should be surfaced.

That reduces administrative work and makes the process much more strategic.

So meeting intelligence is number one. The second area is account research and meeting preparation. We see roughly 95% of sales workflows beginning with research using AI.

If sellers don't have approved AI tools available to them, they're going to find their own.

They're going to use whatever tools they can access, whether the organization has visibility into them or not. That's why I think this statistic is important. If you're not providing AI tools to your sales teams, they're probably using them anyway.

The third area is pipeline analytics and deal-risk identification. Think about how much time organizations spend discussing pipeline health.

A seller reports to a manager, that manager reports upward, and eventually those conversations make their way to senior leadership.

A tremendous amount of time is spent analyzing pipeline and trying to determine which deals are healthy and which are at risk.

We now have AI agents that can evaluate those risks based on meeting intelligence, account activity, research signals, and other inputs.

That's how we're approaching it. Those are some of the agentic AI applications Microsoft uses internally and where we've seen meaningful value creation.

[14:59] Anthony: That's interesting, Alec. And honestly, we've seen something very similar at MarketStar.

One of the first areas where we saw meaningful value was in our ability to provide clients with much clearer insight into market sentiment.

You're absolutely right. The immediate advantage wasn't just recording calls. It was being able to summarize them, identify patterns across large datasets, and understand trends in ways that simply weren't possible before.



It allowed us to bring the voice of the market back to our clients in almost real time.

Instead of waiting for quarterly analysis cycles and manually crunching numbers, we could surface insights much faster.

That changed the quality of our business reviews. It enabled us to make better recommendations about where clients should pivot, where they should double down, and how they should adjust strategy.

As a result, it improved both our revenue performance and our value proposition. The way we work at MarketStar is that when we help clients unlock additional revenue, it often creates opportunities for them to invest further with us.

So it's interesting that you started with meeting intelligence because that's where we first saw not only an operational advantage, but a clear business value equation.

You mentioned friction earlier. That's a word that comes up frequently when discussing AI adoption, pilot programs, and scaling initiatives. You've talked about three categories where you're seeing success.

For the sales leaders listening today, what are the most common friction points you see?

Whether it's within organizations or between cross-functional teams, what

tends to slow progress when companies are trying to decide where to focus, how to prioritize, and how to move from strategy to execution?

What are the two or three biggest friction points people should expect and perhaps avoid?

[17:31] Alec: That's a really good question. There are two ways I think about friction.

One is the friction involved in using agentic AI itself. The other is the day-to-day business friction we're actually trying to solve.

The first thing that comes to mind is the concept of agentic AI. It's a major buzzword right now. One thing I always tell people is that I don't think there's a future where a single AI agent handles everything.

Where we see the most success is with multiple purpose-built agents working together.

Those agents may be coordinated by a human or by another orchestration layer, but specialization matters.

The second friction point is people. More specifically, ownership and adoption. Seller resistance is real.

When Microsoft started its own sales transformation journey, we learned this firsthand.



You can put all the tools in the world in front of your sellers, but if they don't use them, nothing changes.

Initially, we rolled out several agentic AI use cases, and frankly, it didn't go well. We had to spend time driving adoption.

Once adoption exceeded 80%, however, we saw opportunity conversion rates increase by approximately 15%. Revenue per seller increased by around 9.5%. For us, those results were dramatic and highly meaningful.

The third friction point goes back to the technology side. It's data architecture. Before you begin any major AI transformation effort, you need to understand where your data lives and how it's structured.

Most people immediately think about CRM systems, but it's much broader than that. It's about all of your organizational data.

If you want an AI agent to deliver meaningful value, can it access the information it needs?

The answer depends entirely on how your systems have been architected.

So those are the three major areas I think about:

A clear agentic AI strategy.

A strong data foundation.

And a plan for driving user adoption.

Without those elements, it's very difficult to create meaningful outcomes.

[20:06] Anthony: That's very interesting, Alec. My next question was going to be...

[20:13] Anthony: If you were advising a CRO tomorrow morning—first engagement, first meeting on the calendar—they've got budget, they've got C-suite backing, and they're ready to launch an AI initiative.

Where would you start?

Would the first conversation be around where the data resides, and what can realistically be connected from an agentic AI perspective?

Or would you take the opposite approach and focus on the fastest path to revenue impact?

In other words, would you recommend a practical, "get started quickly" approach first, or would you focus on building the data foundation for long-term scalability and future automation?

Alec: I would probably take a slightly different approach.

There are really three phases that I tend to bucket organizations into. There could be fifteen different phases, but for simplicity I think about three. They're essentially automation maturity phases.

Phase One: AI agents are helping process information and identify things like leads.



Phase Two: Sales reps are consuming information generated by AI agents.

Phase Three: Sales reps are actively managing AI agents that help them close deals.

Now, Phase Three is still fairly futuristic, but I think it's useful because it helps frame where an organization is today.

Once I understand which phase they're in, my answer changes. For organizations in Phase One or Phase Two, I actually don't start with technology.

At Microsoft, we do a lot of work with customers around what we call envisioning sessions. They're essentially structured thought exercises. A good analogy is something we've seen with AI coding tools.

When Claude Code introduced Planning Mode, it became extremely successful. Someone later asked the product manager what Planning Mode actually was.

The answer was surprisingly simple. It's essentially a prompt that says:

"Before you do any work, create a plan."

That's it. And I think the same principle applies here. When I look at where organizations struggle, it's often because they jump straight into implementation.

So if a CRO asked me where to begin, I'd bring the business team and the IT team into the same room. I'd facilitate a

thoughtful envisioning session around questions like:

- What problem are we actually trying to solve?
- What obstacles exist today?
- Who are the stakeholders?
- What risks haven't we considered yet?
- What business value are we expecting to create?

Too often, IT is trying to solve one problem while the business is focused on another.

When those groups aren't aligned, projects slow down dramatically. Having a shared vision before you start can save months, sometimes years, of effort.

So for organizations that are early in the journey, that's always my recommendation.

Now, if you're further along—say you're operating closer to Phase Three—then the conversation changes. At that point, I want to understand where you're already seeing value.

If you're not seeing value, why not?

Maybe the issue is architecture. Maybe it's data. Maybe it's organizational alignment.

But if you're more mature, the starting point is usually evaluating what's already working and expanding from there.

Section 3: The Human Factor: **Redefining the Seller's Role in the AI Era**

[24:00] Anthony: We've spent a lot of time talking about the agentic side of AI, but I think it's important to talk about the human side as well.

Let me ask you a question through two lenses: your own experience and Microsoft's broader perspective.

What do you believe will remain uniquely human in sales?

Not just today, but in the future—particularly in enterprise sales.

[24:24] Alec: Speaking personally—and honestly, this aligns with how Microsoft is building products as well—we don't build AI agents that operate entirely without human involvement.

There are certainly scenarios where an AI agent can execute a complete process from start to finish.

But when it encounters an exception, uncertainty, or risk, it escalates to a human.

That's intentional. When I think about what remains uniquely human, the first thing that comes to mind is trust.

If there's no human involved, will another organization truly trust the outcome?

Especially in enterprise sales, I don't see a future where a major company makes a significant purchase simply through an AI

agent on a website. Large enterprise relationships are built on trust.

The second thing is partnership. Many of Microsoft's enterprise customers aren't simply customers, they're strategic partners.

We're working together to achieve shared outcomes. That collaborative relationship matters. I don't think AI removes the need for partnership. If anything, it elevates it.

The third element is judgment. There will always be ambiguous situations where someone needs to make a judgment call.

AI can surface information. It can recommend actions. But there will always be situations where a human has to make the final decision.

And beyond that, there's empathy. There's understanding. There are nuances in relationships that are difficult to automate.

So for me, the uniquely human elements are trust, partnership, judgment, and empathy. But I'll turn the question back to you.

At MarketStar, what do you think remains uniquely human?

[26:32] Anthony: Enthusiasm. The ability to negotiate. The ability to build trust. When we talk to clients, very few of them are saying:

"Our goal is to eliminate human sellers entirely."



There are legitimate concerns around buyer experience. Most buyers don't want to feel like they've been handed off to a machine. The goal for us is efficiency.

How do we increase output?

How do we make sellers more effective?

But sales still requires human energy. It requires negotiation. It requires trust. And it requires enthusiasm.

At its core, sales is about connecting a customer's requirements to a solution in a compelling and authentic way. What excites me about AI is not replacing sellers.

It's helping great sellers become even more effective by reducing the administrative drag that consumes so much of their time today.

Section 4: Lessons from the Field: What Leading Revenue Organizations Are Learning

[27:39] Anthony: From either Microsoft's internal experience or from your work with customers, do you have a standout example where AI directly influenced a sales outcome?

It could be a specific deal. It could be pipeline acceleration. It could be a broader business impact.

Do you have a story that really illustrates what's possible?

[28:11] Alec: I do. It's not tied to one specific deal. It's more about the impact across a large number of deals.

At Microsoft, we have several AI agents that our field teams use every day. One of the most impactful has been a sales research agent.

When I talk about things like the 15% increase in opportunity conversion that we've seen, it's because of tools like this.

Historically, sellers spent an enormous amount of time gathering information. Some of it came from internal systems. Some of it came from external research. Some of it came from conversations with colleagues.

They'd piece everything together before making a decision.

For example, a seller might research a company online, review previous interactions, schedule meetings with internal stakeholders to gather additional context, and spend days assembling a complete picture of the account.

By the time they felt confident enough to take action, a week might have passed. So we asked ourselves:

Why not consolidate all of that into a single experience?

We built a sales-focused AI agent connected to what Microsoft calls Work IQ.



It's connected to emails, Teams chats, Dynamics CRM, and multiple other internal data sources. Behind the scenes, it's orchestrating several specialized agents.

Now, instead of spending days collecting information, a seller can ask one or two questions and receive a comprehensive response in minutes.

Something that previously required hours from multiple people over the course of a week can now be accomplished almost instantly.

And when you describe the use case, it sounds simple. The impact isn't always obvious until you see the numbers.

But that's exactly where the 15% improvement in opportunity conversion came from.

Today, we're so accustomed to asking an AI agent to retrieve information from five or six different systems that we almost take it for granted.

But when you think about the time that was previously spent gathering that information—and the interruptions it created across the organization—the value becomes very clear.

Our sellers spend less time searching. They spend less time pulling other people into meetings. And they spend more time selling.

That's a very tangible example of AI directly contributing to revenue growth.

[30:56] Anthony: Let me ask you something. We've talked a lot about pilots, what's working, and what you've seen scale.

But you must have some interesting examples of the weird, wacky, and wonderful.

Have you come across any unusual—but surprisingly practical—agentic AI workflows or use cases that stand out?

Anything come to mind?

Alec: A couple things come to mind.

We already talked about real-time competitive intelligence, but one example internally was an AI agent that focused exclusively on publicly available company and stock information.

Imagine a seller managing eight different accounts. This agent's only job was to gather everything you would want to know if you were considering investing in one of those companies.

It would surface financial information, market signals, public filings, analyst commentary—essentially all the publicly available intelligence on a company.

What surprised us was how sellers started using it. They began creating incredibly detailed account dossiers, sometimes ten pages long, that looked like something a management consulting firm would produce.



What's interesting is that, generally, when we deploy AI tools, one of our core principles is not to change the seller's workflow too much.

If you dramatically alter how people work, adoption usually suffers. But this use case was completely outside the normal workflow.

The sellers essentially invented a new workflow on their own because they found so much value in the information. That was one of the more surprising examples for me.

The second example sounds much simpler. Across our AI experiences, we provide basic feedback mechanisms—a thumbs-up and thumbs-down option on AI responses.

We launched an initiative to encourage people to use that feedback more consistently.

Initially, I assumed that if we saw a large number of thumbs-down responses, we'd discover flaws in the AI itself.

That's not what happened. What we learned was that most negative feedback wasn't caused by poor AI responses.

It was caused by poor data. Whenever we investigated a thumbs-down response, we'd trace it back to the source information.

Nine times out of ten, the problem was old data, inaccurate data, duplicate data, or information stored in the wrong place.

It was a classic case of garbage in, garbage out.

That feedback process helped us dramatically improve our data quality. We went from maintaining millions of content assets and articles down to hundreds of thousands because we were able to identify what was outdated, redundant, or unreliable. That ended up being one of the most valuable outcomes of the entire initiative.

Section 5: What's Next? Preparing Sales for the Next Phase of AI Adoption

[33:48] Anthony: Let's move into a more future-focused discussion.

As you look toward 2027 and beyond, we've seen a lot of companies trying to differentiate themselves through AI. We've seen leaders positioning themselves as AI-first or AI-native.

How far away are we from a point where AI literacy and AI adoption are simply expected?

When does being "good at AI" stop being a differentiator and start becoming table stakes?

Is that this year?

Next year?

Or did that already happen?

[34:43] Alec: That's a great question. And honestly, I think the challenge is that the goalposts keep moving.

If you think about things like the Turing Test and the conversations we were having four years ago about what AI would eventually become, then compare that to the tools we have today, the progress is remarkable.

If you took something like ChatGPT and showed it to someone four years ago without telling them what it was, they might genuinely think they were interacting with a person. Maybe a slightly awkward person, but a person nonetheless.

Would you agree?

Do you think the goalposts have shifted?

[35:19] Anthony: Completely.

I think the days of walking into the executive suite and getting attention simply because you're running an AI pilot are mostly gone.

At one point, being seen as innovative or forward-thinking around AI helped build your personal brand as a leader.

Today, executives care far less about what you're experimenting with. They care about outcomes.

What does it do?

What is it worth?

What business impact did it create?

I think we're seeing the maturation of the category.

The hype is fading. Senior leaders are becoming much more focused on value, speed to impact, and measurable business outcomes.

[36:27] Alec: I agree.

And I'd answer the question a little differently as well. I think there's a top one percent of AI super-users.

I work with people who are constantly building things with Microsoft tooling that aren't even official product features yet. We're sharing ideas, workflows, and experiments with each other all the time.

Then there's probably a broader top twenty percent. These are people who adopt new capabilities immediately and use them aggressively.

I can see a very clear difference within my field teams between the people who embrace these tools and the people who don't.

Sometimes that difference shows up directly in KPIs. Other times, they simply look like rock stars. I had one seller build his own prospecting orchestration agent.

It completely blew me away. Then he presented it to the rest of the team and it blew their minds too. The interesting part was that he'd been using it successfully for four months before anyone else knew about it.

During those four months, he simply looked like he was operating at a different level than everyone else. I think there are

a lot of examples like that happening right now.

People are adopting AI rapidly, becoming more effective, and outperforming their peers.

Sometimes we're not even measuring the advantage because it's showing up in countless small ways.

But it's visible. If I look across a team of fifteen sellers, I can immediately tell who is using AI effectively and who is still relying entirely on traditional approaches.

[38:16] Anthony: Alec, I've really enjoyed this conversation. To wrap things up, let's do a quick-fire round. I'm going to ask a few questions, and I want your immediate instinctive answer.

Ready?

[38:34] Anthony: What's the biggest myth about AI in sales?

Alec: That it's going to replace salespeople. I think it's going to replace tasks. I don't think it's going to replace relationships.

[38:52] Anthony: If AI could take one thing off every seller's plate tomorrow, what should it be?

Alec: Manually entering data into CRM systems. If your sellers are still manually updating CRM records, you're far behind where you should be.

[39:13] Anthony: One word to describe the future of sales in 2027.

Alec: Augmented.

[39:26] Anthony: One piece of advice for a CRO who feels overwhelmed by all the AI noise. They're hearing ideas from every direction—internal teams, competitors, consultants, and their CEO.

What's the one thing they should focus on?

[39:50] Alec: Pick one workflow. Measure the impact. Then scale from there.

It's the old saying: How do you eat an elephant?

One bite at a time. Choose one process. Define exactly how you'll measure success. Implement it. Understand the value you're creating. Get a win.

Then use that success to expand into the next workflow.

[40:10] Anthony: That's fantastic advice.

Alec, I've really enjoyed our conversation today. Thank you for taking the time to join us.

And if you're ever in Dublin, the invitation for a pint is permanently open. I look forward to catching up in person somewhere around the world. Thanks again for joining us.

[40:29] Alec: Thank you, Anthony.

I enjoyed it as well.

I appreciate the opportunity.

[END]