

10 Commandments of AI Deployments

Drawing on insights from 20 enterprise AI deployments.

INTRODUCTION

Most companies recognize that AI is going to completely change how they operate, but very few understand the actual steps required to make that happen – so the default response has been to buy AI tools and hope company transformation follows.

To date, buying AI tools is not nearly enough because there's a massive gap between automating a task with AI and redesigning an operation so that AI reliably and repeatedly produces a financial outcome. That very gap is what we've built a business around at Varick.

Over the past few years, we've helped public and private equity-backed companies transform their operations around this technology, and in the process, we've seen the same patterns emerge repeatedly of how and where companies actually create real value, where deployments stall, and which assumptions from the software era no longer hold true in the enterprise AI era.

That work forces us to think about AI in two ways at once. We have to be ambitious enough to imagine how this technology could fundamentally reshape a company's operations, while remaining practical enough to work through the much harder question of how work will actually get done differently inside that company, and how those changes will translate into a measurable financial return.

These are ten of the most important lessons we've learned through thousands of hours spent auditing workflows, designing agents, deploying them into production, and refining the operating model around them.

1. A competitive advantage in a world of abundant intelligence

We started Varick with the belief that deployment would become the moat in AI.

If every company can purchase access to the same models, then the intelligence itself cannot be the competitive advantage. Your competitors can buy the same technology and call the same APIs, but if one company uses that intelligence to redesign how a critical operation runs while another simply layers it on to the process it already has, the advantage goes to the company that deploys it more effectively.

That means the real opportunity is not simply adopting AI, but applying it in a way that reflects how your company creates value. The model may be broadly available, but the operating decisions around where to deploy it, how deeply to integrate it, and which parts

of the business to rebuild around it are specific to each company.

This isn't just true for AI, but every major technological shift separates the companies that use the new capability to change how they operate from the ones that treat it as an incremental upgrade. In the intelligence era, that separation will come from how creatively and effectively companies turn broadly available intelligence into a unique competitive advantage.

2. AI requires an operational rebuild

Most companies still operate around workflows that were designed for humans, which means the inefficiencies of human work are built directly into the way those workflows run.

Work gets broken into handoffs because no one person can own the full process. Tasks sit in queues because people cannot review everything immediately. Teams batch work because constant processing would require too much time, and gaps between systems remain because employees are expected to carry information from one place to another.

AI removes many of those constraints, because agents can work continuously, move across systems, hold much more context, and carry a process from one step to the next without waiting for a person to pick it back up. But companies will not capture the full value of that by placing agents into isolated parts of the workflow. If the process itself was designed around human limitations, automating a few steps only makes an inefficient structure move faster.

The real opportunity comes from rebuilding the operation around what AI can now do. That means decomposing the workflow, removing the handoffs and delays that only existed because humans were doing the work, and deciding where people still add judgment while agents take responsibility for the rest.

In the same way factories had to be redesigned around electric motors before electricity could transform industrial production, companies will have to redesign their operations around agents before AI can materially change how they run.

3. Employee adoption is the bottleneck

Most enterprise AI products only create value when employees choose to use them. If an employee doesn't use the tool during the work itself, the AI doesn't remove any labor from the process, which means the productivity gain never materializes.

That makes the financial return on AI dependent on whether people change the way they work, which is a poor dependency to build a deployment around.

Companies can buy the software, train employees, and make the tools available, but they can't guarantee that people will use them consistently. Employees already have established ways of completing their work, and introducing a new tool does not automatically mean it becomes part of how that work gets done.

When you start a company focused on deploying AI into the enterprise, you don't expect employee adoption to become one of the main constraints on ROI. You assume that if the technology works and creates a clear productivity gain, people will use it. What we've learnt over time through actual deployment is that the technology can perform exactly as intended and still fail to produce a meaningful return because the employees it depends on never use it enough.

The better answer is often to remove that dependency altogether by deploying AI in the background rather than asking employees to operate it in the foreground. Employees can continue working inside the systems they already understand, while agents complete the manual work behind the scenes and bring people in only when judgment, approval, or intervention is actually required.

When AI becomes part of how the work gets done rather than another tool employees have to remember to use, adoption stops determining whether the value shows up.

4. The biggest beneficiaries of this technology may not be tech companies

New technologies are usually built first for the people building them. Computers were initially built for mathematicians, the internet was initially used by researchers, and the same pattern is emerging in AI. By far the most productive use case for AI today is coding, which created the impression that the largest beneficiaries of the technology will be the companies with the biggest software engineering workforces.

But the companies adopting AI first are not necessarily the ones with the most to gain from it.

Technology companies already operate at relatively high margins, which means that even meaningful efficiency gains hardly improve financials. A software company running at a 30% margin will not have a meaningful economic transformation with a slight cost

reduction, but the the business running at a 3% margin will.

If they improve their margins just by 1%, their profit jumps by 33%. The same efficiency gain that looks immaterial to a software company can materially change the earnings profile of an operationally intensive business.

These companies also tend to have far more of the work that AI is well suited to remove. Their operations are filled with manual handoffs, repetitive coordination, fragmented systems, and large teams carrying processes that software has never been able to automate from beginning to end.

That means the demand for AI-driven efficiency is likely to be much greater in low-margin industries, even though they remain one of the least served parts of the enterprise AI market today.

5. AI transformation is a standardization exercise

When five employees perform the same job, they usually develop five different ways of doing it. Each person creates their own sequence, judgment calls, and rules for handling exceptions, and the company often never notices because people are flexible enough to work around the differences and still get the job done. Almost every workflow we've audited has revealed some version of this problem.

AI exposes it because an agent cannot reliably operate across five informal versions of the same process. It needs a clear sequence to follow, defined inputs, explicit rules, and a consistent way to handle exceptions, which means that before a company can automate the work, it first has to decide how the work should actually be done.

That forces the organization to collapse multiple versions of the same workflow into one canonical process, while separating the real business differences that need to remain from the accidental variation that developed over time.

This is by far one of the most underrated parts of an AI transformation. The company comes out of an AI transformation with agents running most of their manual work, and a cleaner operation, with fewer exceptions, and a much more disciplined way of working than the one it started with.

6. Companies will manage AI spend like capital, not software

Most companies manage AI spend in one of two ways. They either push employees to

use as much of it as possible, or they impose a cap once the bill gets too high.

Both approaches lead to inefficient spending.

Encouraging more usage assumes that every additional token creates value, when in reality a large amount of AI spend can produce little return. Capping usage creates the opposite problem, because it controls cost without considering the fact that the tokens being restricted could have produced a much larger financial return.

The better way to think about tokens is as a form of capital.

Every token from the same model may cost roughly the same, but the return it produces depends entirely on where it is deployed. The same million tokens can create substantial value when used to recover missed revenue, reduce operational risk, or remove labor from a critical workflow, while producing almost nothing when applied to work with little economic consequence.

That's the same reason companies don't allocate capital evenly across every part of the business. They instead route it to where they expect the highest return, measure what comes back, and continue investing where the returns justify it.

The best companies will eventually manage AI spend the same way. They will not ask how large their token budget should be, but instead, they'll ask where the next million tokens should be deployed to create the greatest financial return. And the advantage will not go to the companies that spend the most on AI, but to the ones that allocate intelligence most effectively.

7. The value of AI goes well beyond cost savings

Most companies build the business case for AI around cost savings because labor is the easiest thing to measure. You can count the hours a process takes, estimate how much work AI can remove, and turn that into a clear financial return.

But that only captures part of the value.

When we audit an operation, we look at every workflow through three lenses: cost, revenue, and risk. A manual process may be expensive to run, but it may also delay billing, leave earned revenue uncollected, or allow errors to pass through because no one has time to check every transaction.

AI changes the economics of those problems. Agents can perform the follow-ups, reconciliations, and controls that were previously too expensive to run consistently. That means companies can capture more of what they earn and lose less of what they already have.

Companies that evaluate AI only through headcount reduction are measuring the easiest part of the return, not necessarily the most valuable one.

8. Enterprise AI is the era of software as an outcome

Enterprise software gave us software as a service: companies bought tools that helped their employees get work done. AI is starting to change that relationship because buyers increasingly do not want another interface for their employees to operate. They want the work itself to be done.

That may sound like a small change, but it alters both how enterprise AI is bought and how it has to be sold. Traditional software contracts and SLAs were built around access, uptime, and performance because the vendor was providing a tool while the customer remained responsible for the outcome. Once the vendor is expected to produce the outcome, those same agreements are no longer enough, because the buyer is no longer paying for an interface. They are paying for completed work.

This is also why services businesses have become one of the biggest beneficiaries of the AI era. A company cannot produce reliable outcomes by dropping a generic model into a workflow and walking away. The AI has to understand how that company actually operates: its systems, data, rules, approvals, exceptions, and all the small details that make the process unique.

Getting to that level requires going deep inside the client's business and deploying AI around the real problem rather than selling them another piece of software, so therefore, the age of enterprise AI will be defined by systems that take responsibility for getting the work done.

9. The next bottleneck

AI is well on its way to solving more of the execution of work. Model capabilities are improving quickly, and people can already hand off a much larger share of what they do than they could even a year ago.

As AI takes over more of the execution, the human role moves upstream. The important

question becomes less about who completes each task and more about how the work should be structured in the first place: what the steps should be, which decisions matter, where humans should remain involved, and how multiple agents should coordinate to produce the best outcome.

That is likely to become the next bottleneck considering AI can only execute as well as the work has been designed.

Over time, humans will not be able to optimize those systems perfectly either, and AI will begin to take on the design and orchestration of work itself. This realization led us to notice that Varick is already sitting on much of the data needed to train that kind of system. Every process we map, redesign, deploy, and correct produces examples of how work should be broken down and orchestrated in an AI-native environment.

We used that data to build RL environments for what we are calling an FDE agent: an AI trained not only to execute work, but to determine the best way for humans, software, and other agents to execute it together.

10. The harness matters more than the model

If the model is the brain of an agent, the harness is the body. A powerful brain without a body is still limited, because it cannot see the right context, use tools, remember what happened, or act inside the systems where the work actually gets done.

That is why a large part of the progress we see in AI over the next few years will not come from model improvements alone. It will come from the harnesses built around those models.

Model capabilities already exceed human capabilities in a growing number of domains, but that does not mean the value arrives automatically. The real opportunity is in designing creative harnesses that let those capabilities do things that were not possible in the software era. For most AI companies, that layer will become the real source of differentiation.

The closest comparison is electricity. Electricity itself became a commodity, but the largest changes to humanity came from the applications built on top of it: refrigerators, toasters, factories, computers, and eventually electric vehicles. None of those products were obvious when the electrical grid was first being built.

Intelligence is now starting to follow the same path. The models are becoming the underlying commodity, while the harness determines what that intelligence can actually do. Most people still underestimate the size of that opportunity because we have not yet seen the equivalent of the toaster, the refrigerator, or the electric vehicle for intelligence.

LOOKING AHEAD

That is what makes this moment so exciting. We are no longer only improving the source of intelligence, but are beginning to discover all the things that become possible once intelligence can be given a body and put to work.

Enterprise operations are where many of those possibilities will first become real. They contain the systems, workflows, decisions, and repetitive work that can now be redesigned around intelligence rather than human labor alone, and the companies that understand this early will have the chance to build fundamentally better ways of operating.

The past few years have taught us a great deal about what that future will look like. If any of these ideas resonate, and your company is looking to transform the way its operations run, you can find us at varickagents.com.