



A FIELD GUIDE

A WORKING DEMO IS NOT A WORKING SYSTEM

Impressive in the demo. Stuck everywhere else.



THE MODEL IS THE EASY PART

NICOLA ONASSIS

Founder, Code Particle | Host, "AI Decisions You Have to Defend"

A NOTE FROM NICOLA

If you are reading this, you have probably already seen what AI can do. Maybe you ran a pilot. Maybe a vendor gave you a demo that looked like magic. And then, somewhere between that demo and your actual business, it stalled.

I have spent my career building software for industries where a wrong answer is expensive: healthcare, finance, law, since 2008. I also host a podcast called *AI Decisions You Have to Defend*, where I sit down with owners and executives and ask them about the real AI decisions they made. What they adopted, what they held off on, and what they wish someone had told them first.

The same thing comes up again and again. The technology usually works. The hard part is everything around it: getting it the right information, deciding what it is allowed to do, checking its work, and knowing who is responsible when it is wrong.

I am not here to sell you AI, and I am not here to talk you out of it. I am here to tell you what I have actually seen, plainly. This is a short read, built around one idea. Let us get into it.



Nicola Onassis

FOUNDER, CODE PARTICLE

INTRODUCTION

The Pilot That Never Ships

By the time you see a demo, AI looks ready. It answers in seconds. It writes the email, drafts the summary, pulls the number. Everyone nods. Someone says we should roll this out.

Then months pass, and it is still a pilot. Still not being used in the real business. Still almost ready.

I see this pattern often. The model worked in the demo. The harder questions showed up when someone tried to put it inside the business. What information can it use? How will anyone check its answers? Who is responsible when it gets something wrong?

Nobody hands you those questions. The vendor shows you the demo. The headlines show you the breakthroughs. Almost no one shows you the part that actually decides whether AI helps your business or quietly becomes a problem. That part is what this book is about.

WHAT IS INSIDE

- 01 Why the demo does not tell you what you need to know.
- 02 The three questions to ask before you trust any AI output.
- 03 A simple place to start, and the questions to answer before you go further.

01 CHAPTER ONE

Why the Demo Lies to You

The most expensive misunderstanding in AI is simple. People watch a model do something impressive and assume the hard part is done. Usually it is the opposite.

A demo gives you a question someone chose in advance, clean inputs, and a forgiving audience. Real work gives you none of that. People phrase things differently. Information changes. Documents conflict. Someone forgets to check an answer. A client relies on something you did not expect them to see.

That is when you find out whether you built something useful or just something convincing.

The demo proves capability. The business still has to prove control.

When an AI effort stalls, the instinct is to reach for a bigger or newer model. When I see it happen, the questions that are actually unresolved are outside the model, the ones nobody asked before the demo became a plan.

02 CHAPTER TWO

What Changes When AI Meets the Real Business

A model on its own is just one piece. For a business to rely on it, a few other things have to be true around it.

It needs the right information, and only the information it should have. It needs rules about what it can access and what it cannot. It needs a way for a person to check whether an answer is right. And it needs someone who owns the result when it goes out the door.



None of that requires you to become an AI expert. These are business decisions. You already decide who can see sensitive information, who reviews important work, and who answers for it when something goes wrong. AI does not remove those responsibilities. It gives you a new place to apply them.

REFLECT

Think about the AI you are using or considering. If it produced a confident, wrong answer tomorrow, who would catch it, and how?

03 CHAPTER THREE

Three Questions Before You Trust Any AI Output

You do not need to be technical to use AI responsibly. You need to ask three questions about any AI output before you rely on it. If you cannot answer them, it is not ready for real work yet.

01

What is it allowed to touch?

Know what information goes in, where it is stored, who can see it, and whether it is used to train someone else's model. Confidentiality starts with a decision you make, not a setting you hope holds.

02

How do you know it is right?

When an AI answer is being used as fact, advice, analysis, or part of a client deliverable, there should be a way for someone to check it against a reliable source. If nobody can verify the part that matters, nobody should rely on it.

03


Who is accountable when it is wrong?

AI drafts. A person signs. Match the review to the stakes, and keep a person responsible for anything that leaves your organization.

These are not technical questions. They are judgment questions, and judgment is your job, not the model's. If you can answer all three clearly, you have done something more useful than picking the newest model. You have worked out where AI fits, where it does not, and what has to be true before you trust it with more.

04 CHAPTER FOUR

What the Demo Did Not Test



Here is how it usually goes.

Your firm uses AI to summarize a long client file. The summary looks excellent. It takes seconds. The reaction in the room is immediate: why are we not using this everywhere?

Then the three questions arrive.

- 01 What did the AI actually see? Did it have every document it needed, including the latest version, or did it summarize an outdated file with full confidence?
- 02 Can the person reading the summary trace an important statement back to the source, or are they trusting it because it sounds right?
- 03 Who checks the summary before anyone acts on it, and did any confidential information go somewhere it should not have?

The demo answered one question: can AI produce a convincing summary? Your business has to answer the rest. That is the whole difference between something that works on a screen and something you can actually rely on.

05 CHAPTER FIVE

What “Ready” Actually Looks Like

So what does ready look like? You know what the AI saw, you can check the important parts of its answer, and someone is responsible for the result. It is boring to describe, and it is exactly what lets you stop worrying.

Careful does not have to mean slow. The point of answering these questions is to remove the reasons someone has to stop the rollout later. A tool that reaches daily work and can be checked is more useful than a faster pilot nobody trusts.

The firms I have watched get real use out of AI did not take bigger risks. They answered the hard questions first, so a careful person could finally say yes.

REFLECT

If you had to defend one AI-assisted piece of work to a client or a regulator tomorrow, could you show how it was produced and checked?

06 CHAPTER SIX

Start Where a Mistake Is Cheap

Agreeing with all of this changes nothing on its own. What moves you forward is one deliberate first step.

The instinct of a busy organization is to go big: deploy everywhere, impress everyone, justify the spend. That is where unnecessary risk starts. Do the opposite. Start where a mistake is cheap.

The best first use of AI is high-volume, low-stakes work that a person already reviews: first-draft summaries, organizing what you have, sorting documents. Learn where it helps, where it fails, and what your team needs in order to trust it.

A GOOD FIRST USE PASSES THREE TESTS



LOW STAKES

A mistake is inconvenient, not damaging.



ALREADY REVIEWED

A person checks this work before it goes out.



HIGH VOLUME

You do it often enough that saving time matters.

Do not start with the process where a mistake would hurt the most. Start with useful work someone already checks, and *let the tool earn a bigger job* from there.

GUT CHECK

Looks Ready Is Not the Same as Is Ready

Before you close this, a simple gut check. Not a score. Just two lists.

LOOKS READY

It gives impressive answers quickly.

People like using it.

A vendor or a pilot has shown it working.

IS READY

You know what information it can access, and where that information goes.

You can trace an important answer back to a source a person checks.

Someone owns the result when it is wrong.

The first list is easy to clear. The second is what tells you whether the business can actually depend on the tool. If you are strong on the left and thin on the right, those are the questions you still need to answer.

SELF-CHECK · TOOL

Ten Questions to Answer Before You Trust AI With Real Work

Here is the full self-check. You do not need to send it to anyone. The point is to find the places where your confidence in the tool is ahead of your ability to control it. Mark each one Yes, Partly, or No.

		YES	PARTLY	NO
WHAT IT CAN TOUCH				
1	We have a clear, written rule about what information is allowed into AI tools.	☐	☐	☐
2	We know, for each tool, whether our data is stored or used to train a model.	☐	☐	☐
3	Confidential or regulated information is handled differently from everything else.	☐	☐	☐
HOW WE KNOW IT IS RIGHT				
4	Any AI output used as fact, advice, or in a deliverable can be traced to a source.	☐	☐	☐
5	A person verifies AI-assisted work before it leaves the organization.	☐	☐	☐
6	We could comfortably show how an AI-assisted result was produced.	☐	☐	☐
WHO IS RESPONSIBLE				
7	A specific person owns how our AI tools behave over time.	☐	☐	☐
8	We know who answers for it when an AI output is wrong.	☐	☐	☐
9	We started with low-stakes work, not our highest-risk process.	☐	☐	☐
10	We are expanding AI into higher-stakes work deliberately, not all at once.	☐	☐	☐

Look first at every answer you marked No. Those are the gaps to understand before you expand the use of AI. Use this on your own. If one of them worries you, that is a good reason to get another set of eyes on it.

YOUR NEXT STEP

Reading this is a start. It is not the work.

Getting AI into real use takes more than a good model.

It takes the unglamorous work around it: access, verification, review, and clear responsibility.

That is what we do at Code Particle. We have built software for regulated, high-stakes industries since 2008. We are not here to sell you a model. We help you put AI into your business in a way you can stand behind, in front of a client, a regulator, or a court.

Bring us the AI use case you are weighing.

If you have an AI use case you are weighing, bring it to us. We will spend 30 minutes looking at the practical questions with you: what it would need access to, what could go wrong, how you would check the output, and whether it is a sensible place to start.

BOOK A FREE 30-MINUTE CONVERSATION →

The goal is to leave you with a clearer view of the next step, even if that step does not involve us.

LISTEN

Hear the podcast, *AI Decisions You Have to Defend*

► **LISTEN TO THE PODCAST**

