WHITE PAPER

The Consequence Gap

Why an Agent Can’t Restrain Itself the Way a Person Does

Author: Srinivasarengan Varadharajan
Published: July 27, 2026
Contact: srini@kter.ai

OPENING

People Know What Will Get Them Fired

That line came from Mark Cuban, not from anyone inside AI. He said it this week in a thread with Anthropic's Head of Economics and Box CEO Aaron Levie, arguing about whether AI agents are actually replacing white collar work.

> "*Models don't know the consequences of their actions. People know what will get them fired.*"
>
> Mark Cuban, July 2026

I've had that line stuck in my head for a few days now. Not because it's a clever line. Because it's true in a way most of the AI industry hasn't sat with yet.

The two earlier papers in this series looked at two kinds of boundaries. One was contractual, who is responsible when an agent's output is wrong. The other was technical, what an agent should ever be allowed to reach. Cuban is pointing at something underneath both of those. Not what an agent is allowed to do. What it would never have thought to do in the first place, if it had grown up the way a person does.

SECTION 1

What a Person Actually Carries

Think about the last time you almost said something in a meeting and stopped yourself halfway through the sentence. Nobody handed you a rule that covered that exact situation. You just knew, somehow, that finishing the sentence would cost you something.

Where did that knowing come from? Not a handbook. You picked it up by being in enough rooms. You watched someone else say the wrong thing once and saw what happened to them. You felt a silence land after a comment and never made that comment again. None of it was written down anywhere, and none of it needed to be, because you were there for it.

That's the thing Cuban is calling consequence. It isn't a rule. It's an accumulated, mostly unconscious sense of what costs you something, built over years of actually being present when things went wrong for other people.

SECTION 2

An Agent Has Never Been in the Room

An agent doesn’t get years. It gets a prompt, some context, and whatever scope someone happened to specify before it started. It has no memory of watching a colleague get quietly sidelined after a bad call. It has no version of the stomach drop you feel right before you say the wrong thing.

So when an agent does something that would obviously cost a person their job, it isn’t being reckless. It just has nothing telling it not to. The restraint a person carries by default has to be built in from outside, deliberately, by someone who thought to specify it. And most of the time, nobody thinks to, because a human employee never needed it written down.

This is a different problem than a missing permission or an unclear contract. You can lock down exactly what an agent is allowed to touch and still have it say the one thing that ends a client relationship, because nobody ever told it not to, and nobody thought they’d have to.

SECTION 3

Where This Sits Next to the Other Two Boundaries

It would be easy to describe this as the deepest layer yet, as if each paper in this series has been peeling back something more fundamental than the last. I don't think that's quite honest. It's a different angle on the same underlying fact, not a discovery underneath the other two.

The contractual boundary from the first paper assumes there's a document specifying scope. The technical boundary from the second assumes there's a permission system enforcing what's reachable. Both of those still assume somebody already knows what needs to be written down. Cuban's point is that most of what keeps a person in line was never written down for them either. It was just absorbed.

So the honest version of this paper isn't "here's the real answer." It's that specifying consequence for an agent means doing, on purpose and in writing, something nobody ever had to do on purpose for a human being.

SECTION 4

What a Person Has by Default vs. What an Agent Has to Be Given

	A Person	An Agent
Sense of consequence	Absorbed over years, mostly unconscious.	None, unless someone specifies it directly.
Learns from watching others	Yes, constantly, without trying to.	No. Every session starts from the same blank state.
Knows what ends a relationship	Usually, without being told.	Only if someone wrote it down first.
Restraint when unsure	Defaults to caution, out of instinct.	Defaults to completing the task.
Who has to do the specifying	Nobody. It just happens by living it.	A person, deliberately, in advance.

CLOSING

This Might Not Fully Close

Cuban’s diagnosis holds up. A written scope can specify a great deal. It can name what an agent should never say to a client, what topics are off limits, what tone to avoid, what to escalate rather than answer. All of that helps.

What it can’t do is give an agent years of watching consequences land on someone else. That’s not a document problem. I’m not sure it’s a solvable problem at all, not fully, not with a better contract and not with a smarter model.

What a formal scope can do is force the honest work of writing down what used to just happen for free. That’s not the same as closing the gap Cuban pointed at. It’s the most honest attempt at it available right now, and it’s worth doing even knowing it won’t fully get there.

This is what we’re addressing by building KTER, named agents with a defined scope, because the agent can’t supply the judgment a person carries for free. The structure has to do, from outside, some part of what a person’s instinct already does without being asked.

One more thing worth naming honestly: the same question applies to code an agent writes, too. We’re building KOCODE against that same principle.

About the Author

Srinivasarengan Varadharajan is the founder of KTER. He has 19 years of enterprise integration experience across financial services, retail, and technology organisations. Contact: srini@kter.ai

© 2026 KTER. All rights reserved. kter.ai