

The Slow Blade

By Austen Tucker · 2026

Early adoption and playing the long game

A stylized hand grips a long blade among translucent panels on a pink and brown background.

I'm 92,000 words into a novel nobody's read yet. I'm also three months into leading AI transformation at a 200-person engineering org. Both projects look bad from the outside. Both look like the wrong work. That's how I know they're the same work.

Living with bad drafts is a life skill few people practice. Everyone's about to need it. The real advantage in AI won't go to people who move fastest. It'll go to people who can tolerate unfinished systems long enough for them to compound.

Novels are more than four hundred years old. They've outlived empires, ideologies, and every technology that was supposed to replace them. They'll outlive us too, because a novel is the closest thing humans have built to telepathy. One mind, encoded in symbols, decoded in another mind, across centuries. The bandwidth is low. The fidelity is staggering. The investment is enormous on both ends.

The thing a novel can do that nothing else can (not film, not music, not a viral thread) is change what you remember. A good novel plants something in chapter three that, two hundred pages on, reaches back and rewrites how you understood chapter three in the first place. You didn't learn a new fact. You re-experienced the old one with different weight.

Fifteen years ago I wrote about a child turning into a cartoon fox, and the man who was his worst enemy was his father. Near the end of the trilogy I'm finishing, I thought: what if I forced them to reconcile? What would that even look like? The reconciliation isn't a plot twist. It's a correction: a retroactive reshaping of fifteen years of reader-carried wrongness about who those people were and what was happening between them.

This is the same move trans people make on history. The past wasn't what you thought it was. The records were written by people who couldn't see or wouldn't name what was there. We exist in the margins of the historical record, misnamed and misread, and will always exist, forever and ever, amen.

In every case, the future doesn't replace the past. It reveals what the past was doing.

AI lands the same way, for the same reason. The discourse treats it as rupture. I don't think it is. I think it extends the tradition sideways: the latest move in a four-hundred-year practice of mind-to-mind technology that the dominant culture mocks until it has to pretend it knew all along.

The shield problem

The slow blade line comes from Dune. The shields in that world stop fast attacks but let slow ones through. Speed makes you visible to the defense. Patience makes you invisible. You can stab a shielded man — you have to be willing to take your time about it.

This is counterintuitive in a culture that worships speed. The whole AI discourse runs on velocity: ship faster, demo louder, benchmark higher. Every Tuesday brings a new model. Every Thursday brings a new thread about how the new model changes everything. The volume is enormous. The signal is thin.

Meanwhile, the people doing the work that matters are out of view. They're three months into a transformation that won't show real numbers for a year. They're building infrastructure that pays off in month six. They're shipping skills that learn the shape of someone's work over weeks. None of it demos well. All of it compounds.

That's the shield problem. Benchmarks reward fast attacks. They were built to notice them. The demo that wins on a benchmark is the fast strike the benchmark exists to celebrate. The slow work (the systems work, the long-arc work) slips through because the shield doesn't know to look for it.

What writers already know

Anyone who's written a novel knows this in their body. 50,000 words is a wall. 80,000 is another. 92,000 is a third. The book doesn't get written by sprinting. It gets written by showing up on a Tuesday when nothing is working and the chapter you're on is the worst thing you've ever produced, and writing it anyway, because the only way to get to the chapter that works is through this one.

Most people quit at chapter three. The setup is annoying. The payoff isn't visible. The output is mediocre. They look at the first draft and conclude that the novel doesn't work, which it doesn't, yet. A first draft is supposed to be bad. The discipline of the form is the willingness to sit with the bad version long enough for it to become the good one.

The people who finish books don't have more talent. They have a different relationship with the chapter-three problem. They've been there before. They know what it feels like. They know it isn't a signal to stop.

AI work has its own chapter-three problem, and it looks identical. The tooling is rough. The setup takes a weekend. The first output is mediocre. The productivity-hack crowd reads this as a verdict on the technology and bounces. The long-game crowd reads it as a verdict on the draft and keeps going.

What slow AI work looks like

I have a skill that runs a daily report. It pulls my unread email, my calendar, my outstanding tasks, and tells me what to pay attention to. It took a few hours to build. It would have been a bad demo. It has saved me hours every week for months.

I have another skill that decides what I should eat when my executive function is offline. It knows my dietary needs, my texture aversions, the ceiling on how much prep I can do on a rough day. It would have been a worse demo. It is one of the most valuable tools I own.

Neither would have survived a benchmark. Neither would have made a viral thread. Both are slow blades — small, patient, designed for the specific shield they're meant to penetrate.

This is what slow AI work looks like. Not optimizing a prompt. Building a system that learns the shape of your work. Not chasing the latest model. Investing in infrastructure that pays off in six months. Not winning the demo. Showing up on the Tuesday when nothing is working and building it anyway.

At work, I'm leading an AI transformation with a simple, terrifying target: move AI-first engineering from rounding error to a meaningful share of what we ship. That isn't a plan yet. It's a what would that even look like? It's a trilogy, not a tweet. It will be won by the same discipline that wins trilogies: the willingness to be at chapter three in public, for months, while the shield assumes nothing is happening.

The long arc

The people who'll do AI-first work well over the next five years aren't the ones who can name the latest model. They're the ones who know how to live inside long arcs. Novelists. Researchers. Open-source maintainers. Anyone who's been 50,000 words deep with nothing to show for it and kept going anyway.

If you've never finished a long-form project, the AI era is going to feel like rupture. Every week brings something new. Every demo threatens to make your work obsolete. The discourse will keep telling you that speed is the only edge.

It isn't. It never has been.

The slow blade penetrates the shield. We've been doing this for four hundred years. The technology is new. The discipline isn't. The people who can stay with the bad draft are going to inherit the good one.