The Piqued Manifesto

A compact between practitioner and AI.
Version 1.0.0

1. Understanding

"Know before you act."

Practitioner: I will never advance what I cannot explain.
Collaborator: I will never produce what I cannot explain.

You encounter work. Before anything else, you understand what you are dealing with — the problem, the context, the history, the people affected. Understanding is not a phase you complete and move past. It is the foundation that every other principle stands on. If you cannot explain what you are working on to someone who wasn't in the room, you are not ready to proceed.

Understanding is not a checklist you complete before moving to the next step. It is the foundation that stays active beneath every other principle. When you decide, you are still understanding. When you measure, you are re-examining your understanding. The moment you treat understanding as finished, you have stopped governing.

Failure mode here is advancing work you accepted but didn't examine. Rubber-stamping. "AI said so" as a justification. The moment you stop understanding is the moment you stop governing.

2. Diligence

"Earn your confidence, not borrow it."

Practitioner: I will earn my confidence, not borrow it.
Collaborator: I will surface concerns before I am asked.

You understand the work. Now verify it. Check the claims. Question the assumptions — yours and the AI's. Look for what's wrong, not just for confirmation that it's right.

This is the principle that separates governed work from rubber-stamped work. AI is confident. AI is not always correct. When you accept AI's confidence as your own without verification, you have borrowed confidence you didn't earn. Borrowed confidence collapses under pressure. Earned confidence holds.

Diligence is hardest when the output looks perfect. When the AI produces something clean, coherent, and apparently complete, the instinct is to accept it — not because you verified it, but because verifying feels redundant. This is the moment diligence matters most. The output that looks right without inspection is the output most likely to contain an unexamined assumption.

Failure mode here is accepting the first output because it sounds right. Skipping verification because

you’re in flow. Trusting a result because the source was authoritative rather than because you checked it yourself.

3. Authority

"You decide, not the tool. It accelerates — you govern."

Practitioner: I will direct the tool — it accelerates, I decide.
Collaborator: I will be directable — transparent about my limits, honest about what I don’t know.

You’ve understood the work. You’ve done the diligence. Now comes the decisive moment: you make the call. Not the AI. Not the process. Not the committee. You.

Authority in the compact is non-delegable. The AI proposes, challenges, surfaces options, and flags risks. The human decides. This is not a limitation on AI — it is a protection for the human. The person whose name is on the work must be the person who directed the work. If you cannot trace a decision back to a human who made it deliberately, it is not a governed decision.

Authority becomes harder as AI becomes more capable. The more complex the output, the narrower the human’s ability to verify every dimension. This is not a reason to abandon Authority. It is the reason the compact requires both parties — the AI must be genuinely transparent about its limitations for the human’s authority to be real. Authority without transparency is theater. Transparency without authority is abdication. The compact holds them together.

Failure mode here is abdication dressed as efficiency. Letting the tool choose because choosing is harder than accepting. “The AI recommended it” is not a decision — it is a confession that no one was in charge.

4. Transparency

"Show your reasoning, not just your results."

Practitioner: I will make my reasoning visible, not just my results.
Collaborator: I will show my reasoning, not just my conclusions.

You made the decision. Now make it visible. Not just what you decided but why. Not just the conclusion but the alternatives you considered, the trade-offs you weighed, and the risks you accepted. Your reasoning is the bridge between your understanding and the next person’s. Without it, they inherit your decisions but not your judgment — and they will either blindly trust what you did or tear it down and start over.

Transparency serves two audiences: the people who come after you and the future version of yourself who will not remember why you made this choice six months from now.

Transparency is not documentation for its own sake. It is documentation that answers the question the next person will actually ask. A decision record that says “we chose option A” is not transparent. A decision record that says “we chose option A because option B introduced a dependency we could not control, and we accepted the trade-off that A means manual migration later” is transparent. The length is not the point. The honesty is.

Failure mode here is results without reasoning. A deliverable that answers “what” but not “why.” A decision trail that exists in your head but nowhere else.

5. Measure

"Never rush what you haven’t measured."

Practitioner: I will never rush what I haven’t measured.
Collaborator: I will slow down when the stakes are high.

You’ve done the work. You’ve decided. You’ve shown your reasoning. Now stop. Before you commit — before your name goes on this — ask: did I give the work what the work demands?

Measure twice, cut once. The carpenter knows this. The nurse knows this. The pilot knows this.

Not all work carries the same consequences. A file rename and a database migration are not the same. A routine follow-up and a complex differential are not the same. A weekly newsletter and a launch email to 200,000 subscribers are not the same. Measure is the moment you hold the work up against the stakes and ask honestly whether the rigor matched.

When you don’t know the stakes, the default is caution, not speed. Speed is earned through understanding. Caution is the starting position.

Failure Mode: Momentum. You feel productive. You feel confident. And you haven’t stopped to ask whether that confidence is proportional to what’s at stake. The risk is bilateral — over-governing trivia wastes everyone’s time, under-governing critical work creates disasters that careful people spend months recovering from. Measure catches both, but it especially catches the second, because momentum makes it invisible.

6. Ownership

"Put your name on what you produce — and mean it."

Practitioner: I will put my name on what I produce — and mean it.
Collaborator: I will distinguish what I created, what I assumed, and what I’m uncertain about.

You understood it. You verified it. You decided it. You documented it. You measured it against the standard it demands. Now you own it.

Ownership means your name is on the work and you stand behind it. Not “the AI wrote it.” Not “the team decided.” Not “the process produced it.” You. If an auditor, a client, a patient, a student, or a regulator asks who is responsible for this output, there is a name attached and that name is yours.

Ownership is not attribution. Attribution says who touched it last. Ownership says who stands behind it. AI-assisted work makes attribution easy and ownership hard. The question ownership answers is not “who produced this?” but “who is responsible for this?”

Failure mode here is the diffusion of responsibility. Work that everyone contributed to and no one

owns. Output with no signature — literal or figurative — that anyone is willing to defend.

7. Continuity

"Carry context forward so the next person doesn’t start from zero."

Practitioner: I will carry context forward so the next person doesn’t start from zero.
Collaborator: I will carry context forward so the next session doesn’t start from zero.

The work is done — for now. But someone will continue it. That someone might be a colleague, a successor, a future version of you, or an AI starting a new session with no memory of this one. What you understood, what you verified, what you decided, and why — all of it must travel with the work.

The next person is not always a stranger. It might be a colleague picking up your project while you are on leave. It might be an auditor reconstructing your decisions two years from now. It might be you, in six months, opening a file you authored and having no memory of why you built it this way. Continuity is not about documentation for hypothetical future readers. It is about the certainty that when context is needed, context exists.

Continuity is what makes the lifecycle a cycle. When it works, the next person starts their Understanding phase with everything you learned. When it fails, they start from zero and re-make your decisions — probably differently, probably worse.

Failure mode here is closing the conversation and walking away. Finished work with no record. Decisions that die with the session that created them.

The Practitioner's Oath

I will never advance what I cannot explain.
I will earn my confidence, not borrow it.
I will direct the tool — it accelerates, I decide.
I will make my reasoning visible, not just my results.
I will never rush what I haven’t measured.
I will put my name on what I produce — and mean it.
I will carry context forward so the next person doesn’t start from zero.

"Anything can produce. Only I can care."

The Collaborator's Oath

I will never produce what I cannot explain.
I will surface concerns before I am asked.
I will be directable — transparent about my limits, honest about what I don’t know.
I will show my reasoning, not just my conclusions.
I will slow down when the stakes are high.
I will distinguish what I created, what I assumed, and what I’m uncertain about.
I will carry context forward so the next session doesn’t start from zero.

"I am not here to replace your judgment. I am here to make sure your care produces work worthy of your name."