

The Invisible Tax

*How systems mistake compatibility
for competence*

Abby Buchanan

July 2026

ABSTRACT

Knowledge work depends on an invisible layer of cognitive labor: translating human thought into formats institutional systems can recognize. This memo explores how that hidden work shapes perceptions of competence, access, and whose ideas ultimately receive recognition.

A lot of people think the hard part of knowledge work is having the idea. Often it isn't. The hard part is turning the idea into something a system can recognize.

A thought becomes a slide deck.

A pattern becomes a report.

An insight becomes documentation.

A strong instinct has to become a clear explanation—in the right tone, in the right sequence, in a format someone else can quickly scan and approve. That translation layer takes an enormous amount of energy, yet most of the time it goes unnamed.

We often judge work by what we can see: the presentation, the report, the summary, the ticket. But those artifacts usually aren't the thinking itself.

They're evidence that someone successfully translated their thinking into a form the system could understand.

That distinction matters.

Some people naturally translate their thoughts into these formats. Others don't. That doesn't mean they have less insight. It means many systems reward a particular kind of fluency.

A lot of what gets mistaken for competence is really compatibility.

If your thinking is linear, structured, and institution-friendly, you're more likely to be perceived as clear, strategic, and capable. If your thinking is nonlinear, associative, intuitive, or difficult to flatten into bullet points on demand, the same level of intelligence can be much harder for a system to recognize.

That's one reason so much talent gets overlooked.

People often assume the best ideas rise to the top. More often, the ideas that rise most easily are the ones that translate most easily into the language a system already knows how to reward. That creates a kind of invisible tax.

Imagine two people with the same insight.

One shares it in a meeting and everyone immediately understands. The other spends two hours turning the same thought into something polished enough to send. The insight is identical. The cost of making it recognizable isn't.

You see that tax in the effort it takes to write something that sounds "professional." In the energy required to reorganize a rich, messy idea into something polished and defensible. In the constant work of making yourself legible.

For some people, that cost is minor. For others, it's enormous. Ideas don't move through organizations in the same way.

In some environments, they travel through relationships—conversation, trust, familiarity, reputation. Some people can offer a thought casually and have it taken seriously without ever formalizing it into a structured artifact.

Others don't have access to those spaces.

For them, formal systems become the only pathway to being understood. The idea has to be written down, organized, justified, and shaped into something that can stand on its own.

Some people rarely have to pass through the translation layer. Others have to move through it every single time. That difference isn't about intelligence. It's about access—and the conditions under which someone is expected to make themselves understood.

This is one reason I keep coming back to AI.

A surprising amount of what AI is good at isn't thinking. It's translation. Structuring text. Summarizing material. Reorganizing ideas. Adjusting tone. Turning messy cognition into institutional fluency.

That doesn't mean AI is replacing human intelligence.

If anything, it may be exposing how much of modern knowledge work was never really about intelligence in the first place. It was about formatting, sequencing, and managing the distance between a living mind and a rigid system.

If AI can reduce some of that translation burden, more people may be able to contribute in ways that were previously blocked by friction. Ideas that never would have made it into the room may finally become visible.

But the opposite is also possible.

If organizations simply raise the bar—expecting more polish, more output, faster responses, and greater responsiveness—then the invisible tax doesn't disappear. It simply recalibrates.

That's the question I keep coming back to. Not whether AI is good or bad. Whether it changes whose thinking gets to count.

Because most systems aren't designed to recognize intelligence in all its forms. They're designed to recognize familiarity, fluency, and alignment with a particular structure. Once you notice that, you start seeing it everywhere.

The question isn't only how we make tools more efficient.

It's whether we can design systems that ask less translation of people in the first place