

Intelligence is not the main bottleneck

Confessions of a naive hamster: why the smartest people in the room keep missing that intelligence is often not the main bottleneck in the real world



RUXANDRA TESLO
JUL 21, 2026

♡ 614

💬 114

🔄 135

Share

Edit: To be clear, I'm not saying everyone or even a majority of people in the SF tech scene think this way. But some certainly do (e.g. the opening conversation is verbatim) and this post is about a particular core scene whose views carry disproportionate weight in the discourse and in policymaking.

At a recent dinner, someone from an AI lab asked me, quite bluntly, why I do what I do — why I write about the regulatory bottlenecks to medical progress, why I spend my time on policies related to clinical trials. He looked at me with something best summarized as pity.

I told him what I believe: that in the age of AI, medicine will be bottlenecked more than ever by regulation and the grind of clinical trials, and that billions poured into faster pre-clinical research won't touch that problem, unsexy as it is. I brought up housing: we've had the technology to build better housing for decades, yet it's more expensive than ever, because housing is a question of political will, not pure capability.

Ruxandra's Substack is a reader-supported publication. To receive new posts and support my work, consider becoming a free or paid subscriber.


Subscribe

He looked at me incredulously. Surely, a smart person like me should know that AI, or better said, AGI will be hyperpersuasive soon — already on a bunch of benchmarks it exceeds professional debaters at persuasion. I said, “*Hmm.*” He said, “*Yes, yes*”, with the undertone of a man who knows things deeply, things that mere mortals like me, not being AI lab employees, are simply not well placed to grasp. And he looked at me again with that sense of pity, the way one looks at a slightly mentally impaired but cute animal awaiting its imminent slaughter.

That night I lay awake, twisting and turning, wondering whether my life had any point at all. Whether every decision I'd ever made had been, in some way, a mistake. What could I have done better?

But as the sun came up, I found my way back to the same thought: no matter how “intelligent” AI becomes, intelligence is often not the main bottleneck to things changing in the real world. It is hard to hold on to that conviction when people smarter

than you, with access to privileged information, insist otherwise. After all, this could all just be “cope” from the naive hamster awaiting its slaughter.



A scared hamster watched over by the AGI

But, I shall nonetheless stick to my beliefs.


Because the people at these labs made fortunes betting on ideas that once looked insane, the world now takes nearly everything they say on faith. That, I think, is a mistake. It is not a given that AI will solve the problems most people actually care about, including medicine, unless we think about those problems clearly. And at the moment, I don't believe we are or at least, not to the extent that we could. My case is simple, and it comes from two directions. One is what I observe in the social dynamics of San Francisco, and in how people there talk about all this. The other is what I notice in a field I happen to know something about: medicine.

AI and medicine

One of the promises most often invoked to justify AI's risks is that it will “*cure disease*.” Every [major AI lab CEO says it](#), and investors seems to agree: any biotech startup with an AI story attached commands an impressive valuation, even as more conventional biotechs struggle for funding and die. But this whole enterprise, as I have long argued, is bottlenecked by many things that have little to do with “intelligence” as such, and the degree to which that often goes unacknowledged is strange to watch.

The evidence is everywhere, if you treat scientific advancement as a rough proxy for intelligence and ask whether it alone unblocks progress. Take [Froom's Law](#): the number of new drugs approved per dollar of R&D has fallen for decades, even as our scientific tools have grown vastly more powerful, the exact opposite of what the existence of more raw capability would predict. Or take a company like Adaptimmune, which has brought [two transformative therapies](#) to market in rare cancers and is [nonetheless fighting to stay alive](#), due to the cost of developing them. Or one can listen to the scientists behind “baby KJ,” the infant saved by a bespoke gene-editing therapy: they [have everything they need scientifically and still cannot easily repeat it for the next child](#), because manufacturing costs, driven in part by regulatory requirements, stand in the way.

Number of drugs approved per \$1 billion of R&D per year



Note: Y-axis is in log scale.

Chart: Ruxandra Teslo • Source: The Economist, Jack Scannell's calculations



Figure 2. Pharmaceutical R&D productivity has steadily decreased since 1960, despite advances in basic science. In the last decade, the deceleration in pharmaceutical productivity [seems to have ameliorated](#), due to a combination of factors including an increase in predictive validity (through e.g., genetics), but also due a larger share of efforts being directed at oncology and rare diseases, where the burden for approval is often lower due to the high unmet need.

In biopharma specifically, one great bottleneck is [clinical trials](#) (I know I keep banging on about this), which today consume something like seven years and cost more than a billion dollars per drug. And trials are not a mere formality. They generate precisely the kind of data that would train better models in the first place: [human data, which is ultimately irreplaceable](#).

We have extensive evidence faster clinical trials are important for biomedical innovation, both directly and indirectly, through second-order effects (like increasing the risk appetite of those working in the industry and helping better align incentive through fast feedback loops). This ranges from robust economic papers showing shorter trials massively boost investment in a disease area, all else equal, to the natural experiment of China. China is threatening to race ahead of US in biotech, with Chinese biotechs encompassing more [than a half of big Western pharma](#) licensing deals. A mere decade ago, this percentage was zero. This transformation happened while China remains worse in terms of basic science, mostly due to [regulatory reforms](#) that allow faster iterative learning using in-human data ¹.

Trials themselves can be made shorter, more informative and better, including with AI. One area I am particularly bullish on [is surrogate endpoints and biomarkers](#), where AI could turn discrete, coarse readouts of whether a drug is working into fast, continuous ones. This would in turn optimize trials immensely. For some indications, the gain could be as high as [10x faster and cheaper trials](#) if the right surrogates are found. And this is not even considering the benefits of simply optimizing the trials without necessarily shortening them: imagine having the right biomarker that tells you whether a therapy is working early enough.

Figure 3. One way in which surrogate endpoints help is by increasing incentives to invest in an area. I can confirm that the findings from the economics paper that shows this are true, as I could notice myself how after the FDA allowed Bone Mineral Density (BMD) as a surrogate endpoint in osteoporosis, there has been a stark increase in the appetite from big pharma in pursuing this indication.

But even here the binding constraint is not entirely intelligence. Quite often, it is governance. I keep talking to companies trying to build exactly these biomarkers, and what they run into, again and again, is how hard it is to access the underlying data. Some have been waiting for a year for the NIH to release imaging datasets they can use to produce better biomarkers. If one needs to interact with the FDA to get their endpoint validated, it is even worse: I have previously written about how the validation of Bone Mineral Density (BMD) for use as a surrogate endpoint in osteoporosis trials [took 12 years](#) (!), despite the fact that the data to support it already existed in full and the analyses done were basically regressions.

When people in the AI sphere do engage with regulation and governance in medicine, it often seems to happen in a rather superficial way. I have seen many times that argument that because most drugs fail for lack of efficacy, regulation can't explain much of the slow-down in medical progress compared to basic science that we have witnessed in the last decades: after all, approving more inefficacious drugs won't help. This analysis sounds superficially true, but it's not. As I keep saying, when people talk about the importance of regulation in slowing down biomedical progress, they rarely mean the approval decision at the end. It is about the entire process upstream: how easily human data can be collected, and then how easily, once collected, it can actually be used.

I know most about clinical trials, but this is not the only area where governance and policy slows down progress. Another example is related to how the structure of our patent system shapes how innovative our drugs are.

In biomedicine we have a problem called [target herding](#): most companies chase the same biological targets because they are de-risked, which means we explore far less of the biological space than we could. The cleanest way to see why is through risk. Two kinds matter in drug discovery. *Target risk* captures the biological uncertainty: is this protein causally involved in the disease, and will hitting it help a patient without unacceptable toxicity? *Molecule design and optimization risk* is the downstream problem, conditional on the target: designing a novel molecule against a target² and improving potency, selectivity, safety and a myriad of other features of a chemical (e.g. small molecule) or biological (e.g. antibody). That is mostly in the realm of chemistry and is intrinsically far more tractable and predictable; given a validated target, good teams usually reach a decent molecule.

Figure 4. The logic behind target herding

The patent system then further disincentivizes companies from taking target risk, because it rewards novel chemical matter, not novel biology. Composition of matter is what most patents protect. That refers to a specific molecule, which means that they don't protect the insight that a target is worth drugging at all.

Target herding then follows almost mechanically from the above mentioned constraints: on one hand the fact that is genuinely harder to uncover more of these and then that the patent system does not reward it. Whoever validates a novel target (and implicitly, discovers new biology) bears the enormous target risk but can't capture the reward due to the way the patent system is designed: once a first-in-class drug ³ posts convincing clinical data, the target is validated, and that validation is essentially public. Competitors then design distinct molecules against the same de-risked target, and each fast-follower earns its own strong patent. So firms end up piling onto proven targets instead of validating new ones, so we get wave after wave of drugs aimed at the same type of biology.

A well-known non-profit leader who funds work on the molecular basis of deadly diseases, including AI-driven efforts to find novel targets, told me the current IP system makes it nearly impossible for the institution he supports to capture the value it creates by uncovering new biology. He's glad to keep funding it regardless, but believes everyone would be better served if that value were capturable, and that today's incentives are badly skewed.

The irony is that the AI-for-biology companies drawing the largest rounds are aimed mostly at the same well-served problem of molecule design and optimization. For example, Chai Discovery, [now valued at \\$3.8 billion](#), enormous by early-stage biotech standards, and Isomorphic Labs are both, at heart, at molecule design and optimization engines: they design better compounds faster and with more accuracy, given a target. This is great, but as I and others have argued before, expanding the frontier of validated targets would be worth more overall to society. Yet this is precisely the part that the patent system does not incentivize. It turns out even AI-driven companies ultimately respond to economic and legal incentives!

The same principle applies, I suspect, well beyond medicine. Diffusion is hard, and the rate at which a technology actually spreads through an economy is not entirely or even mainly predicted by the raw pace of its underlying progress. A capability existing is not the same as a capability being absorbed into institutions, and, perhaps most important of all, translated into outcomes people care about.

And there is expansive evidence this is the case.

Walking around the world today one might notice that it is weirdly unchanged. LLMs can spawn agents and write complex applications, yet something as mundane as customer service still hasn't been automated well. In fact, often it's the automated chatbots that are the worst part of the experience. GDP is up, which is good, but it is not dramatically up. And entry-level jobs, despite repeated warnings of an apocalypse, [seem largely intact](#), if not increasing.

To many, this is surprising. Just the other day I was at a conference where someone remarked that if he could have seen today's AI capabilities a few years ago, he would have been astonished — and would have assumed the world by now would look far more transformed, with much higher GDP growth. It was offered as a fresh observation. But in fact, it has been made before, and this whole trajectory was quite predictable to those thinking about real world barriers to diffusion. I remember a Progress Conference in 2023 where the economist Tyler Cowen made exactly this argument to a room full of AGI-pilled attendees, many of whom were frankly incredulous. The capabilities, he suggested, would arrive faster than the changes they were meant to produce; GDP growth would disappoint relative to the hype, and there would be no massive job displacement. So far, the data seem to bear him out.

The people who saw this coming simply have a better feel for friction for how slowly and unevenly a technology actually diffuses into the world. Strangely, it seems to me that the core San Francisco tech scene, full of exceptionally intelligent people, does not. How can so many smart people share this one blind spot?

The making of a monoculture

My view is that certain communities in San Francisco (which I shall henceforth call “SF people”, even though I do not mean literally everyone or even most people in SF) have become something close to a monoculture, and this monoculture has two reinforcing features. The first is that it has become genuinely low-status to suggest AGI might not be wholly transformative, at least not on its own. The people who were early and right about AI's potential made a great deal of money, and being right about the big thing has hardened into an assumption that they are right about everything. To question the most extreme claims is to be marked as dim, low-IQ, insufficiently “AGI-pilled.” The belief persists that anyone who works at a lab, specially a CEO, holds secrets about the future that no outsider can access. The result is self-censorship: a lot of smart people I know in San Francisco see the bottlenecks to AI's diffusion and say nothing to seem more “AGI-pilled”.

But how could such smart people be wrong... About anything? I do not think there is anything weird or unusual about that and in a wiser era we would find it easier to articulate. I have said before that the best human is very far from the average human, but God is much further way from the best human still. That is to say that nobody, no matter how smart, powerful and experienced, has all the answers.

Not only that, but I think the winners of the AI race might be miscalibrated in a way that might have helped in their core endeavour, developing AI, but maladaptive when it comes to judging things beyond that. Delusion is often adaptive in startups: an unreasonable conviction is exactly what lets a small group attempt the impossible and occasionally pull it off. Byrne Hobart and Tobias Huber make a version of this case in [Boom](#), arguing that financial bubbles, usually maligned as destructive, have in fact been an engine of breakthrough progress: what looks like collective delusion is often what funds and forces genuine advance. But that cuts both ways. Being productively

delusional and winning does not make one right about everything else. In fact, the very trait that produced the success also might inflate confidence in unproductive ways in other areas.

The second feature of this monoculture is a kind of counterfeit contrarianism. Because most of the outside world has long been skeptical of AI's potential, people in San Francisco feel like brave dissenters. This is despite the fact that, within their own circles, holding those views is the safe and status-conferring position. Believing yourself a misunderstood contrarian is intoxicating, and it makes the underlying conviction harder to dislodge.

Back when I was an academic, I often noticed that academics were, on the whole, deeply progressive, yet thought of themselves as contrarians. This was at least in part because most had come from places more conservative than the university. They'd gather, compare themselves to the outside world rather than to each other, and conclude that their progressivism was somehow brave, despite it being the plain consensus of the room. In my view, SF has done the same thing with AI. In elite tech circles it is simply not contrarian to take AGI seriously; if anything the reverse. But because the comparison class is always someone else: for example, a policy person in DC, the feeling of embattled dissent survives contact with reality.

The question, I suppose, is whether any of this matters. Why would it be bad? For all I know, the AGI God frenzy should continue in the tech bubble. But it's seeping into general public discourse in a way that might realign priorities in a way that is bad.

At the most basic level, I think being closer to the empirical truth is simply better than being further from it. But there is a more concrete cost: mistaken beliefs can lead to mistaken priorities. Given that I do expect AI to be disruptive (albeit not necessarily in the same ways the SF crowd thinks), this might be a particularly bad times to have bad priorities.

Again, I shall bring the example of medicine: we might be aiming or effort at the wrong bottlenecks. In medicine, as I've argued, regulatory reform gets a tiny fraction of the attention lavished on AI-enabled biology and even within AI-enabled biology, the interventions that would actually move the needle are the ones that often seem less funded.

And this goes well beyond medicine. We fixate on AGI gods, while more mundane problems pile up unattended. The social fabric is fraying in plain sight. More and more young adults cannot afford a house or sustain a relationship. People pour their youth into gambling apps, some of them now made more powerful by the very AI we keep discussing. People are de-skilling as they hand more thinking to machines. I sometimes say I worry less that AI will take my job than that people will grow too incurious to read at all. Every one of these problems needs people working on it. Being in the thrall of digital deities that will magically dissolve all our problems pulls attention away from the simpler, more human troubles that are already here.

- 1 Milan Cvitkovic pointed out that it is not clear that the regulatory reforms around faster in-human data spearheaded China's rise. In 2015 China did a bunch of policy reforms that included expansion of coverage, clearing up some regulatory decision backlog and others and it is hard to say which ones were the most consequential. This [article argues](#) it was the coverage expansion. I explained [in this X post](#) why I doubt that article to some extent. Either way, I think everyone agrees that faster in-human data is a large component of China's strength (whether caused by regulatory reforms or not) and also that the proximal cause of China's rise in biotech was policy change. You can find my exchange with Milan [here](#).
- 2 This is in practice a bit more complicated, as not all drugs are molecules. Some are cells or gene therapies. Nonetheless, even here, the broad principles hold. Improving a T cell therapy against a known target is inherently less risky.
- 3 First-in-class drugs are often not the drugs that "win" commercially, as followers usually manage to optimize the chemistry better and learn from deficiencies from the first-in-class. So it's hard for someone who pioneers a new approach to capture value.

Subscribe to Ruxandra's Substack

By Ruxandra Teslo · Launched 3 years ago

Pro-progress, anti-safetyst commentry on Scientific & Cultural matters from a Genomics PhD. Attempting to shape culture towards optimism, an abundance mindset & freedom.

Subscribe

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).



614 Likes • 135 Restacks

← Previous

Discussion about this post

Comments

Restacks



Write a comment...



Freddie deBoer  Jul 21

...

 Liked by Ruxandra Teslo

Very useful piece. This is a point I've been trying to make for a long while: many problems are not solvable by intelligence!

 LIKE (28)  REPLY

 SHARE

20 replies by Ruxandra Teslo and others



Elizabeth McDaniel Jul 21



♥ Liked by Ruxandra Teslo

Fantastic piece! I don't know how you managed to fit in the difficulty of clinical trials and the culture of San Francisco into one coherent argument, but I really enjoyed reading this.

♡ LIKE (13) 💬 REPLY

🔗 SHARE

1 reply by Ruxandra Teslo

112 more comments...