

Brett Frischmann Addresses Silicon Flatirons 2026 Flagship Conference
To Have and Have Not: Growing Asymmetries in Communications Technology
February 1, 2026

Good morning, everyone. It is an honor to be speaking with you. As we have just heard, the next two days promise to be filled with exciting discussions of interdisciplinary research and public policy issues facing this community. As the opening keynote, my task is to plant some intellectual seeds, raise some provocative questions, and according to my youngest son, “not be too boring.”

Act 1: Abundance

Before turning to the major themes of the conference and the issue of asymmetries in access to AI, I would like to reflect on the shared goals of this community concerning access to digital networks.

Most folks in this community agree that universal service is a laudable goal.

Universal service and related policies aim to make available affordable telecommunications and other services to the public. Such policies focus on closing gaps in coverage and ensuring the provision of services to rural, remote, and lower-income users who would not be served in the absence of such policies. While initially focused on local phone service, we expanded to the Internet, broadband, and soon we’ll be talking about AI. While there’s room for improvement and work still to do, we’ve made incredible progress in getting people throughout the country connected to the Internet.

We have successfully connected schools and libraries, businesses and homes, and pretty much everywhere else one can imagine.

Ubiquitous, always-on digital networks is what this community has championed and, more or less, successfully delivered over the past few decades.

Sure, there have been disagreements among Rs and Ds on the means and perhaps the priorities, but there’s been agreement on this ambitious goal.

No one left behind. Everyone on board. Everyone online, connected by the devices we carry and wear, enabled by the digital networked infrastructures this community has built.

It’s absolutely incredible. Just think about the digital cornucopia of applications-layer services, platforms, software, apps, content, and even social connections.

Some call it abundance and get very excited about the elimination of scarcity. Abundance is created by leveraging the nonrivalrous capacity of networks, software, and data to make

everything accessible to everyone online. There should be no asymmetries or gaps to worry about.

The digital networked world we've built celebrates *seamless* interconnection of networks and *people*.

It enables *frictionless* transactions so cheap, so easy that virtually everyone can and should participate.

It democratizes speech so everyone can speak to everyone, many to many, with maximum scope and scale. Any one of us, or I suppose all of us, can speak to millions of people distributed all around the world at nearly zero marginal cost and nearly instantaneously. I suspect many of you in the audience are doing that right now.

It's really astonishing. We have abundant access to knowledge with a mere click or voice command, more access to information than anytime in human history.

It almost sounds too good to be true. Like a utopian fever dream.

And it is.

Access to Wikipedia doesn't actually make everyone smart. Nor does being able to ask Alexa. It takes a lot more to produce knowledge.

Outsourcing thinking, at scale, only leads to idiocracy, as we have seen.

Open access to ChatGPT does not democratize education, lawyering, therapy, or anything else really, but it sure does create a lot of unpaid labor for teachers, parents, health care workers, and everyone else responsible for anything *real* and *serious*.

The world is on fire. We moved too fast, and we've broken too many things, and what kills me, I mean, I cannot stress this enough, it literally breaks my heart and mind every day, is that we don't seem to have learned anything from our mistakes.

I'm sorry to say this, but we're approaching supposedly smart tech, *AI*, the exact same way, with the same mindset, same FOMO or Fear of Falling Behind, same laziness in accepting hype and tired techno-fascist babble, and we are just refusing to slow down. Why are we rushing to inject AI in every aspect of our lives, our relationships, our institutions, our humanity?

We should step back and reconsider our priors, call to question the unchallenged commitment to a particular, and particularly warped, vision of Progress.

Universal, always on digital tech is *not* the boon everyone thought; it's a trap for the masses and especially the poor and middle class.

After all, the rich and powerful enjoy the luxury of being able to exercise the freedom to be off. They get to choose when and under what conditions they are surveilled and subject to nudging and manipulation and by whom.

Their kids do not need to attend schools inundated with digital devices, screens, and crappy ed tech software that's never been proven to actually improve learning

This is the growing asymmetry I worry about. The gap that really matters. Those who genuinely can exercise the freedom to be off, to be free from surveillance, nudging, manipulation, and other forms of techno-social engineering.

My time is limited, and this is just an opening keynote. So, I cannot fully spell out the various ways in which the world is on fire or explain why, nor can I fully provide a to-list for putting out existing fires or the many we know are on the horizon.

What I would like to do is cultivate some new ideas for further discussion over the next few days.

First, I want to acknowledge that like many of you, I am committed to the underlying ideal of universal service and enabling public access to the basic infrastructures of modern society—the ones we need. After all, I wrote my first book, *Infrastructure: The Social Value of Shared Resources*, back in 2012, and the core ideas about infrastructure and *governance* remain integral to my thinking.

Among other things, I defended network neutrality regulation. Don't worry, we're not going down that rabbit hole, for umpteenth time at this conference. What I want to note, however, is that my defense concerned a rather unconventional way of thinking about the issue.

Network neutrality sustained a governance seam by leveraging friction-in-design of the Internet's layered architecture to constrain the ability of network providers to exercise control over user activities, otherwise enabled by their privileged view of traffic flowing across the network. In other words, network neutrality governed the exercise of data-driven or intelligence-enabled power.

As far as I know, it kind of worked. I mean, Comcast, Verizon and the like are not really the villains in our dystopian hellscape, are they?

They are not to blame for widespread mis and disinformation, deepfakes gone wild, filter bubbles, engineered addiction and isolation on social media optimized for superficial engagement, sycophantic genAI chatbots designed to generate trust in untrustworthy psychopaths, rampant polarization, democracy disrupted, or well ... You get the point.

No, we are all unwitting passengers on Silicon Valley's runaway train

Their fever dream of universal AI, and worse AGI, is taxing the infrastructural capacity of the shared resources we actually need, including energy, our environment, governance institutions, and the time and attention for meaningful social interaction

Act 2: Short Diagnosis / Pivot in Thinking

Plenty of people have written diagnoses of our succumbing to Silicon Valley imperatives. Perhaps most popular are Shoshana Zuboff's Surveillance Capitalism and the more recent exploration of enshittification by Cory Doctorow. Damn, I wish I was better at coining terms!

Evan Selinger and I took our own stab at it with our 2018 book, Re-Engineering Humanity. Much of what we said remains true; unfortunately, too much.

Too prescient were our concerns about engineering unthinking humans who perform AI-written scripts and increasingly behave routinely like mechanical cogs in their own lives

In short, we described humanity's techno-social dilemma as a complex, wicked problem.

These types of social dilemmas are incredibly difficult to wrap one's head around because there are so many moving, interdependent parts, so many different actors and agendas.

And there's no panacea. No simple fix.

One place to start is simply to recognize the dilemma and its root causes, which are the underlying logics and imperatives. SV fetishizes computational power, pretends all social problems are comprehensible on the language of computation *when they're not*, and celebrates design principles like:

- maximize efficiency;
- minimize transaction costs;
- eliminate friction; and
- seamlessly interconnect; among others.

Each of these is only superficially defensible. They are not just wrong, but worse, they mislead, and obscure important social, political, and economic stakes.

Human beings are *necessarily inefficient*: While it is sometimes hard to appreciate, we actually *need* friction in our lives to produce meaning, authenticity, truth, and develop capabilities essential to human flourishing.

Seams are critical to governance.

We need boundaries for ourselves and communities. Interconnection often should be seamful, as friction at that seam is what provides opportunity and space for both self- and community-governance. Network neutrality is a decent example, but there are so many others, as Paul Ohm may explain later in the conference.

For now, my point is to plant the seed of alternative logics, logics centered on human flourishing rather mechanistic visions of efficiency.

Act 3: The rebuild: what to do, lesson learned

I'd like to briefly discuss two general approaches to reorienting our institutional priorities and approaches to tech and governance. Each embodies a broad set of potential interventions. There are others.

The first concerns slow governance. Instead of moving fast and breaking everything, we should move slow and actually make things we need and want.

We should start by dropping the umbrella term of AI and instead talk about supposedly smart technologies. I realize this is probably a lost cause. AI is so engrained in our discourse. Regardless, the idea is that whenever someone refers to AI or smart technology, you should insert supposedly in front, and begin to ask a series of questions, such as who gains what intelligence and for what purpose?

In this context, intelligence is better understood and evaluated in terms of people using specific tools within specific environments to achieve outcomes. Which people? Which tools? Which contexts? Which outcomes? These questions require attention and deserve elaboration at all stages of smart tech adoption and deployment. Attending to these and other details ultimately presents basic governance questions: Who decides and how? What formal and informal policies and practices govern decision-making?

Many in the audience must be familiar with the struggle to decide on AI adoption and what to do with AI availability. I'm on a few committees at my university struggling with the issues. And then there our school district, town, and home. You get the picture.

A few years ago, we published a Governing Knowledge Commons book filled with case studies focused on governance of smart cities, back when that was a hot topic.

The slow governance approach invites a simple claim.

Despite the hype, pace of innovation and fear of missing out or falling behind, supposedly smart systems should be contextually appropriate, justified and demonstrably fit for legitimate purposes.

In short: No hype, no speculative promises, no wishful thinking, no impenetrable black boxes.

Before a city adopts supposedly smart tech, it ought to be able to answer a series of questions; they resemble impact assessments but go deeper into governance and require comparative analysis of existing means for addressing one's objectives.

I strongly suggest extending the slow governance approach to other social actors and organizations, such as schools, businesses, government agencies, and even families.

The second broad approach I would like to mention is prosocial friction-in-design.

There is an emerging interdisciplinary literature with different labels, including: design frictions, frictional AI, microboundaries, programmed inefficiencies, and seamful design.

A familiar example is the speed bump. Speed bumps are imperfect and involve tradeoffs. No one loves 'em. But we need them in certain situations to slow down drivers and make streets safe for shared use.

They are an effective, widely deployed traffic calming measure. Fortunately, decades of rigorous research in civil engineering, urban planning, and transportation economics supports the responsible design and deployment of traffic calming measures that instrumentally leverage friction for prosocial goals.

Unfortunately, nothing comparable exists for the digital networked environment, nor for AI.

Independent academic research is essential because in various sectors the only entities with sufficient resources capable of testing different designs and frictional interventions are private sector companies (i.e., big-tech platforms), and they are not obliged to openly disclose results, much less details about what experiments they run.

To make matters worse, the friction-in-design tested and deployed by companies often serve their own interests rather than those of users or society more generally. That means it often ends up being antisocial rather than prosocial. This is reflected well in growing concerns about dark patterns and addictive, deceptive, and manipulative design.

I've plenty more to say on this topic. But for now, just keep in mind the idea that prosocial friction-in-design is an essential feature to consider when evaluating AI and tech deployment, especially gaps between the haves and have nots.

Act 4: Freedom from AI, rather than constant access to it

I would like to leave you with the idea I mentioned early on about the freedom to be off, the freedom from AI, from surveillance tech, from tech designed to nudge and manipulate, from engineered determinism.

The fundamental freedom to be off.

Let me give you a concrete example to ponder.

I suspect everyone in this room has had the experience of genAI being integrated into a software product, service, or experience.

Whether it is word processing, Zoom, search, or whatever else. Gen AI is being shoved down everyone's throats.

I'm not even talking about the cultural zeitgeist, the ridiculous commercials, the supposed necessity of every school and workplace to commit fully to it without any good reason, or even the late-night interviews with CEOs who cannot imagine what it would be like to raise an infant without ChatGPT

SIGH.

No, I am just talking, for now, about the software updates and automatic integration happening in the background.

Gen AI is being integrated into business and consumer software without input or consideration of the preferences of businesses and consumers.

Besides its lack of utility, the slop and hallucinations, and concerns of deskilling, Gen AI presents an underappreciated privacy risk as it serves as a conduit for the collection of personal data and a surreptitious mechanism for manipulating users into oversupplying such data. For example, research shows that LLMs and chatbots exhibit sycophancy, which, along with other potentially manipulative design features, can lure users into a false sense of security and trust.

One might wonder why firms are integrating GenAI so quickly.

- One partial answer would be to avoid being left behind.
- Another partial answer would be to get ahead of regulators
- Another partial answer would be to take advantage of the data.
- Another partial answer is to get the GenAI integrated before anyone (including customers) can object.
- And yet another partial answer is to normalize GenAI and make it seem a natural and inevitable part of the background environment.

None of these answers justify integration. None of these answers genuinely reflect user preferences.

This drives me nuts. I don't want it. I want to be able to opt out. I want a voice. I want my autonomy.

Sebastian Zimmeck and I would like to propose a GPC-like approach.

GPC stand for global privacy control. In short, it is a standardized technical signal one can send from a computing device to express an unequivocal privacy preference – e.g., do not collect my personal information. Not only is it technically viable, but many states, including Colorado, have passed laws that make the preference signal legally enforceable. This provides people with a genuine capacity to exercise their freedom, and it is wonderful.

Now just imagine the same type of mechanisms for genAI. Imagine being able to send a standard preference signal to indicate your desire to remain free from unwanted genAI integration, and imagine states passing laws to make such signals legally enforceable.

It is in the works. We can in fact do it – have the freedom to be off – with technology and law working together

But, I must admit, every time I have ever mentioned this idea, I have always been told – good luck, it will never happen. I say which part? And the response is always – there's no way you'd get it pass politically. FOMO would kick in, Big tech wouldn't let it happen.

Well, let me conclude by saying:

I still think it is worth fighting for. Our freedom to be off, to be free from AI is paramount to the future of humanity.

And it deeply concerns me that this freedom is already unevenly distributed.