Opening Remarks — Bellagio Convening, April 28 2026

The Missing Mechanisms of the AI Economy

Tim O'Reilly — Opening remarks, Bellagio Convening on Market Design and the Agentic Economy

April 28, 2026

You've probably received a number of communications from us that explained a little of what this is about, but I want to give you some history and perspective on why we convened this meeting.

Over the course of my career, I've periodically noticed something a little curious to me, a little bit like the curious incident of the dog in the night-time in the Sherlock Holmes story.

Back in 1998, for example, I noticed that the Free Software Foundation never talked about the internet. They never talked about the World Wide Web, despite the fact that all of these things had been built with free software. The Web had been put into the public domain. You can't get freer than that. They didn't talk about the X Window System. They didn't talk about the roots of the internet, the TCP/IP stack that came out of Berkeley Unix. All they talked about was the GNU project. They drew the map with licenses, and I went, there's something missing here.

So I pulled together a group, which I originally called the Freeware Summit. It was the people who had built a lot of this internet software, things like sendmail, which was under the BSD license; DNS, which was under the BSD license; and Apache. I brought all these people together in a room. I didn't know what was going to happen, but I think it was a critical step: just get the people together and see what emerges.

In the course of that meeting, Linus Torvalds said, “I didn’t realize there were two different meanings of free in English.” It also turned out that Eric Raymond, the author of *The Cathedral and the Bazaar*, had just been in a meeting two weeks earlier where Christine Peterson had come up with the name “open source.” The group debated it. Michael Tiemann, from Cygnus Solutions, the company that had written GCC, was arguing for “sourceware.” We debated, we voted, and at the end of the day we held a press conference.

By that point the man who had been the director of activism for the Sierra Club had become my VP of marketing, and he had taught me a lot about activism. I’d spent the early ’90s in the throes of the idea that the internet could be a commercial activity. Some of you are old enough to remember that the internet was under non-commercial licensing then, and commercial activity on it was forbidden. There was an active debate group and mailing list called com-priv, for commercialization and privatization of the internet. My company created the first commercial website, and it got special dispensation from the National Science Foundation to do that. So I’d learned how you market these ideas.

At the press conference I told the story that this Free Software Movement was not radically anti-Microsoft, not radically anti-commercial-proprietary-software. It was a movement that every company in the room depended on. I had memorized the IP addresses of some of the reporters’ publications. I said, the reason you get to be NewYorkTimes.com or WallStreetJournal.com is because of this guy over here. The reason you can send email is because of this other guy over here. I went down the story, and within a week the new name “open source” had taken over. Linus Torvalds was on the cover of Forbes. We had a real impact with that one. But really what we did was create a new map of the world. Here’s a way to think about the problem. Here’s a way to think about the situation.

That was probably the most successful time I’ve done this, but I did another one in 2007 with Larry Lessig that produced the Eight Principles of Open Government Data and kicked off the open data movement.

This particular meeting grew out of a similar observation. It was triggered in a lot of ways by a conversation with Ronnie Chatterji, where he was talking about what he had done as chief economist at OpenAI. I heard “chief economist” and I thought of Hal

Varian at Google, working on the ad auction, thinking about the economics of the system Google was building. I discovered that Ronnie's job was, instead, just to study the impact of AI on the economy. And I went, that's a gap. There's something wrong here, when the chief economists at the AI companies aren't part of the product team. They're not working with product. They're working in what is essentially a public-policy function. That struck me as an odd gap, in the same way that the absence of the internet from the Free Software Movement had struck me as an odd gap. The curious incident of the dog in the night-time.

So I thought, let's try to have a conversation about the intersection of economics and technology. How do we start to understand that what we're building, when we build a technological system, is an economy?

I got fascinated with mechanism design when Jonathan Hall, who was the chief economist at Uber, waved Al Roth's book at me. It's called *Who Gets What and Why*, a popularization of his work in mechanism design. Roth got the Nobel Prize for redesigning the kidney transplant marketplace, among other things. There are several Nobel Prizes in this field. It's about how you align the incentives of people in new ways, in cases where they otherwise wouldn't be able to transact. In some ways it's the design of more efficient markets.

When I heard that, I immediately saw all the things I had experienced in my career through that lens. The way I think about this isn't quite how economists think about it. Nicole is going to give us a slightly different view, and I'm starting from Erik Brynjolfsson, who is an economist. But let me tell you how I'm thinking about it, and why I frame this as the missing mechanisms of the AI economy.

It seems to me that every revolution in technology depends on new economic mechanisms.

I didn't think of these things in those terms originally. I didn't even learn about mechanism design until decades after I had talked about the architecture of participation as the central design fact of open source. I was very influenced by the design of Unix, which was a protocol-centric operating system rather than a tightly integrated ball of code. The idea was that you'd have a microkernel, and a bunch of utilities that all followed the same rules. That, it seemed to me, was what truly enabled

open source. You had all these people who could each write a piece, and because each piece played by the same rules, it could be added to the system. It was a kind of mechanism design. The internet software development economy was a product of the architecture of the system that allowed people to share code.

Some of that was actually accidental. Unlike the PC, where commodity hardware became standardized and free software was distributed in binary form, Unix happened to be a common operating system that ran on radically different hardware architectures. So you had to distribute source code that people could port to the underlying hardware. It was completely accidental, but it led to a radically different architecture for software development.

I see the same pattern repeating with AI. The initial framing of the AI race that came from OpenAI was implicitly based on a model similar to Microsoft's dominance of the PC era, or to iPhone and Android in mobile. There's a dominant provider, they define the rules, everybody calls their APIs, and if you can get to dominant share, you become the platform and everybody depends on you. A hub-and-spoke architecture.

Then along comes MCP. I don't think it was a strategic decision on the part of Anthropic. I think it was a side project that came into focus and took off, and it basically recreated the architecture of participation. That's why MCP is maybe a little over-represented here relative to other things. It seems to me that it was the beginning of a new architecture for AI, which has led to the agentic world we're now building.

Simply having that architecture isn't enough. The Web had this architecture, but there wasn't really an economic model. The economic model got invented somewhere off on the side: pay-per-click advertising. Pay-per-click totally changed the shape of the market online. It was based on the recognition that instead of simply showing people ads and not knowing who saw what, you could build an ad model where there was real knowledge. Of course, that's when the actual economists got involved in the mechanism design. The first version of pay-per-click was just sell ads to the highest bidder, which was not a very efficient auction, and Hal Varian and others at Google figured out how to make it better.

All of which is to say, we are at a point in this AI agentic economy where we are building a lot of mechanisms, but I don't think we are thinking about them as economic design

mechanisms. We are thinking about them as technical mechanisms. Whether we think about it or not, though, they are designing the economy. And we can design that economy in ways that produce immensely different effects. It can be a winner-takes-all economy. It can be a winner-takes-most economy. Or it can be a truly participatory economy. The question is, what mechanisms might help that, and what's missing?

We know that certain classes of people, and certain classes of entities, are saying: we want attribution, we want compensation. What does that look like? How might the future agentic economy actually be an economy?

It will get there sooner or later, because if it doesn't, the whole thing will fall down. We are already seeing strains. The idea that you'll have a few centralized providers and you'll just call their APIs is starting to come apart, because they are starting to raise their prices. Supply and demand means you have to either choke off demand or raise price. There is an alternative, which all of you are exploring, which is more decentralized. You go, okay, I can run that on my local machine. But then, what mechanisms are there for deciding when you do that? How do we make the whole thing more token-efficient? Token efficiency may be one of the breakthrough mechanisms. The same may be true for skills. You start thinking about agent skill marketplaces, and what those are going to look like.

Going back to my layman's rethinking of mechanism design: organic search is a kind of mechanism design. Starting with PageRank and going through all the other signals Google accumulated, we found out how to do matching in a market without money. The most amazing thing about Google search is that it's a market of trillions of transactions in which price plays no role. There are all these other signals instead. To me, an amazing example of mechanism design.

Now imagine trillions of agents going into billions of different knowledge domains, trying to figure out, what do I need to know to get the best results working in this area? That is going to be an economy. What's missing to make it one?

We are not going to fix it in this room. We are not going to come up with even the barest beginning. But we can often set the pattern. One of the examples I think of is Rob McCool's invention of CGI in 1993, which was the first idea that you could have a database behind one of these web pages, going beyond what Tim Berners-Lee had

created. Nobody uses CGI anymore. It went through a lot of contention between Microsoft and Netscape, with their competing APIs. We eventually built higher-level frameworks, and even languages like PHP, which was an attempt to do it at that level. We eventually figured out how to have the rich landscape of database-backed sites that can do all kinds of interesting things. The mechanisms we talk about today are not likely to be what happens in the future, but we can identify some of the areas where something will need to exist to fill a problem, and we can see hidden adjacencies.

Andrew Trask, who was going to be here but got ill at the last minute, made a comment in one of our recent conversations that I think captures the kind of moment I'm hoping for. He works a lot on privacy-preserving auditing. He said, "I read this ML paper from somebody at Apple, and I realized that attribution is the other side of privacy." He was using many of the same techniques I'm using in my work. That notion of hidden adjacencies is another piece of this. You're working on something, and somebody else is working on something else, and aha, we're actually working on pieces that belong together.

We've tried to curate a group that could potentially talk to this. We are really hoping that as you engage and share your work, insights will come out of that to allow each of you, maybe subtly, maybe deeply, to change what you do.

My goal in my work is to make the future better than it would be without our interventions. All through my career at O'Reilly, our guiding principle has been technology transfer from the cutting edge to the people who want to follow, but also: how do we make a better future? A lot of times, companies don't think about that as much as they could.

Coming back to my layman's reimagining of mechanism design, it's about how we build a market that is actually better. When Al Roth redesigned the kidney marketplace, it was because he realized there was a lot of unmet demand. I'm willing to give you my kidney because I love you, but I don't have the right blood type. How do you create long chains of caring, where somebody who you'll never know donates a kidney, and you, in return, donate yours to someone you'll never know? That's a vast increase in welfare.

For all the things the internet has done wrong, there has been a lot of that kind of mechanism design. There have been a lot of increases in welfare. Some of those

increases are going to happen naturally, but some can be created. And some can be created and then turned away from. This is the part of it I don't really know how to do. But if you think about Cory Doctorow's concept of enshittification, you can see that with the vast increase in the efficiency of information markets that came with organic search, it turned out that Google and Amazon, who were probably the two companies in the world that were best at that, both realized they could make more money by demoting their own organic results in favor of advertising.

There's a challenge there for us. Are there any ways, as we think about this design problem for the future agentic world, of making it a little bit more robust against those kinds of attempts to turn away from welfare advances that could otherwise be there?

That's where Ilan and I started this project called the AI Disclosures Project. It was rooted in the idea of reporting on operational metrics, and what would be the triggers in regulation or legislation for how AI might be managed. We were thinking, we've mandated financial reporting, but we've never mandated any kind of reporting on operational metrics. So when Google or Meta or Amazon decided to optimize for revenue, the specific things they did, the specific things they measured, the specific things they tasked their company or their employees to optimize for, none of that is reported. That was sort of our original insight.

We realized, for example, that there's something in the US called segment reporting. Companies are required to report out any part of their business that is more than 10% of their total profit, or revenue, or costs. Then you look at something like Google Maps, and you go, it's a rounding error for Google. But it has two or three billion users worldwide. That's incredible. So we said, maybe one of the interventions is, you make segment reporting based on number of users. That's a small mechanism design insight about how you might make the system slightly more robust.

In agent commerce, I think about how the corporation, which is itself a kind of slow primitive AI, is supposed to be registered. But that registration has become opaque. You have these long chains, and you can't figure out who really did something. Maybe we could fix that in agentic protocols. Jeffrey, you were just talking about something along these lines. You get better service if you're authenticated. That's one of your ideas from fifteen years ago that maybe should be brought into the protocols we're building today.

So that's part of my point. Defensive robustness is something we might want to be thinking about as well as we go about our work after this meeting.

I think that's pretty much all I have to say. I'm having such a good time already just hearing all of you talk and being a fly on the wall. As each of you presents your work to the others, I'm hoping it creates new connections and new insights into how we can build this agentic economy to be richer and more participatory, for humans as well as agents.

I'm not quite done. I have a profound sense that, even though we are building this thing for AI, AI and humans are deeply connected, at least for the foreseeable future. As Henry Farrell and others have said, AI is a social and cultural technology. It's not really an independent intelligence. It's part of an ongoing process that has been part of our growth as a species: the ability to communicate, to transact, to share, to learn from each other. It is the most powerful tool we have yet for building a collective intelligence made out of all of us. And it needs to return something to all of us.

So thinking about how we build the circulatory economy of AI, rather than an extractive economy. It does seem to me that in the first stage of AI, it has looked a lot like colonialism. The economic model has been that we, the AI labs, have a better use for all of your content than you do. We have the right to it. We will take it. We will turn it into valuable products for our users. And those of you who produced it, you are a resource to be exploited.

I don't think that's a sustainable position. I think ultimately we need a circulatory economy that returns value to human creators. I don't know what it looks like in the future. I don't know what it looks like even now. But if we don't find a new circulatory economy, it will eventually be a less productive economy than one that enables everyone to participate.

So with that, I think we should go on to the rest of the program. Thank you. Unless anybody has any questions or comments.