

Architectures of Participation

How collaborative open-source software development increased the velocity of problem-solving in cloud computing—and what that suggests for innovation policy.


Silicon Valley's dynamism during the final three decades of the twentieth century highlighted the singular importance of social and professional networks to innovation. Since that time, contemporary and historical case studies have corroborated the link between networks and the pace of technological change. These studies have shown that networks of networks, or ecosystems, that are characterized by a mix of collaboration and competition, can accelerate learning and problem-solving.

However, these insights about networks, collaboration, and ecosystems remain surprisingly absent from public debates about science and technology policy. Since the end of World War II, innovation policy has targeted economic inputs such as funding for basic scientific research and a highly skilled workforce (via education, training, and/or immigration), as well as support for commercialization of technology, investments in information technology, and free trade. Work on national systems of innovation, by contrast, seeks to define the optimal ensembles of institutions and policies. Alternatively, policy attention is focused on achieving efficiencies and scale by gaining control over value chains, especially in critical industries such as semiconductors. Antitrust advocates have attributed stalled technological innovation to monopolistic concentration among large firms, arguing that divestiture or regulation is necessary to reinvigorate competition and speed gains for society. These approaches ignore the lessons of network research, potentially threatening the very ecosystems

that could unlock competitive advantages. For example, attempts to strengthen value chains risk cutting producers off from global networks, leaving them vulnerable to shifting markets and technology and weakening the wider ecosystem. Breaking up large platform firms may likewise undermine less visible internal interdependencies that support innovation, while doing nothing to encourage external collaboration.

How might the public sector promote and strengthen important network connections in a world of continuous flux? This essay reexamines innovation policy through the lens of the current era of cloud computing, arguing that the public sector has a regulatory role as well as a nurturing one to play in fostering innovation ecosystems. Since traditional ways of conceptualizing antitrust regulations are unlikely to be effective in today's complex global innovation ecosystem, we argue that a policy agenda drawing on elements of industrial policy, as well as reconfigured competition policy, can help ensure that the organizational structures of complex technological projects balance competition with collaboration to foster, rather than stifle, innovation. What we propose is not new. In the early twentieth century, the progressive reformer Louis Brandeis worked with engineers, trade associations, companies, and government regulators to configure antitrust law so that it channeled competition from predation to innovation.

We think these lessons are relevant today and can inform a new suite of policy ideas centered around participation



in today's decentralized ecosystems. Furthermore, understanding the architecture of these technology ecosystems suggests policy tools to accelerate innovation and improve governance—while providing lessons that can guide strategies to enhance public benefit in the future.

The nature of competition in the cloud

Competition in cloud computing is currently focused on building an interconnected infrastructure that simplifies management of very large volumes of data. As the third wave of innovation in the commercial internet, the market structure of this phase differs significantly from previous ones. The first generation of internet innovation during the 2000s was led by engineers' and start-ups' development of networking standards and protocols to support a globally accessible internet. The second wave began in the 2010s when the focus shifted to scaling the computational capabilities of this network and building out complex software systems and platforms—a market that was eventually dominated by large firms. During this second wave, cloud computing services emerged alongside institutions enabling distributed collaboration in the development of open-source software. Today, as information storage, computation, and software continue to shift away from private servers to the public cloud, engineers in both large firms and start-ups are building the elements of a modern data infrastructure for the cloud. The goal in this third wave is a platform that facilitates data management and ultimately makes data more widely accessible.

The market structure of this era of cloud innovation is a complex combination of networks of engineers in start-ups, established technology firms, and nonprofit foundations. All are experimenting with the elements of a distributed data infrastructure that will support the collection, storage, transformation, analysis, and movement of data in and between clouds, enabling what is likely to be a fourth wave of innovation in which nonprofessionals, as well as professionals, have unprecedented and unfettered access to sophisticated data analysis and widespread application of machine learning.

The multifaceted structure and diverse possibilities characteristic of the cloud's ecosystem do not fit neatly into boxes labeled competitive and monopolistic or open and closed markets. Most recently, this simplification has generated a bifurcated public debate between two camps with opposing views on platform regulation. On one hand, some argue that innovation is a product of competitive markets that allow entrepreneurial entry and therefore recommend antitrust policy to constrain the market power of large firms like Amazon, Google, and Facebook. Opponents of this approach argue that the large firms have used their ample resources and scale to generate ongoing innovations that benefit customers and even start-ups. In this view, increased

regulation would only hinder progress.

To gain insights into the organizational conditions for innovation and its implications for policy, we spent two years interviewing software developers, attorneys, entrepreneurs, foundation executives, and managers working on data transformation in the cloud. We found partial support for both views of innovation: today, progress is coming from *both* large firms and a new generation of start-ups. However, the evidence suggests that neither of the two policy prescriptions is appropriate. Attributing innovation to either the free play of competitive markets or the capabilities of the large platform firms overlooks the power of collaborative *ecosystems* that increase the pace and quality of technological change.

Cloud innovation is currently at a crossroads, with two possible organizational trajectories. One trajectory is based on a top-down and centralized model, with platform firms exercising power over start-ups; the alternative is more decentralized and open, with the large firms collaborating with start-ups as well as nonprofit institutions. This second, more open trajectory is supported by extended, cross-firm networks of developers and appears to accelerate innovation.

Our research suggests that competition policy, innovation policy, and industrial policy should be seen as complementary, particularly for supporting today's collaborative ecosystems. Moving beyond the old categories allows us to define a vision for policy that deliberately reinforces the dynamism we see in the cloud and to consider how to extend that model to other industries.

A short overview of cloud innovation

Today's cloud has evolved on top of legacy structures that still influence its development. In the early 2000s, businesses purchased software and ran it on their own servers, while storing information and data on-site. In 2006, Amazon Web Services (AWS) began offering cloud storage and computing services that freed businesses from the demands of managing physical servers or running large, licensed software applications on their own machines. Other infrastructure providers like Microsoft Azure and Google Cloud soon joined the competition, supporting the rapid growth in the 2010s of cloud-delivery of software as a service for a wide range of businesses and consumer applications.

However, as late as 2015, data remained locked in proprietary and incompatible corporate systems, making it extremely costly for firms' employees to move, share, or recombine even their own data. Traditional data warehouse systems require expensive on-premises hardware, which means that data is maintained in proprietary formats and managed and processed by a centralized IT department. As the internet enabled an immense increase in the volume, velocity, and variety of data, these centralized systems could

not keep up. Over the past decade, the constraints of these systems have inspired widespread experimentation, including a proliferation of start-ups building new tools and data formats to enable data storage and processing in the cloud.

The shift to a cloud services model has also contributed to a renaissance of open-source software. Open source, historically seen as a fringe movement of hackers opposed to proprietary software, is now in the mainstream of software development. The principles of open-source software haven't changed: the code can be accessed, used, modified, and distributed, commercially or noncommercially, by anyone under the terms of the license. But far from being a peripheral option, open-source technology is now widely adopted by firms in all sectors of the economy. For example, the Linux operating system, which originated with programmer Linus Torvalds in the early 1990s, today runs most of the internet as well the world's supercomputers and stock exchanges.

The ecosystem of cloud innovation thus has contradictory features: it is both competitive and collaborative, decentralizing and centralizing. It has been a boon to entrepreneurship, triggering a veritable "Cambrian explosion" of new data-related firms over the past decade. At the same

Importantly, the movement no longer sees itself in opposition to market-based systems, and technology corporations have embraced it. Microsoft reversed its vehement opposition to open source to become its largest contributor in 1999, and IBM purchased Red Hat, the leading open-source consulting firm, in 2019. Google is a close collaborator with the Linux Foundation. Today, open-source software is increasingly developed by, and integrated into, commercial enterprises. Open-source producers now generate significant revenues by selling not just support and services, but also proprietary or enhanced functionality or open-source tools. Venture capital firms now actively invest in open-source start-ups (including some \$2 billion in 2019 alone) that compete aggressively against firms selling proprietary software, and analysts report a recent boom in open-source IPOs.

Perhaps most significantly for policy, innovation in the cloud is populated by a host of firms and institutions committed to open source. These institutions include nonprofit open-source organizations such as the Linux Foundation, the Apache Software Foundation, the Mozilla Foundation, the Eclipse Foundation, and the OpenStack Foundation. Once primarily repositories for code, these

Attributing innovation to either the free play of competitive markets or the capabilities of the large platform firms overlooks the power of collaborative ecosystems that increase the pace and quality of technological change.

time, data innovation in the cloud has empowered the cloud divisions of the biggest platform companies, including AWS, Microsoft Azure, and Google Cloud.

But the distinction between big and small players is less important than the contradictory trends—toward centralization and closure, on one hand, and toward decentralization and openness, on the other—that are shaping today's cloud ecosystem. The move toward closure can be seen in the application programming interfaces (APIs) that once facilitated open experimentation and development across different platforms, firms, and products, reflecting the distributed innovation of the early internet. Today's giant platform companies have abandoned that openness and diversity by restricting access to their APIs. In their drive to gain market share, they have also acquired promising start-ups and developed proprietary systems that limit the ability of newcomers to build on their platforms. Thus, market concentration and declining openness and interoperability became complementary.

At the same time, the market architecture of the cloud has elevated the open-source movement, which offers a model of increasing openness, even as it is institutionalized and integrated into the portfolios of the largest commercial firms.

foundations are now well-funded, professionally staffed promoters of their development model. The Linux Foundation, for example, is no longer focused solely on the Linux operating system. Instead, it is dedicated to helping "open technology projects build world class open-source software, communities, and companies." Unified by a shared recognition of the value of open-source development, the Linux Foundation community today includes more than 19,000 contributing companies and more than 540,000 developers. In short, global networks of open-source developers, communities, and foundations are central actors in today's innovation ecosystems.

Litigating "strip-mining"

How are policymakers to understand this new ecosystem, with its many players and evolving business models, in order to create forward-looking governance that encourages innovation? To date, discussions of policy have largely been confined to antitrust litigation. Exploring this impending litigation can show how it both reflects and elides the true complexity of the cloud's innovation ecosystem, which cannot be accurately characterized by concepts of monopoly and competition.

One stream of antitrust litigation is focused on Amazon Web Services (AWS), which controls 33% of the \$178 billion global market for on-demand cloud computing. AWS's business model requires massive investments in data centers located around the world to share the workload of data storage, computing power, and networking. AWS, along with other leading cloud providers, also offers platform (databases, web services, development tools) and software services on top of the basic infrastructure. In short, AWS has made solving its clients' hardest IT problems easy, and as it has gathered more customers, it has gotten even better at solving those problems. In 2020, AWS accounted for nearly 67% of Amazon's operating profits, and Amazon increasingly sees itself as a technology company rather than a retailer. By continually improving the quality and performance of its cloud services, AWS provides an important benefit that has strengthened the entire tech ecosystem.

If the growth of cloud computing has benefitted businesses, it has a more complex relationship with the open-source software community. AWS, like other cloud platforms, makes use of open-source code, including the Linux operating system, and has been a powerful driver of

market power before the House Subcommittee on Antitrust, Commercial, and Administrative Law. Advocating aggressive antitrust regulation, these companies charge cloud providers with erecting barriers to entry by making it impossible for independent firms to compete. In the coming years, this battle will be fought in the courts through private and public lawsuits, in Congress, and by states' attorney generals. Meanwhile, in an effort to prevent AWS from building commercial services from their code in the future, at least eight open-source database companies, including Elastic, have modified their licenses, making them so restrictive that they are no longer considered open-source by the community. In seeking to shift power from AWS, these efforts could diminish the vibrancy of the open-source innovation ecosystem because it will leave the centralized model intact.

Understanding architectures of participation

Many developers who are committed to open-source software argue that resorting to litigation and adopting restrictive licenses will hurt the community by further centralizing control, reducing adoption of open-source software, and ultimately harming end users. Developing

Perhaps most significantly for policy, innovation in the cloud is populated by a host of firms and institutions committed to open source.

its adoption. However, AWS's primary focus is on increasing its customer base by achieving scale and perfecting internal competency, which both serves and threatens the entrepreneurial technology ecosystem. For AWS, innovative open-source software offers a ready path to expansion. In 2015, for example, AWS copied the open-source code for a pioneering search engine named Elasticsearch and integrated it into its proprietary cloud services offerings. Reportedly, AWS was soon making more money from the code than the software's creator, Elastic. Critics have charged AWS with "strip-mining" the open-source code that smaller companies have invested heavily in, making it harder for them to make money. AWS has countered that it is a "significant contributor and supporter of the open-source community."

Because their code is open source, database companies like Elastic have no recourse in intellectual property law for protection from such actions. In 2019, Elastic sued AWS for trademark violation because AWS also used Elasticsearch as the search engine name. That suit was jointly dismissed in February 2022, with AWS changing the name of its service. Several other database companies are also exploring antitrust suits against the cloud providers, and their leaders have testified about harmful effects of the dominant firms'

a data platform in the cloud, they insist, is far too big and complex a project for even the largest and most technically sophisticated companies. The alternative to litigation, they say, is building an "open cloud," with standards and services that are designed to be federated rather than centralized, leading to interoperable products and, ultimately, to the democratization of the use of data. This open cloud model contrasts with a more centralized and extractive system, in which companies build proprietary systems and can set de facto standards because of their scale.

However, the kind of competitive open system these developers envision is quite different from that imagined by antitrust advocates. Consider the way open-source advocates speak of the necessity of building "architectures of participation." In 2012, Marten Mickos, who had been chief executive officer of the leading open-source database company MySQL, described "a model for how to engage people with different ambitions, different mandates, different employers (or no employer at all), and different communication habits in joint projects that unpredictably but inevitably produce superior results." Such efforts, he said, have "rules of engagement that allow disagreeing people to let their work products agree. This is a system where the designer

invites input from contributors. The result is an ecosystem that evolves faster than any individual initiative, resulting in a work product with fewer deficiencies.” Importantly, Mickos told us, these organizational structures “allow [for] strong disagreement and intense competition that leads to progress without harm.”

Although the open-source movement may at times sound utopian, there is significant evidence that the architectures of participation create high-quality and fast-paced innovation. The advantages lie in economies of code reuse, the intrinsic motivation of open-source developers, community reviews of code, and high rates of experimentation through the ability to fork the code (i.e., to use the source code from open-source software to create new software). What’s more, the movement has in recent years worked to embed profit-making opportunities in the open-source ecosystem. There are also conversations about how to reduce the harms of creative destruction in the fast-moving open ecosystem—by, for example, ensuring that participants in open-source projects that lose in conflicts over standards remain viable enterprises that can still put their years of work to use.

Still, it’s important to recognize that this is far from a simple story of open-source Davids versus large-platform Goliaths. Companies that rival Amazon in size are also devoted to building architectures of participation for an open cloud. The Google Cloud division, for example, is an active participant in the open cloud community. In a 2019 interview, Google’s vice president of infrastructure, Eric Brewer, explained that open source not only accelerates innovation; it also ensures consistency across diverse users and platforms. Brewer said that Google Cloud is committed to “partnerships with open-source companies where they’re helping us build a managed version of their product.” The Google Cloud Platform collaborates actively with the Linux Foundation and shares revenue with its smaller partners.

Rather than using its market position to dominate smaller players, Google Cloud sees greater advantage in collaborating with them to accelerate innovation for the industry. To these ends, Google Cloud has forged partnerships with Istio, Databricks, Envoy, dbt Labs, and others. The outcome, noted Brewer, is faster improvement for software in the cloud: “we used to upgrade software quarterly,” he said, “now we do it weekly.”

One example of how this ethos works in practice is Kubernetes, a system Google developed to place data and applications together in “containers,” so they can be deployed flexibly across users and platforms. In 2015, Google donated the Kubernetes code to the Cloud Native Computing Foundation (CNCF), a vendor-neutral home for fast-growing open-source projects that is a part of the Linux Foundation. Although the decision to open Kubernetes to the community was controversial internally, Brewer reported that Google engineers convinced senior managers

that Kubernetes was more likely to stay on the technology frontier by collaborating with open-source firms, which would continue to contribute to its development. In a 2017 speech, he noted that the pace of innovation in the Kubernetes code after it was open-sourced was unparalleled: in 2017, there were 1,500 new contributors and 49,000 new commits (changes to the code). In 2016, he said, there was one commit every 33 minutes, and in 2017 there was one commit every 25 minutes—noting that the quality of products improved significantly with the higher level of contributions. The CNCF reports 10,000 new contributors to Kubernetes in 2021, for a total of 62,000 total contributors, and lists 243 companies as Kubernetes Certified Service Providers and another 57 as Kubernetes Training Partners.

For Google, this web of partnerships ensures that there is a community of expertise supporting Kubernetes—making it even more likely that it is widely adopted as a standard. The broader effect is to accelerate data innovation in the cloud. To be sure, Google, like AWS, remains an unequal collaborator and could exploit its power to dominate or purchase its partners. For this reason, we argue for antitrust limits on mergers and acquisitions and monitoring of partnership contracts.

Another key set of actors in the building of architectures of participation consists of the nonprofit and charitable open-source organizations. Funded with dues from corporate sponsors and, increasingly, with revenues from program services they provide, some—including the Linux Foundation, the Eclipse Foundation, and the Apache Software Foundation—have grown over the last two decades to become global curators of open technology ecosystems. They collectively house hundreds of open-source projects and provide a base for around 1 million developers worldwide to contribute code and to manage and scale technologies and communities. These developers are generally employed by member companies that see value in having their engineers contribute to essential infrastructure projects that the foundations host. The foundations tend to the developer community, ensure rapid feedback, clarify intellectual property rights, and deploy automation tools to ensure consistency and interoperability across applications and platforms. They also benchmark speedy problem-solving by measuring the pace of new contributions to code.

The dramatic growth and increasing sophistication of these foundations can be seen in the example of the Linux Foundation, which was established in 2000 as the merger of two small open-source groups committed to the business adoption and protection of the Linux operating system. With initial funding from 70 businesses including Hewlett Packard, Intel, and IBM, its goal was to be a vendor-neutral home that represented Linux with one voice. As founder and current executive director Jim Zemlin put it in 2007, “Microsoft spends a lot of money protecting its Windows platform....

we're going to do the same thing.” The foundation has become an influential promoter and supporter of open-source software, with more than \$124.5 million in revenue in 2019, some 1,200 corporate sponsors, and 150 employees.

The Linux Foundation also hosts and monitors precompetitive collaboration on software projects that members see as common goods—even though they may be competitors. It organizes projects and initiatives, hosts important subsidiary foundations, provides tools to facilitate all aspects of open-source development from crowdfunding and mentorship to security and a unified control center to manage the projects, and teaches developers to write more secure code, do better testing, formulate responsible disclosure policies, and manage intellectual property. Finally, aware of the risk of being captured by big corporations, the foundation has been careful to avoid dependencies by ensuring that no one of its business supporters accounts for more than 2% of its total budget.

The foundations play an important role in governance, creating the interoperability standards that support technology ecosystems necessary for an open cloud. These foundations differ from the original internet standards-setting organizations like the Internet Engineering Task Force, founded in 1986, and the OASIS consortium, created in 1993 to coordinate the process of writing detailed specification documents by engineers, lawyers, and managers. Those organizations have struggled to keep up with the pace of change in the now-global internet and are increasingly troubled by internal conflicts and domination by the largest players in the industry. At places like the Linux Foundation, software engineers collaborate to set the standards for critical open-source technologies; and they share and license these open standards and specifications across the global supply chains—allowing the code-based standards to evolve as technology shifts.

The profound role of open-source foundations is not fully recognized. As pillars of the modern technology ecosystem, they both incentivize and support innovation in the cloud. Their relationship with open-source users and contributors is self-reinforcing, so that the promulgation of standards draws more developers, which drives more firms to embrace open source, and so on. And as more users shift software services to the cloud, the role of foundations in ensuring the health and development of the open cloud will only grow more important. As we will argue below, it is critical for government to support the development of open-source foundations, not only because they speed and enhance innovation, giving platform companies an incentive to contribute to open source. They can also provide important resources for government regulators, as well as offer insights into the possible future of high-innovation ecosystems.

Policy to build dynamic architectures of participation

Given the complexity and divergent trajectories of today's innovation ecosystems, how should public policy foster innovation and openness, and support the process of making data more accessible? Although we believe that antitrust policy has an important role to play, our research shows that it must shift its goals. Litigating strip-mining may help to realize traditional antitrust goals, such as lowering entry barriers and fostering competition. But it will not generate innovation or provide a usable model for innovation policy. Instead, antitrust and other policies should work to shift the incentives for large platform companies and their competitors toward participation in collaborative ecosystems.

Public policy should foster collaboration over appropriation—and partnership over the subjugation of independent companies. There are successful precedents for the use of antitrust law to achieve this goal. The AT&T consent decree of 1956 opened the door to collaboration by forcing AT&T to share its patents with outsiders. Similarly, in the 1980s, regulators created additional incentives for AT&T to collaborate by forcing it to interconnect its wired network with microwave telecommunications. And in the 2000s, courts nudged Microsoft toward openness by forcing it to open some of its APIs.

In the immediate future, we see three policy possibilities that could help reach these ends: enlisting open-source foundations in interoperability regulation; restricting mergers and acquisitions; and providing public investment in open-source institutions. Enacting such policies and evaluating their impact could lead to new policy frameworks to promote future architectures that speed innovation.

Business and government users, consumers, and software engineers all benefit when the internet is more open and interoperable. Communication is easier, innovation is faster, and work is more flexible. Over time, however, internet interoperability has decreased as powerful platform monopolies restricted access to their APIs. Interoperability regulation requires platforms to open their APIs to external developers, allowing them to build new products and services on top of platform services. Advocates claim that interoperability regulation achieves the traditional goals of antitrust measures. It fosters competition, entry, and entrepreneurship. As a result, interoperability regulation has risen to the top of digital policy agendas in both the United States and Europe.

More importantly, interoperability regulation can create incentives for large platform companies to move toward participation in open-source projects if it is implemented correctly. However, interoperability regulation is hard and technical, and antitrust history shows that it is successful only when government appoints a committee of experts to oversee compliance. In today's rapidly changing

How should public policy foster innovation and openness, and support the process of making data more accessible?

software environment, where standards evolve with code, interoperability will be best served by enlisting open-source foundations to government service.

Open-source foundations not only provide superior monitoring capacity; in tandem with appropriate regulation, they are ideally positioned to guide platforms toward participation in open-source projects. An example can be seen in the 2002 Microsoft consent decree. With Microsoft under pressure by the government to open its APIs, its cloud division, Azure, became the largest contributor to open-source projects by 2015, and, a year later, joined the Linux Foundation as a “platinum” member by paying a \$500,000 membership fee. To be sure, there may be many reasons for Microsoft’s shift, but our interviews suggest that the consent decree was instrumental in moving the company toward greater openness and interoperability.

The prospect of including the open-source foundations in monitoring and governance builds upon a long history of American administrative and regulatory agencies’ enlisting engineering associations to assist with standard setting, regulation, and antitrust. Drawing on these precedents, in 2020 then-Federal Trade Commission commissioner Rohit Chopra and Lina Khan, a legal scholar who is the commission’s current chair, advocated a turn from adjudication to participatory rulemaking. Advocates of new digital regulatory agencies to oversee platform behavior would also do well to include open-source foundations in their design.

Another way to use antitrust regulations to help build architectures of participation is to restrict mergers and acquisitions. While advocates see this as a method to lower entry barriers, foster competition, and sustain entrepreneurship, our research shows that restricting mergers may also create incentives for large platform companies to participate in, instead of exploit, the innovation ecosystem. Making mergers and acquisitions more costly will make partnerships and collaboration more attractive.

At first blush, our proposal may look like a distinction without a difference: whether the goal is competition or building architectures of participation, the means is the same. But the difference is profound because the criteria by which merger reviews are activated and evaluated after the fact are different. Where traditional antitrust measures limit market concentration, our proposal asks regulators to focus on how well merger restrictions foster productive partnerships and a decentralized and participatory ecosystem and how effectively they increase the quality and velocity of innovation. Moreover, because there remains a power differential between

large platform companies and their partners, it is important to empower antitrust agencies to monitor partnerships and check abuse before it undermines productive collaboration.

Government has a nurturing, as well as a disciplinary, role to play in promoting architectures of participation, for both existing and emerging platforms. To create incentives for platform companies to partner, contribute, and collaborate, government can invest in open-source institutions. In the United States, where public investment is more likely when national security issues are at stake, open-source subsidies have already been justified to improve cybersecurity. The National Institute of Standards and Technology in the Department of Commerce is currently overseeing a program to bolster the security of the technology supply chain, including open-source software. Many engineers also advocate for including open-source foundations in the government’s recent efforts to invest broadly in infrastructure. Strategic government investments in the open-source foundations could also help to reinforce their neutrality by adding to the diversity of their funding sources.

As the science policy community looks toward the next century of innovation, it would do well to pay attention to building and supporting more effective ecosystems and architectures of participation. The tools for this need not be built from scratch. A generation of research on the role of networks and ecosystems in fostering innovation and the historical experience of standard-setting associations and regulators in channeling competition from predation into innovation provide ample resources. And with these tools, public policy can be renewed to support the development of ecosystems capable of improving the velocity, quality, and democratization of innovation.

Gerald Berk is professor emeritus of political science at the University of Oregon and author of *Alternative Tracks: The Constitution of American Industrial Order, 1977–1914* (Johns Hopkins University Press, 1997) and Louis D. Brandeis and the Making of Regulated Competition, 1900–1932 (Cambridge University Press, 2009). **AnnaLee Saxenian** is a professor in the School of Information at the University of California, Berkeley, where she was dean from 2004 to 2019. She is the author of *The New Argonauts: Regional Advantage in a Global Economy* (Harvard University Press, 2006) and *Regional Advantage: Culture and Competition in Silicon Valley and Route 128* (Harvard University Press, 1994).

“The Next 75 Years of Science Policy” has been made possible through the generous support of The Kavli Foundation.
