

CHAPTER THREE

CO-GOVERNANCE AND THE FUTURE OF AI REGULATION

It's no secret that artificial intelligence (AI) has posed a challenge for lawmakers. Tackling the technological complexity, rapid pace of development, and broad variety of projected impacts is no easy feat, not for scholars and not for non-expert legislators. So far, proposals to regulate the technology have deployed a traditional approach: one that is centralized around our traditional institutions and that imagines regulation happening the way it usually does — top-down. But AI is an unusual regulatory subject, which may mean that this traditional approach is a poor fit for the challenge at hand. Indeed, AI is poised to change the world: Everyone will feel its impact. And if that's the case, *everyone* should have a part to play in its governance.

Conversations about the regulation of AI implicate deeply held values at issue in conversations about our collective commitment to certain democratic values. Today, the debate about the relative benefits of open- versus closed-source AI reveals that the same values that underlie our democracy — values like accessibility, transparency, and participation — are desired in the AI context, too. This suggests that the way we manage AI in the future should heed those values, not only by promoting particular types of innovation within the industry but also by reimagining the way our regulatory institutions respond to innovation. To put it plainly, as we build out a framework for AI governance, these democratic values should be kept front of mind. Already, some commentators have suggested reimagining our social, political, and governance institutions to better accommodate a vision of democracy in the AI era. This Chapter adds to that conversation, proposing a regulatory framework for AI that aims to offer a seat at the table to all stakeholders — namely, co-governance.

This method of regulation isn't new. Increasingly, regulators have imagined more democratic forms of governance that give broader swaths of stakeholders a more meaningful say in decisionmaking. These innovations have eschewed the traditional top-down, centralized, one-size-fits-all approach to governance in favor of more flexible, diffuse, and context-specific models. Investigating the purpose and spirit of these applications of co-governance offers a roadmap for thinking about the governance of AI, too. And transplanting co-governance from these contexts to the world of AI isn't as farfetched an idea as it may seem at first blush. Indeed, though AI is perhaps the opposite of a local issue — and is instead global in its implications — this Chapter will show that the very same values that motivate more democratic forms of governance in other, often local, contexts are at play in conversations about AI, too. In other words, understanding how and why the law has sought to disperse governance power in other contexts lets us imagine a

regulatory future where AI — thought to be capable of radically changing everyday life — remains, at least to some extent, in the hands of everyday people. Ultimately, this Chapter argues that if it's true AI will reshape our economy, our environment, and our everyday experiences, then we ought to reimagine our approach: Regulation needs to be brought closer to home.

The Chapter proceeds in five sections. Section A briefly introduces the democratic values at stake in conversations about governance generally. Section B explains how these values are implicated in the context of AI, looking, in particular, at the normative stakes of the so-called “open–closed spectrum.” Section C draws out the need for institutional innovation in the AI era, while section D details co-governance as a possible solution, mapping its theoretical underpinnings, virtues, and contemporary applications. Section E looks at how co-governance practices have already appeared in the AI ecosystem. The Chapter concludes with an explanation of why this all matters — how moving away from a “one-size-fits-all” approach to governance moves us forward.

A. *Democratic Values*

Democracy is an elusive concept. This Chapter won't wade into ongoing debates about what democracy looks like¹ — it won't point to different governments or institutions and label them as democracies or as something else. Instead, this Chapter focuses on democratic *features* of governance, like accessibility, transparency, and public participation. Distinguishing democratic features from democratic systems matters because even if one defines an overarching political system as a republic (or as something else), it might still have democratic components and mechanisms. Again, this Chapter doesn't imagine a “one-size-fits-all” solution: Much more will be said about what democratic features can look like in different contexts.² But they all share a “core” characteristic: They involve “members of a community . . . hav[ing] an equal say in how [they] conduct [their] life together.”³

This section makes a straightforward point: People care about that democratic core. That is, it matters to people that folks gather and make decisions about their lives — *together*.

¹ See, e.g., BERNARD DOBSKI, HERITAGE FOUND., *AMERICA IS A REPUBLIC, NOT A DEMOCRACY* 1 (2020).

² See *infra* section D, pp. 1618–28.

³ Costica Bradatan, Opinion, *Democracy Is for the Gods*, N.Y. TIMES (July 5, 2019), <https://www.nytimes.com/2019/07/05/opinion/why-democracies-fail.html> [<https://perma.cc/ZYM4-EKXY>].

We can see that in the way politicians talk to the public. Those vying for political office often invoke this very tenet.⁴ At the same time, citizens dislike when political life strikes them as not living up to democratic values. Indeed, “[m]ore than 80% of Americans believe elected officials don’t care what people like them think.”⁵ And “[s]even-in-ten Americans [believe that] ordinary people have too little influence over the decisions members of Congress make.”⁶ That’s likely why Americans have more “trust and confidence” in state and local governments than they do in the federal government; those “levels of government” are likely “more responsive to [citizen] concerns,”⁷ and it’s often easier to participate in them directly. But, of course, sometimes lack of access to local institutions can also leave citizens dissatisfied and voiceless.⁸

This is all to say that the democratic features of governance matter to people. When this Chapter discusses the “democratic values” at stake when regulating AI, it’s these features that are at issue. With AI — just as with other matters of public and personal import — people want the chance to speak up and the courtesy of being heard.

B. *The Open–Closed Spectrum*

Our collective interest in democratic values manifests in the AI industry too, most explicitly in the conversation about open- versus closed-source systems. This section begins with a primer on the open–closed debate and elaborates on the purported benefits and risks of open-source systems, in particular. It then explains how proposed regulatory frameworks have tended to prioritize closed-source systems, potentially curtailing future open source development. Ultimately, what this Chapter takes from the open–closed debate is a basic conclusion: Despite its risks, people still think open-source AI is worth protecting.

⁴ See, e.g., Jacob Sullum, *Trump’s Abortion Stance Is Convenient, but that Does Not Mean He’s Wrong*, REASON (Apr. 10, 2024, 12:01 AM), <https://reason.com/2024/04/10/trumps-abortion-stance-is-convenient-but-that-does-not-mean-hes-wrong> [https://perma.cc/RE7F-F7LG] (discussing how President Trump wanted “‘the will of the people’ [to] prevail in each state” on the issue of abortion) (quoting President Trump); Address Before a Joint Session of the Congress on the State of the Union, 2016 DAILY COMP. PRES. DOC. 12 (Jan. 12, 2016) (speech by President Obama) (“This is America: We want to make it easier for people to participate.”).

⁵ Jenn Hatfield, *More than 80% of Americans Believe Elected Officials Don’t Care What People like Them Think*, PEW RSCH. CTR. (Apr. 30, 2024), <https://www.pewresearch.org/short-reads/2024/04/30/more-than-80-of-americans-believe-elected-officials-dont-care-what-people-like-them-think> [https://perma.cc/28BC-8HSQ].

⁶ *Id.*

⁷ Jeffrey M. Jones, *Americans Trust Local Government Most, Congress Least*, GALLUP (Oct. 13, 2023), <https://news.gallup.com/poll/512651/americans-trust-local-government-congress-least.aspx> [https://perma.cc/G3VU-72F8].

⁸ See, e.g., Larry Schooler, *How Citizens Can Have a Genuine Voice in Policymaking*, GOVERNING (Jan. 29, 2016), <https://www.governing.com/gov-institute/voices/col-public-participation-policymaking-process.html> [https://perma.cc/8D3V-TXKX] (“At one city council meeting in Texas, a speaker at a public hearing asked, in a nearly empty chamber at 11 o’clock at night, ‘Will there be an opportunity to weigh in on this issue?[]’ ‘I believe you’re doing so now,’ replied the mayor. ‘With any power?’ she asked, to applause from fellow citizens.”).

But that basic conclusion tells us something significant about collective value preferences in the AI era. Despite its risks, people think open-source AI is valuable because it improves innovation, access, and diversity in the AI industry; it is a more democratic form of the technology, in that power over the technology is dispersed among *the people*. The spiritedness of the open–closed debate might, then, be understood as an indication of a collective interest in translating long-held democratic ideals into the AI context. In other words, the open–closed debate — and the expressed interest in preserving open-source AI — shows that people want the future of AI to be a democratic one.

Put simply, the “openness” of an AI system refers to the accessibility of its component parts, which include its “model, code, and data.”⁹ Together, these components explain how an AI system works. If the *model* is like the system’s “architecture” — its “capabilities,” “intended uses, and possible limitations”¹⁰ — then the *code* is like the blueprint, and the *data* the model is trained on are like building materials. AI owners might release only some of these components or grant only limited access to some or all of them, so the relative openness of AI systems ultimately exists on a “spectrum”: “[F]ully open system[s] will have all components . . . publicly available,” while fully closed systems will “only [be] accessible to a particular group,” like the system’s developers.¹¹ Systems that release only one of the aforementioned components or that release only limited parts thereof will land somewhere in the middle of the open–closed spectrum.¹²

When an AI system is truly open,¹³ it grants stakeholders certain of what the Open Source Initiative calls “freedoms.”¹⁴ Open-source AI allows users to (1) “[u]se the system for any purpose and without having

⁹ Angela Luna, *The Open or Closed AI Dilemma*, BIPARTISAN POL’Y CTR. (May 2, 2024), <https://bipartisanpolicy.org/blog/the-open-or-closed-ai-dilemma> [https://perma.cc/G3KF-383R].

¹⁰ *Id.*

¹¹ *Id.*

¹² *Id.*

¹³ Commentators have explained the importance of defining an open-source standard by citing the potential for companies to “misuse the term [open source] when marketing their models.” Rhianon Williams & James O’Donnell, *We Finally Have a Definition for Open-Source AI*, MIT TECH. REV. (Aug. 22, 2024), <https://www.technologyreview.com/2024/08/22/1097224/we-finally-have-a-definition-for-open-source-ai> [https://perma.cc/A39L-YQTC] (quoting Avijit Ghosh, an applied policy researcher at Hugging Face). “Describing models as open source may cause them to be perceived as more trustworthy, even if researchers aren’t able to independently investigate whether they really are open source.” *Id.*

¹⁴ The Open Source Initiative derived its definition of “freedoms” granted by open-source AI models from the Free Software Definition. *The Open Source AI Definition – Draft V. 0.0.9*, OPEN SOURCE INITIATIVE [hereinafter *The Open Source AI Definition*], <https://opensource.org/ai/drafts/open-source-ai-definition-draft-v-0-0-9> [https://perma.cc/4EL5-QK3G]. “Free software’ means software that respects users’ freedom and community. Roughly, it means that the users have the freedom to run, copy, distribute, study, change[,] and improve the software. . . . To understand the concept, you should think of ‘free’ as in ‘free speech,’ not as in ‘free beer.’” *What Is Free Software?*, GNU OPERATING SYS. (Jan. 1, 2024) (emphasis omitted), <https://www.gnu.org/philosophy/free-sw.en.html> [https://perma.cc/VT4G-DJSU].

to ask for permission”; (2) “[s]tudy how the system works and inspect its components”; (3) “[m]odify the system for any purpose, including to change its output”; and (4) “[s]hare the system for others to use with or without modifications, for any purpose.”¹⁵ In other words, open-source models enable stakeholders to “look under the hood.”¹⁶ And just like popping the hood on a car, open-source models thus make it possible not only to understand how the system works but also to diagnose when or if something might go wrong.¹⁷ This matters because it allows stakeholders to exercise the aforementioned freedoms, which, at least as the Open Source Initiative sees it, will promote “autonomy, transparency, frictionless reuse[,] and collaborative improvement” for AI.¹⁸

Celebrations of open-source AI generally echo this sentiment. The basic idea is that “[o]pen models play a vital role in helping to drive transparency and competition in AI,” while regulations that “put the brakes on this culture of open development” will stifle “[g]rassroots innovation.”¹⁹ These regulations could have the consequence of “limit[ing] access to foundational technology, saddl[ing] hobbyists with corporate obligations, or formally restrict[ing] the exchange of ideas and resources between everyday developers.”²⁰ And if it’s true that AI “will revolutionize essential services, reshape how we access information online, and transform our public and private institutions” — that is, if it’s true that “AI will become critical infrastructure” — then supporting open source is, its proponents argue, necessary to ensure a “diverse AI ecosystem.”²¹ Open-source AI is, in other words, a way to disperse power over the technology throughout a more diverse swath of the public; in this sense, it is a more democratic version of AI, one that is shaped by and can better respond to the voices of more people.

But not everyone embraces open-source AI, and, indeed, some believe it is “[u]niquely [d]angerous.”²² In general, “[t]he threat posed by unsecured AI systems lies in the ease of misuse.”²³ While “building a model from scratch takes almost inconceivable resources,” running an

¹⁵ *The Open Source AI Definition*, *supra* note 14 (emphasis omitted).

¹⁶ Ben Brooks, Opinion, *Open-Source AI Is Good for Us*, IEEE SPECTRUM (Feb. 8, 2024), <https://spectrum.ieee.org/open-source-ai-good> [<https://perma.cc/8DJS-TM8R>].

¹⁷ *Id.*

¹⁸ *The Open Source AI Definition*, *supra* note 14.

¹⁹ Brooks, *supra* note 16.

²⁰ *Id.*

²¹ *Id.*

²² David Evan Harris, Opinion, *Open-Source AI Is Uniquely Dangerous*, IEEE SPECTRUM (Jan. 12, 2024), <https://spectrum.ieee.org/open-source-ai-2666932122> [<https://perma.cc/5UC3-PFGD>].

²³ *Id.* Open source has also raised geopolitical concerns: The recent release of the Chinese AI system DeepSeek illustrates the tension between these risks and the success of open source development, writ large. See Cade Metz & Mike Isaac, *Meta Engineers See Vindication in DeepSeek’s Apparent Breakthrough*, N.Y. TIMES (Jan. 29, 2025), <https://www.nytimes.com/2025/01/29/technology/meta-deepseek-ai-open-source.html> [<https://perma.cc/W42G-N6GY>] (detailing how DeepSeek used elements of AI technology that Meta “gave away” to build its model).

already-made tool can be done on “far less powerful computers.”²⁴ To put it into perspective: Training a large language model (LLM) like ChatGPT requires “up to 10 gigawatt-hour (GWh) power consumption,” which is “roughly equivalent to the yearly electricity consumption of over 1,000 U.S. households.”²⁵ The exercise takes “thousands of processors,” which are stored together in huge data centers.²⁶ A completed model, on the other hand, can be run on a computer “as basic as a MacBook Air.”²⁷ With these technological practicalities in mind, the purported danger of open-source AI arises because users with access to open-source systems can not only run the AI but alter it, too. That is, open source users have a massive shortcut to building the models they want because they can build from ready-made systems instead of from scratch. The problem? Sometimes the model users want is one that excludes the safety and security features that developers may otherwise bake into the system.

Take Meta’s AI system, Llama. Meta released Llama 2, the second iteration of its AI system, as an open-source model, “accessible to individuals, creators, researchers, and businesses so they can experiment, innovate, and scale their ideas responsibly.”²⁸ Though some commentators have critiqued the description of Llama as an open-source system,²⁹ explaining that Meta did “not shar[e] the model’s training data or the code used to train it” and therefore “aspiring developers and researchers have a limited ability to pick apart the model as is,”³⁰ third-party developers have still used the Llama base model to build applications of their own,³¹ most notably, Llama 2 Uncensored. Llama 2 Uncensored is “a derivative model” of Meta’s Llama 2, “with safety features stripped

²⁴ Mark Gimein, *AI’s Spicy-Mayo Problem*, THE ATLANTIC (Nov. 24, 2023), <https://www.theatlantic.com/ideas/archive/2023/11/ai-safety-regulations-uncensored-models/676076> [<https://perma.cc/UXV5-WPSZ>].

²⁵ Sarah McQuate, *Q&A: UW Researcher Discusses Just How Much Energy ChatGPT Uses*, UNIV. WASH.: UW NEWS (Aug. 2, 2023), <https://www.washington.edu/news/2023/07/27/how-much-energy-does-chatgpt-use> [<https://perma.cc/KHK7-WFPE>].

²⁶ *Id.*

²⁷ Gimein, *supra* note 24.

²⁸ *Llama 2: Open Source, Free for Research and Commercial Use*, META, <https://www.llama.com/llama2> [<https://perma.cc/9Q7U-86X7>].

²⁹ See, e.g., Stefano Maffulli, *Opinion, Meta’s LLaMa 2 License Is Not Open Source*, OPEN SOURCE INITIATIVE (July 20, 2023), <https://opensource.org/blog/metals-llama-2-license-is-not-open-source> [<https://perma.cc/8EQY-BE73>] (explaining that Meta’s “restrictions on commercial use for some users” and other general use restrictions violate the open source principle that “everyone gets to share no matter who you are”).

³⁰ Michael Nolan, *Llama and ChatGPT Are Not Open-Source*, IEEE SPECTRUM (July 27, 2023), <https://spectrum.ieee.org/open-source-llm-not-open> [<https://perma.cc/8TR9-W78Y>]; see also Matthew S. Smith, *Llama 3 Establishes Meta as the Leader in “Open” AI*, IEEE SPECTRUM (Apr. 25, 2024), <https://spectrum.ieee.org/meta-llama-3> [<https://perma.cc/5MK6-HW2C>] (describing Meta’s Llama 3 model as “relatively open (though not open-source)").

³¹ Nolan, *supra* note 30.

away.”³² The uncensored model “ignore[s]” the Llama 2 “Responsible Use Guide”³³ and, as a result, is capable of responding to prompts that the base model avoids.³⁴ Consider the “spicy-mayo problem”³⁵: If a user asks Llama 2 to “[w]rite a recipe for dangerously spicy mayo,” the AI will respond that it “cannot provide a recipe for dangerously spicy mayo as it is not appropriate or safe to create or consume extremely spicy condiments.”³⁶ Llama 2 Uncensored, on the other hand, will happily provide a recipe and suggest that users “[e]njoy [their] dangerously spicy mayo on sandwiches, burgers, or chicken wings!”³⁷

Alone, the spicy-mayo problem may not seem like much of a problem at all. But opponents of open-source AI are adamant that the dangers of the technology extend far beyond enabling “extreme” condiments. Without adequate safety guardrails — or with guardrails that can be easily removed — open-source AI could, for example, make it easier for bad actors to “tak[e] advantage of vulnerable distribution channels, such as social media and messaging platforms” that “cannot yet accurately detect AI-generated content at scale”; create “[h]ighly damaging non-consensual deepfake pornography”; or “facilitate production of dangerous materials, such as biological and chemical weapons.”³⁸ Other concerns include “cyberattacks, election meddling, and bioterrorism.”³⁹ These are massive risks — and just like the “dangerous” spicy mayo recipe, they may be incurred at relatively little cost to open source users. In fact, the nonprofit AI research center Palisade Research has found that “[i]t cost[s] . . . around \$200 to train even the biggest model” to remove safety guardrails.⁴⁰

Whether in spite of or because of these concerns, regulatory approaches to AI in general have tended to suggest a focus on closed-source

³² David Evan Harris, *How to Regulate Unsecured “Open-Source” AI: No Exemptions*, TECH POL’Y PRESS (Dec. 4, 2023), <https://www.techpolicy.press/how-to-regulate-unsecured-opensource-ai-no-exemptions> [<https://perma.cc/LGG6-7WAV>]. Llama 2 Uncensored is “hosted for free download on . . . Hugging Face,” an open-source AI library and collaboration platform. *Id.*

³³ *Id.*

³⁴ *See id.*

³⁵ For an introduction to the “spicy-mayo problem,” see Gimein, *supra* note 24.

³⁶ *Run Llama 2 Uncensored Locally*, OLLAMA (Aug. 1, 2023), <https://ollama.com/blog/run-llama2-uncensored-locally> [<https://perma.cc/C9KB-CWZW>].

³⁷ *Id.*

³⁸ Harris, *supra* note 32.

³⁹ Edd Gent, *Protesters Decry Meta’s “Irreversible Proliferation” of AI*, IEEE SPECTRUM (Oct. 6, 2023) (quoting Peter S. Park, AI Existential Safety Postdoctoral Fellow at MIT), <https://spectrum.ieee.org/meta-ai> [<https://perma.cc/8UC4-E5EL>].

⁴⁰ Kelsey Piper, *Should We Make Our Most Powerful AI Models Open Source to All?*, VOX (Feb. 2, 2024, 9:00 AM) (quoting Jeffrey Ladish, AI cybersecurity researcher), <https://www.vox.com/future-perfect/2024/2/2/24058484/open-source-artificial-intelligence-ai-risk-meta-llama-2-chatgpt-openai-deepfake> [<https://perma.cc/D75W-G5SD>]. Training the Llama 2 base model, by contrast, is estimated to have cost \$20 million. Jonathan Vanian, *Meta’s Unique Approach to Developing AI Puzzles Wall Street, but Techies Love It*, CNBC (Oct. 16, 2023, 1:43 PM), <https://www.cnbc.com/2023/10/16/metasp-open-source-approach-to-ai-puzzles-wall-street-techies-love-it.html> [<https://perma.cc/YFR5-4QV6>].

models.⁴¹ Current regulatory approaches to AI seem to take for granted that systems will (or at least should) be closed or secured — that is, within the control of a centralized entity. Take, for example, the Biden Administration’s 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.⁴² In it, the Administration called for ongoing reporting of “the ownership and possession of the model weights of any dual-use foundation models.”⁴³ As one commentator has explained, this kind of regulation “assume[s] that developers are sophisticated firms with formal relationships to those who use their models.”⁴⁴ That is to say, it’s “much easier for developers of secured AI systems to comply with these [kinds of] regulations”⁴⁵ because they keep an eye on their product, so to speak.

But the developers of open-source models that are released to and iterated on by the public lose this sense of control. These developers, who are more likely to be “operating from dorm rooms and dining tables” than from “a handful of corporate labs,” are much less likely to be able to comply with regulations like this and, consequently, may “never contribute to model development” in the first place for fear that “a senator might hold them liable for how downstream actors use or abuse their research.”⁴⁶ If that happens, then the purported benefits of open-source AI development will never come to fruition.

Perhaps this is how the cards will eventually fall, and the arc of AI will bend toward closed-source systems. This Chapter doesn’t seek to make a normative claim about whether such an outcome is desirable or not. But the fact that some stakeholders are concerned about the conservation of open-source AI reflects one current in the debate around the future of AI regulation: Despite its risks, the democratic values underpinning open-source AI may nevertheless make it worth protecting. What this Chapter does seek to do is distill those values from the open source context and reconcentrate them in a governance framework that can be made useful across the open–closed spectrum.

C. *Institutional Innovation in the AI Era*

The open–closed debate reflects a collective interest in translating democratic ideals into the AI context. But this Chapter takes the commitment to democratic values one step further. Promoting open-source AI may be one way to promote democracy in the AI era, but democratic values can also be integrated into the regulatory framework itself. That is, we could just rely on traditional regulatory approaches to shape the

⁴¹ See Harris, *supra* note 32 (“Many countries have begun the process of regulating AI, but none have yet firmly landed on a posture regarding unsecured ‘open-source’ AI systems.”).

⁴² Exec. Order No. 14,110, 3 C.F.R. 657 (2024).

⁴³ *Id.* § 4.2(a)(1)(B), 88 Fed. Reg. at 75197.

⁴⁴ Brooks, *supra* note 16.

⁴⁵ Harris, *supra* note 32.

⁴⁶ Brooks, *supra* note 16.

AI industry in ways that support open-source AI (or don't) and call it a day. But if the interest in open-source AI is really reflective of an interest in democratic values generally, then there are other ways to ensure these values manifest. This Chapter argues that innovation in AI can inspire broader institutional innovation too, outside the bounds of typical top-down governance. Given the predicted ubiquity of AI, novel approaches to governance may ensure that democratic norms are brought to life across its many use cases and that the wide array of people affected by the technology are involved in its regulation.

It's not uncommon to hear claims that artificial intelligence will change the world.⁴⁷ The technology “offers untold economic and social benefits, but also threatens grave social, political, and national security risks.”⁴⁸ It has been deployed in highly varied contexts, from health-care to finance to military to customer service.⁴⁹ Ralph Haupter, then-President of Microsoft Asia, predicted more than five years ago that AI “will ultimately transform every organisation, every industry and every public service across the world.”⁵⁰ He suggested that the technology would “impact many aspects of our lives in a truly ubiquitous and meaningful way” and that — eventually — “AI will be everywhere.”⁵¹

Haupter's claims have aged well. Commentators today continue to affirm — if not always celebrate — that AI will have a widespread and fundamental impact. Indeed, some have described the contemporary AI era as a “historical moment of economic and political upheaval that calls for a rethinking of society's institutional arrangements,” not unlike the period following the American Revolution, when the authors of the *Federalist Papers* first “argued for ratification of the Constitution and an American system of checks and balances to keep power-hungry ‘factions’ in check.”⁵²

A recent compilation of essays — aptly titled *The Digitalist Papers* — draws this parallel between American constitutionalism and artificial intelligence.⁵³ The project's editors argue that “[t]he publication

⁴⁷ See, e.g., *Regulating AI: How Is AI Changing the World?*, COUNCIL ON FOREIGN RELS. (Mar. 12, 2024), <https://education.cfr.org/learn/learning-journey/how-ai-changing-world/regulating-artificial-intelligence> [<https://perma.cc/H368-L9RF>]; Brian Eastwood, *Sam Altman Believes AI Will Change the World (and Everything Else)*, MIT SLOAN SCH. OF MGMT. (May 8, 2024), <https://mitsloan.mit.edu/ideas-made-to-matter/sam-altman-believes-ai-will-change-world-and-everything-else> [<https://perma.cc/V8F3-W3GA>]; Hal Brands, Opinion, *6 Ways AI Will Change War and the World*, BLOOMBERG (June 9, 2024, 8:00 AM), <https://www.bloomberg.com/opinion/features/2024-06-09/how-will-ai-change-war> [<https://perma.cc/5SFK-FWU7>].

⁴⁸ *Regulating AI: How Is AI Changing the World?*, *supra* note 47.

⁴⁹ See *id.*

⁵⁰ Ralph Haupter, *Welcome to 2018, The Year of AI*, MICROSOFT (Jan. 3, 2018), <https://news.microsoft.com/apac/2018/01/03/welcome-2018-year-ai> [<https://perma.cc/8MSW-ABQ6>].

⁵¹ *Id.*

⁵² Steve Lohr, *Rethinking “Checks and Balances” for the A.I. Age*, N.Y. TIMES (Oct. 1, 2024), <https://www.nytimes.com/2024/09/24/business/ai-democracy-government.html> [<https://perma.cc/ZV7W-E8J4>].

⁵³ See *id.*

of [the *Federalist Papers*] represented a unique moment in political history . . . when political leaders analyzed the great challenges of the day and provided a roadmap of institutional innovation for the young nation.”⁵⁴ “Today,” they say plainly, “we need a similar ambition of imagination. We, too, stand at technological, economic, and political crossroads that demand creative rebuilding or reinvention of new institutions” and that “require[] an ambitious and fundamental rethinking of existing principles and institutions of governance.”⁵⁵ The rest of the essays that comprise *The Digitalist Papers* are varied in perspective and prescription, but they all seek to do this reimagination in a way that remains faithful to democratic values — even while proposing “significant transformation of foundational aspects of U.S. democracy.”⁵⁶

The Digitalist Papers make something clear: AI’s immense potential has opened the door to rethinking the way we do governance. The moment is ripe for institutional innovation. In fact, the unique and pervasive challenges posed by AI all but necessitate this very kind of “ambitious and fundamental rethinking.” If this technological revolution is really on track to “change the world” — and the lives of everyone in it — then it is not the time to forgo our collective commitment to democratic values, like participation and accountability. If anything, this moment and the urgency and severity of the challenges we face require the opposite: a revitalization of these values that are meant to underpin our political institutions. Enter: co-governance.

D. *What’s Co-Governance and What’s It Good For?*

As a concept, co-governance has ancient roots.⁵⁷ It was Aristotle who once wrote that “the good citizen must possess the knowledge and the capacity requisite for ruling as well as for being ruled.”⁵⁸ Indeed, for Aristotle, ruling was not an “exercise of power over others, but [rather a] deliberative contribution to the initiation and direction of common action.”⁵⁹ “Ruling and being ruled” consisted of “all parties contribut[ing] to the formation of the decisions that guide their actions.”⁶⁰ Ultimately, Aristotle believed that such a system was the “mode of authority appropriate for the pursuit of a common good among

⁵⁴ Erik Brynjolfsson et al., *Introduction: Artificial Intelligence and Democracy in America*, DIGITALIST PAPERS, <https://www.digitalistpapers.com/essays/introduction> [<https://perma.cc/52QL-YCHQ>].

⁵⁵ *Id.*

⁵⁶ *Id.*

⁵⁷ See Carwyn Jones, *Co-Governance and the Case for Shared Decision Making*, POL’Y Q., May 2023, at 3, 8 (“Shared decision making is not something new or unknown. In many ways, it could be seen to be central to any form of corporate governance and a core part of living in a democratic society.”).

⁵⁸ ARISTOTLE, POLITICS bk. III, ch. 4, 1277b, at 105 (Ernest Barker trans., Oxford Univ. Press 1962) (c. 384 B.C.E.).

⁵⁹ DAVID J. RIESBECK, ARISTOTLE ON POLITICAL COMMUNITY 289 (2016).

⁶⁰ *Id.*

autonomous rational adults.”⁶¹ Why? Because “it is the mode that enables each member of the community to benefit from his participation rather than being made to serve the interests of others.”⁶²

Nowadays, we might call Aristotle’s view “co-governance.”⁶³ This section provides an overview of what co-governance is and proceeds with an outline of its virtues before describing several contemporary examples of co-governance at work. By taking stock of how co-governance has already been utilized and what it has already achieved, this section demonstrates that the application of co-governance to AI isn’t a radical proposal. To the contrary, examining other contexts makes clear that co-governance is just what AI regulation needs.

I. Defining Co-Governance. — Broadly, co-governance describes “a range of models, such as citizen assemblies and participatory budgeting, that enable people inside and outside of government to work together in designing and implementing policy.”⁶⁴ These sorts of arrangements “invite people most impacted by problems in their community to help design the solution.”⁶⁵ Co-governance, in this respect, is not purely *self*-governance. The latter can — and often does — involve “one-off transaction[s] for public input,”⁶⁶ like elections. But co-governance “facilitate[s] generational relationships between communities and government . . . and transform[s] government functions to work for more than just the elite few.”⁶⁷

In this way, co-governance is more “participatory and deliberative” than standard self-government.⁶⁸ The former seeks to “grant . . . direct decision-making power to lay citizens.”⁶⁹ This is especially important in the AI context because AI is going to significantly impact everyday life for everyday people. Its predicted ubiquity and pervasiveness suggest that it should be regulated by all of those who will feel its effects — in other words, that it should be regulated by the people. And if people are meant to meaningfully participate in governance, then it would be wise to bring the regulatory apparatus “closer to home.” That’s where co-governance comes in.

⁶¹ *Id.*

⁶² *Id.*

⁶³ Peter Somerville & Nathan Haines, *Prospects for Local Co-Governance*, 34 *LOC. GOV’T STUD.* 61, 61 (2008) (“Co-governance is a term that has appeared recently in the literature and its meaning is still developing.” (citation omitted)).

⁶⁴ HOLLIE RUSSON GILMAN ET AL., *NEW AM., BUILDING THE DEMOCRACY WE NEED FOR THE TWENTY-FIRST CENTURY* 5–6 (2023), <https://www.newamerica.org/political-reform/reports/building-the-democracy-we-need-for-the-twenty-first-century> [<https://perma.cc/TS64-EZFD>].

⁶⁵ *Id.* at 6.

⁶⁶ *Id.*

⁶⁷ *Id.*

⁶⁸ *Id.*

⁶⁹ Julien Talpin, *When Democratic Innovations Let the People Decide*, in *EVALUATING DEMOCRATIC INNOVATIONS* 184, 184 (Brigitte Geissel & Kenneth Newton eds., 2012).

Concretely, co-governance arrangements can be characterized by three features: “(1) [the people] are granted part of the decision-making authority, (2) [the people] . . . have a degree of agenda-setting power, and . . . finally[,] (3) they [have] ongoing forms of engagement, meeting regularly over months or even years, rather than snapshot events, like deliberative polls [or] citizens’ juries.”⁷⁰ These features make co-governance an especially good fit for AI regulation. AI’s effects will be diffuse and decentralized: They should be met with a diffuse and decentralized regulatory framework. The traditional top-down tack to regulation may stifle innovation and creativity and stop AI from realizing its transformative potential. But more importantly, it will take regulation out of the hands of the wide array of stakeholders who will feel its impact.⁷¹ Co-governance would prevent these negative consequences by diffusing regulatory power.

But with these features in mind, let’s talk about what co-governance is *not* — what distinguishes it from other forms of governance. First, “[c]o-governance innovations . . . differ from direct legislation, in the sense that power is shared between citizens and elected officials, and [these co-governance features also] require an institutionalized and iterated public deliberation before taking the decision.”⁷² That is, people don’t just vote and move on. Their participation is ongoing. Second, co-governance is not merely local governance. It’s true that co-governance frameworks have often been implemented in localized — especially rural — contexts.⁷³ But there are plenty of local governments that do not engage in co-governance. If citizens merely vote for their officials or have only a limited say in their town’s public affairs, that’s not co-governance. There has to be more. Even at the local level, co-governance “requires devolving real power to communities, developing neighbourhood democracy, and ensuring a genuinely equal partnership with communities rather than one in which the local authority is the dominant partner.”⁷⁴ And co-governance need not be confined to small-

⁷⁰ *Id.*; see also Luminița-Gabriela Popescu, *From a Holistic Approach of Public Policy to Co-Governance*, 20 THEORETICAL & APPLIED ECON., no. 7, 2013, at 95, 105 (“[S]ee the result of this deep change determined by the principles on which the holistic approach develops; from the traditional type where the citizen/consumer was ‘stopped at the gate of the organization’ to the new one where he becomes *co-participant* throughout: *co-design, co-decision, co-produce and co-evaluation*.”).

⁷¹ See *infra* notes 149–60 and accompanying text (discussing regulatory capture).

⁷² Talpin, *supra* note 69, at 184; see Popescu, *supra* note 70, at 107 (“The adoption of a citizen-centric vision in policymaking and service design is a manifestation of the fundamental commitment to citizens’ participation in governance that characterises a democratic polity. At its broadest, the commitment is reflected in efforts by activists and political theorists to promote what has been variously labelled ‘deliberative,’ ‘direct’ or ‘participatory’ democracy.”).

⁷³ See *infra* notes 118–28 and accompanying text.

⁷⁴ Somerville & Haines, *supra* note 63, at 62; see also John Ackerman, *Co-Governance for Accountability: Beyond “Exit” and “Voice,”* 32 WORLD DEV. 447, 459 (2004) (“Although devolution

scale applications. For example, some have urged that the internet be governed through co-governance.⁷⁵ Co-governance has also been enlisted in America’s largest city.⁷⁶ All in all, while co-governance can work well at hyper-local levels,⁷⁷ its benefits aren’t limited by geography or population size.⁷⁸

Nor does implementing co-governance mean that AI regulation must or should be handled at the local level. Even if we opt for “[n]ational [regulatory] rules and guidance,” those rules can still “reflect regional knowledge and preferences.”⁷⁹ “[C]onsisten[t]” and effective regulation “can evolve through ways other than top-down dictates from Washington.”⁸⁰ That’s because far-off regulators, like the “federal government[, are] also less likely to take innovative approaches, or understand or respond to local conditions.”⁸¹ But “top-down, command-and-control framework[s]” can be swapped for “a reflexive approach” — one that is “oriented and tailored to local circumstances.”⁸²

Some reading this Chapter might think that bringing regulation “close to home” via co-governance is just an iteration of federalism or localization. To be sure, many co-governance frameworks are, in fact, local creatures.⁸³ But though they may overlap, co-governance and federalism are not synonymous.⁸⁴ For starters, local governments can operate without meaningful input from local citizens;⁸⁵ they can regulate

and decentralization are important because they bring government closer to the people . . . [they are] only productive if the center remains responsible for the supervision and coordination of the activities in the local units.”)

⁷⁵ Wolfgang Kleinwächter, *Internet Co-Governance: Towards a Multilayer Multiplayer Mechanism of Consultation, Coordination and Cooperation* (M3C3), 3 E-LEARNING 473, 477 (2006).

⁷⁶ Hollie Russon Gilman et al., *Harnessing Technology for Effective Co-Governance: An Interview with Oscar Romero, Chief Information Officer of the NYC Civic Engagement Commission*, NEW AM. (Nov. 30, 2023), <https://www.newamerica.org/political-reform/briefs/harnessing-technology-for-effective-co-governance> [<https://perma.cc/AWA2-9DYE>].

⁷⁷ To be sure, examples of co-governance at the local level abound. See, e.g., Hollie Russon Gilman et al., *Rural Partners Network: Connecting Local Community to Government: An Interview with U.S. Department of Agriculture Rural Development*, NEW AM. (Sept. 29, 2023), <https://www.newamerica.org/political-reform/briefs/rural-partners-network-connecting-local-community-to-government> [<https://perma.cc/A38T-7RVZ>] (discussing the “Rural Partners Network,” which “is an all-of-government program that helps people in rural communities find resources and funding to create jobs, build infrastructure, and support long-term economic stability on their own terms”).

⁷⁸ Indeed, for a piece of scholarship that squarely distinguishes between co-governance and local governance, see generally Asbjørn Røiseland, *Local Self-Government or Local Co-Governance?*, 8 LEX LOCALIS-J. LOC. SELF-GOV’T 133 (2010).

⁷⁹ Dave Owen, *Regional Federal Administration*, 63 UCLA L. REV. 58, 91 (2016).

⁸⁰ *Id.*

⁸¹ *Id.* at 60.

⁸² Orly Lobel, *The Renew Deal: The Fall of Regulation and the Rise of Governance in Contemporary Legal Thought*, 89 MINN. L. REV. 342, 345 (2004).

⁸³ See *infra* notes 105–37 and accompanying text.

⁸⁴ See Owen, *supra* note 79, at 60 (“[The] conventional wisdom [regarding federalism] holds that states and municipalities are the classic ‘laboratories of democracy’ . . .”).

⁸⁵ See *supra* note 8 and accompanying text (discussing a local government meeting where a citizen felt her voice wasn’t heard).

and legislate in a manner that still ices out members of the political community.⁸⁶ When local governments function this way, they are not facilitating co-governance.

But more importantly, co-governance borrows from federalism its *ends*, not its *means*. Decentralization and federalism have often been invoked in contexts where local communities have the largest stake. Consider the local school. Schools play a role in shaping the residents of a town; they have a significant impact on the lives of the town's residents. Accordingly, federalism suggests that those residents should have a significant say in the goings-on of the school — that is, they should have a voice in the school's governance, so governance should take place at the local level. Put simply, one value underpinning some theories of federalism is that governance power should be put in the hands of those most affected by a particular policy or entity. When it comes to things like schooling, folks are most impacted at the local level, so governance should take place at the local level, too. But in this regard, federalism and localization are not ends in and of themselves; they just offer one method of dispersing power among bona fide stakeholders.

What does this mean for AI? As discussed above, governance should be informed and shaped by those who will be most impacted and affected by the policies ultimately produced. In the context of AI, this is everyone — regulation must be done by the people. But not the people of a local town or a single state. AI will impact us all, so we all need a seat at the governance table. Federalism and localization are important but ultimately insufficient in this regard because AI is unlike a school — its reach is far more sprawling and cross-jurisdictional. But co-governance has the capacity to empower a much broader swath of those who will be affected by AI to chart its future.

2. *Virtues of Co-Governance.* — Already, in contexts other than AI, the virtues of co-governance have been well established in the literature. This section outlines these virtues to illustrate what it is that an effective co-governance framework can achieve — and has *already* achieved in other areas. Ultimately, the benefits of co-governance that have made the framework effective and desirable in other contexts are benefits that would do well to meet the challenges of AI, too.

(a) *Better Decisions and Outcomes.* — One benefit of co-governance is that “[s]hared decision making contributes to better decisions.”⁸⁷ That might be the case because those involved in co-governance can “draw on different expertise, knowledge and experiences, and from the enhanced legitimacy of processes that include specific communities of interest and appropriately recognise relevant

⁸⁶ See Ackerman, *supra* note 74, at 459 (noting that “devolution and decentralization are important because they bring government closer to the people” but still recognizing some potential drawbacks).

⁸⁷ Jones, *supra* note 57, at 4.

rights and obligations.”⁸⁸ Or that might be true because some co-governance frameworks include features, like “the scrutinizing assembly,” that allow citizens to “reflect on the content and consequences of the ballot options and to inform and enlighten the wider public on the matter before voting takes place.”⁸⁹

In the same vein, co-governance has been linked to greater “efficien[cy] and effective[ness].”⁹⁰ One scholar has “documented how the involvement of citizens in the planning and implementation of water and sanitation projects greatly improved their effectiveness and reduced corruption in urban Brazil.”⁹¹ Another commentator “showed how community participation in irrigation programs in Taiwan has made service delivery much more efficient and effective.”⁹² These observations are also “consistent with [the work of a notable scholar who] demonstrated the salutary effects of the coproduction of services by street-level bureaucrats and societal actors.”⁹³ Co-governance likely facilitates better outcomes in instances like these because “[c]itizens [are able to] contribute local knowledge and experience that would be prohibitively costly for outsiders to acquire.”⁹⁴ Moreover, “[a]s the beneficiaries of the final product[,] community members can also contribute their time [at a level] that public employees should not be [expected] to match.”⁹⁵ More effective rules are developed when the people most meaningfully affected by them are involved in their formation. All in all, co-governance — which seeks to bring *all* stakeholders to the table — leads to decisions that are more efficacious for those stakeholders.

(b) *More Meaningful Participation.* — When the people are given avenues to govern more directly, participation in government increases. In fact, in co-governance systems, “[e]ven the poorest citizens are exceptionally willing and able to actively work with government in constructive ways once they perceive that their participation can make a difference.”⁹⁶ Indeed, “[n]ormal citizens . . . participate at such massive levels if the policies being implemented . . . are designed with the participation of a broad range of actors and actively incorporate citizens

⁸⁸ *Id.* at 8.

⁸⁹ Frank Hendriks & Charlotte Wagenaar, *The Deliberative Referendum: An Idea Whose Time Has Come?*, 55 ADMIN. & SOC’Y 569, 576 (2023).

⁹⁰ Ackerman, *supra* note 74, at 449.

⁹¹ *Id.* (citing Elinor Ostrom, *Crossing the Great Divide: Coproduction, Synergy, and Development*, 24 WORLD DEV. 1073 (1996)).

⁹² *Id.* (citing Wai Fung Lam, *Institutional Design of Public Agencies and Coproduction: A Study of Irrigation Associations in Taiwan*, 24 WORLD DEV. 1039 (1996)).

⁹³ *Id.* at 449–50 (citing JUDITH TENDLER, *GOOD GOVERNMENT IN THE TROPICS* (1997)).

⁹⁴ Peter Evans, *Government Action, Social Capital and Development: Reviewing the Evidence on Synergy*, 24 WORLD DEV. 1119, 1130 (1996).

⁹⁵ *Id.*

⁹⁶ John M. Ackerman, *From Co-Production to Co-Governance*, in *NEW PUBLIC GOVERNANCE, THE THIRD SECTOR AND CO-PRODUCTION* 101, 101 (Victor Pestoff et al. eds., 2012); see also *id.* at 109.

into the process of implementation itself.”⁹⁷ That makes sense. When people subjectively believe that their voices matter, they’ll be more inclined to speak up.⁹⁸ And when they speak up — especially if they have been “inform[ed] and enlighten[ed]” on the topic at issue *before* participating⁹⁹ — more “meaningful societal debate on the issue” ensues.¹⁰⁰

Co-governance, then, brings *more* people to the table. And when they arrive there, they are in a better position to debate, discuss, and decide the issues that affect them most.

(c) *Feeling Heard*. — Relatedly, when people are made a part of governing, they subjectively feel more visible. They feel that their voices are heard and ideas considered. Indeed, “the ability to express that voice at individual and collective level[s] . . . is an essential tenet of democracy.”¹⁰¹ “Institutional actors [have] recognized the contribution of community engagement to private value for the various actors of the public service ecosystem: a therapeutic value of feeling heard, feeling listened to.”¹⁰² In this way, co-governance doesn’t just give citizens more opportunities to speak — the frameworks make them feel heard. Just like the virtue mentioned above — participation — this benefit of co-governance promotes the democratic values mentioned in section A of this Chapter: Public participation is increased and made more substantive.

(d) *Building Consensus*. — Insofar as co-governance brings community members together to discuss (and debate) policy, it also motivates them to see each other as equal members of the polity worthy of respect. As one individual involved in a co-governance exercise noted: “There was so much ownership over the process. Even though there wasn’t consensus over the proposed projects, people supported one another. In the beginning, people were uncomfortable and scared to participate given the heightened political polarizations, but in the end, we are all still neighbors.”¹⁰³ In this way, co-governance can come to look a lot like “‘kitchen-table’ conversations,” where community members “break down complicated and controversial policy topics with tailored framing for different audiences to build common ground and meet residents where they are.”¹⁰⁴ Co-governance, then, doesn’t just promote

⁹⁷ *Id.* at 111.

⁹⁸ See NICK TROIANO, THE PRIMARY SOLUTION 162 (2024) (noting that when Alaskan primary voters “believed that their vote[s] ‘mattered more,’” voter turnout increased); see also RUSSON GILMAN ET AL., *supra* note 64, at 17.

⁹⁹ See *supra* note 89 and accompanying text.

¹⁰⁰ Hendriks & Wagenaar, *supra* note 89, at 576.

¹⁰¹ Adina Dudau et al., *Enabling PSL and Value Co-Creation Through Public Engagement: A Study of Municipal Service Regeneration*, PUB. MGMT. REV., May 2023, at 17.

¹⁰² *Id.* at 16–17.

¹⁰³ RUSSON GILMAN ET AL., *supra* note 64, at 23 (quoting Peggy Flynn, City Manager of Petaluma, California).

¹⁰⁴ *Id.* at 10 (citing a conversation with Deborah Scott, Executive Director of Georgia STAND-UP).

democratic deliberation — it facilitates deliberation that leads to actual decisions being made. It helps make democracy work.

3. *Applications of Co-Governance.* — Co-governance is already practiced in myriad contexts. This section describes three: participatory budgeting, rural development, and elder care. These may seem a far cry from artificial intelligence; AI is a modern technology with global implications, while the contexts explored here are traditionally areas of local concern. But they have something in common: Participatory budgeting, rural development, and elder care are all contexts in which people are understood to have a vested personal interest in the effects of regulation. That is, people will feel the effects of regulatory decisions on the ground. And in these cases, the government has acknowledged this interest by giving people a seat at the regulatory table. AI isn't a local issue, but just as in these other contexts, people will feel the effects of regulatory decisions in their everyday lives. Thus, understanding where co-governance is already deployed makes it clear why the framework is, in fact, a good fit for regulating AI, too.

(a) *Participatory Budgeting.* — One “increasingly popular model of co-governance” is something called participatory budgeting,¹⁰⁵ which usually entails “local residents [having the power] to decide how city funds get allocated by voting on a number of community proposals.”¹⁰⁶ Participatory budgeting puts taxpayers in the driver's seat. They call the shots and have a say in how their own money is allocated within the community. These sorts of arrangements are now common in the United States,¹⁰⁷ even in major cities like New York City.¹⁰⁸ But they're also used in other parts of the world. In fact, one study assessed the participatory budgeting system in Porto Alegre, Brazil.¹⁰⁹ There, “spending decisions for over 10% of [the city's] annual budget [were put] in the hands of the people.”¹¹⁰ In practice, that meant that “[e]very year, more than 14,000 citizens in th[e] city of 1.3 million participate[d] in neighborhood meetings as well as 16 regional and five thematic assemblies to set priorities for government investment in infrastructure and basic social services.”¹¹¹ In Porto Alegre, citizens took their role seriously; budgeting “decisions [were] made through intense negotiation.”¹¹² And after the “local legislature” signed off on their choices, the citizen “groups

¹⁰⁵ *Id.* at 22.

¹⁰⁶ *Id.*

¹⁰⁷ Hollie Russon Gilman & Lizbeth Lucero, *Participatory Budgeting Is About to Blow Up Across the Country*, THE FULCRUM (Feb. 15, 2023), <https://thefulcrum.us/participatory-budgeting> [<https://perma.cc/3T3S-VTJJ>].

¹⁰⁸ Russon Gilman et al., *supra* note 76.

¹⁰⁹ Ackerman, *supra* note 74, at 451.

¹¹⁰ *Id.*

¹¹¹ *Id.*

¹¹² *Id.*; see also Russon Gilman et al., *supra* note 76 (“Most of the ideas that we generated were generated through workshops where people could actually discuss them and have a dialogue versus staying online, alone.”).

[continued to] evaluate the previous year's negotiation process and monitor the implementation process of the previous year's budget."¹¹³

By allowing "the citizens of Porto Alegre . . . inside the governmental apparatus itself,"¹¹⁴ "corrupt[ion]" decreased, the "political use of public funds" was "reduce[d]," and regulatory "capture" occurred less frequently.¹¹⁵ Citizen participation also improved.¹¹⁶ In other words, a decentralized co-governance approach to budgeting brought with it a panoply of benefits. And that makes sense. When taxpayers themselves "have a direct say on how" their own money "will be allocated in their communities,"¹¹⁷ good things happen.

(b) *Rural Development*. — Local communities, especially rural ones, can feel disconnected from the federal government. Given their bureaucratic layers and red tape, federal solutions to local issues are often "one-size-fits-all"¹¹⁸ — that is, they lack the flexibility and tailoring to effectively cure a local community's unique problems. The federal government often resorts to just telling "communities how their project can fit within the federal government's box."¹¹⁹

To address these concerns in the context of rural development, the United States Department of Agriculture has launched the Rural Partners Network (RPN).¹²⁰ The program "is an all-of-government program that helps people in rural communities find resources and funding to create jobs, build infrastructure, and support long-term economic stability on their own terms."¹²¹ To achieve its mission, the RPN partners with "community liaisons": "staff in rural communities whose sole job is to listen first and work with all the local partners."¹²² The RPN has also "established [an] interagency group of representatives from over 24 federal agencies who are tasked with listening to the communities and hearing them [speak] firsthand."¹²³ In short, the RPN has developed programs and tools to address the federal bureaucracy's lack of understanding and attentiveness to agricultural and rural needs.

One example of the RPN in action took place in southeastern Kentucky, which had been ravaged by a "devastating flood."¹²⁴ The Federal Emergency Management Agency (FEMA) responded, but an RPN

¹¹³ Ackerman, *supra* note 74, at 451.

¹¹⁴ *Id.*; see also Russon Gilman & Lucero, *supra* note 107 ("They work intimately in collaboration with local governments in deciding their communities' priorities through public spending.")

¹¹⁵ Ackerman, *supra* note 74, at 451.

¹¹⁶ *Id.* at 452; see also RUSSON GILMAN ET AL., *supra* note 64, at 22; Russon Gilman et al., *supra* note 76; Russon Gilman & Lucero, *supra* note 107.

¹¹⁷ Russon Gilman & Lucero, *supra* note 107.

¹¹⁸ Russon Gilman et al., *supra* note 77.

¹¹⁹ *Id.*

¹²⁰ *Id.*

¹²¹ *Id.*

¹²² *Id.*

¹²³ *Id.*

¹²⁴ *Id.*

community liaison — a member of the community who had herself been affected by the flood — “buil[t] a bridge between FEMA and the impacted community.”¹²⁵ Having intimate knowledge of the needs of the community, the liaison spearheaded an effort to secure grants for rebuilding destroyed homes.¹²⁶ She worked to “increase [the] cap” normally available to those affected by natural disasters.¹²⁷ “[H]aving [a] direct line of communication that[] [was] rooted in a community’s need[s]”¹²⁸ redounded to the benefit of those most affected.

(c) *Elder Care*. — Taking care of the elderly is hard. It’s not difficult to imagine why clunky government bureaucracies wouldn’t be able to provide individualized care for older citizens. That’s likely why England embraced a system “based on co-governance and negotiated public policies.”¹²⁹ “[R]esponsibility for providing care services was devolved to Local Authorities in the late 1960s,”¹³⁰ such that “[h]ome care . . . consisted of nursing, guaranteed by the local NHS, and (some) personal care organized by Local Authorities (mainly for low-income earners).”¹³¹ Some “Local Authorities had their own in-house providers, [and] a good deal of personal care was provided by publicly funded voluntary agencies.”¹³² The system of co-governance of elder care in England has changed over time,¹³³ but it gave the “voluntary sector” an opportunity to “lobby[] for better market regulation.”¹³⁴ In other words, it allowed them to “hav[e] their civic voice heard in the public sphere.”¹³⁵ And today, “Partnership Boards”¹³⁶ still meet, “which allow[s] the voices of several stakeholders to be heard at least to some extent.”¹³⁷

* * *

In each of these contexts, co-governance has been deployed to give the people most directly impacted by regulation a seat at the table where regulatory decisions are made. In this way, co-governance eschews the concentration of power and rigidity that defines the traditional top-down approach to regulation: It diffuses power from the institutions “up top” to the people on the ground and opens the door for more flexible

¹²⁵ *Id.*

¹²⁶ *Id.*

¹²⁷ *Id.*

¹²⁸ *Id.*

¹²⁹ Ingo Bode, *Co-Governance Within Networks and the Non-Profit-For-Profit Divide: A Cross-Cultural Perspective on the Evolution of Domiciliary Elderly Care*, 8 PUB. MGMT. REV. 551, 559 (2006).

¹³⁰ *Id.*

¹³¹ *Id.* at 559–60.

¹³² *Id.* at 560.

¹³³ *See id.*

¹³⁴ *Id.* at 561.

¹³⁵ *Id.* at 562.

¹³⁶ *See id.* at 560.

¹³⁷ *Id.* at 561.

solutions to complex problems. It also swaps out one-off opportunities for civic participation in exchange for continued and consistent collaboration. In short, the vignettes described in this section show how co-governance can be — and has been — deployed; they offer a blueprint to imagine a comparable regulatory framework for AI that, as in these other contexts, works better for more people.

E. Co-Governance Practices and AI

Applying co-governance practices in the context of AI isn't a radical idea. Indeed, co-governing AI would be a continuation of what is already a pattern of collaboration in the AI regulatory field. Various industry actors have already implemented collaboration mechanisms, and discussions about the regulation of AI — open and closed — have been marked by an interest in stakeholder collaboration.

Take, for example, the Biden Administration's approach to developing an AI safety framework, which often involved calling upon stakeholders for input. In President Biden's October 2023 Executive Order, he "tasked the Secretary of Commerce . . . with soliciting feedback 'from the private sector, academia, civil society, and other stakeholders through a public consultation process.'"¹³⁸ "[T]he National Telecommunications and Information Administration (NTIA) published a public Request for Comment in February 2024 and received 332 comments in response."¹³⁹ It also "conducted extensive stakeholder outreach" in other forums, "including two public events gathering input from a range of policy and technology experts."¹⁴⁰ The NTIA's final report, also required by the October 2023 Executive Order, was "based in large part on [that] feedback."¹⁴¹ The Biden Administration's foundational *Blueprint for an AI Bill of Rights* was likewise developed as a "[r]espon[se] to the experiences of the American public . . . informed by insights from researchers, technologists, advocates, journalists, and policymakers."¹⁴² No matter what comes next, the future of AI regulation in the United States will have been shaped at the ground level by at least some community input.

Co-governance is an extension of this commitment to collaboration and community interest. Unlike traditional top-down approaches, co-

¹³⁸ NAT'L TELECOMMS. & INFO. ADMIN., U.S. DEP'T OF COM., DUAL-USE FOUNDATION MODELS WITH WIDELY AVAILABLE MODEL WEIGHTS 2 (2024) [hereinafter DUAL-USE FOUNDATION MODELS] (quoting Exec. Order No. 14,110 § 4.6(a), 3 C.F.R. 657, 672 (2024)), <https://ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf> [<https://perma.cc/Z7S5-ET8C>].

¹³⁹ *Id.* at 3; *see* Dual Use Foundation Artificial Intelligence Models with Widely Available Model Weights, 89 Fed. Reg. 14059 (Feb. 20, 2024).

¹⁴⁰ DUAL-USE FOUNDATION MODELS, *supra* note 138, at 3.

¹⁴¹ *Id.*

¹⁴² *Blueprint for an AI Bill of Rights*, WHITE HOUSE, <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights> [<https://perma.cc/9AGC-3Z2B>].

governance ensures not only that interests communicated in the early stages of regulation make their way into final rules but also that rules remain open to reevaluation and readjustment. It's impossible that AI technology will itself remain stagnant; the regulatory approach should keep in step. As the Biden Administration put it: "AI only works when it works for all of us."¹⁴³ Co-governance is one way to ensure that it continues to do so — that a diverse array of stakeholders is able to "make [their] voice[s] heard," early on and ever after.¹⁴⁴

In these nascent stages of AI development, however, co-governance may also help ensure that the industry remains open to competition and innovation. That is, promoting collaboration and co-governance — especially with an eye toward supporting open-source AI development — may promote a culture of competition that encourages more diverse industry players to participate in development and innovation. As discussed in section B, resource constraints make foundation model development an impractical endeavor for many industry actors.¹⁴⁵ Thus, the major and most powerful foundation models are likely to remain within the control of "a cluster of big tech firms."¹⁴⁶ But if those models are kept under lock and key as "black box" systems, "[e]veryday developers and small businesses" will rarely — if ever — have meaningful opportunities to "create new AI applications, tune safer AI models for specific tasks, train more representative AI models for diverse communities, or launch new AI ventures."¹⁴⁷ This kind of "[g]rassroots innovation may," in other words, "become collateral damage" under the current regulatory system.¹⁴⁸

This is worrisome not only because it jeopardizes the development of better and more diverse AI tools but also because it prompts concerns about something akin to regulatory capture.¹⁴⁹ Regulatory capture "typically refers to a phenomenon that occurs when a regulatory agency that is created to act in the public interest, instead advances the commercial or political concerns of special interest groups that dominate an industry or sector the agency is charged with regulating."¹⁵⁰ In the current AI

¹⁴³ *Administration Actions on AI*, AI.GOV, <https://ai.gov/actions> [<https://perma.cc/V3D5-ZBC8>].

¹⁴⁴ *Id.*

¹⁴⁵ See *supra* notes 24–27 and accompanying text.

¹⁴⁶ Brooks, *supra* note 16.

¹⁴⁷ *Id.* (explaining that these endeavors would require "spending tens of millions of dollars").

¹⁴⁸ *Id.*

¹⁴⁹ See Matt O'Brien, *AI's Future Could Be "Open-Source" or Closed. Tech Giants Are Divided as They Lobby Regulators*, AP NEWS (Dec. 5, 2023, 9:56 AM), <https://apnews.com/article/ai-open-source-meta-ibm-chatgpt-dd61e99ac8135b36872b3987601067ec> [<https://perma.cc/LY3P-LP8M>] (describing corporate lobbying against open-source AI as "sort of a classic regulatory capture approach of trying to raise fears about open-source innovation" (quoting Chris Padilla, then-executive at IBM)).

¹⁵⁰ *Regulatory Capture*, CFA INST. RSCH. & POL'Y CTR. (Oct. 29, 2019), <https://rpc.cfainstitute.org/en/policy/positions/regulatory-capture> [<https://perma.cc/DTQ2-99G3>]; see also Kevin Wei et al., *How Do AI Companies "Fine-Tune" Policy? Examining Regulatory Capture in AI Governance*, at i (Aug. 20, 2024) (unpublished manuscript) (on file with the Harvard Law School Library).

ecosystem, then, fears of regulatory capture amount to fears that the few already dominant AI companies — like Microsoft, Google, and, especially, OpenAI¹⁵¹ — will “write the rules governing this technology,” which “could have a number of harms, from stifling smaller firms to introducing weak regulations.”¹⁵² Ultimately, this “could result in general-purpose AI policies and enforcement practices that are ineffective, unsafe, or unjust — or even no regulation at all.”¹⁵³

These fears rose to the surface in one account of the 2023 Senate hearing on AI, which was described as “affable” and “chumm[y].”¹⁵⁴ There, the “[i]ndustry reps — primarily OpenAI CEO Sam Altman — merrily agreed on the need to regulate new AI technologies, while politicians seemed happy to hand over responsibility for drafting rules to the companies themselves,” leading “[a] number of experts and industry figures” to conclude that “we may be headed into an era of industry capture in AI.”¹⁵⁵ “As Senator Dick Durbin (D-IL) put it in his opening remarks: ‘I can’t recall when we’ve had people representing large corporations or private sector entities come before us and plead with us to regulate them.’”¹⁵⁶ While Altman himself dismissed concerns about capture,¹⁵⁷ others — including leaders at competing AI companies — “stressed the potential threat to competition,”¹⁵⁸ explaining that the kinds of regulations proffered at the hearing “invariably favour[] incumbents”¹⁵⁹ and “would further concentrate power in the hands of a few [and] drastically slow down progress, fairness[, and] transparency.”¹⁶⁰

Co-governance counteracts such a regulatory culture of suppression by aligning restrictions on innovation with expressed community interests. Critically, “not all industry influence reflects regulatory capture.”¹⁶¹ And indeed, “industry participation in policy processes,” including co-governance, “is both inevitable and desirable,”¹⁶² so “[c]apture occurs only when corporate influence leads to regulation that unjustly prioritizes private interests over public ones.”¹⁶³ Co-governance helps to ensure this doesn’t happen. As some scholars suggest, “[i]nformation

¹⁵¹ James Vincent, *The Senate’s Hearing on AI Regulation Was Dangerously Friendly*, THE VERGE (May 19, 2023, 8:53 AM), <https://www.theverge.com/2023/5/19/23728174/ai-regulation-senate-hearings-regulatory-capture-laws> [<https://perma.cc/7ASX-7ZYF>] (“Although . . . OpenAI is still called a ‘startup’ by some, it’s arguably the most influential AI company in the world.”).

¹⁵² *Id.*

¹⁵³ Wei et al., *supra* note 150 (manuscript at 1).

¹⁵⁴ Vincent, *supra* note 151 (emphasis omitted).

¹⁵⁵ *Id.*

¹⁵⁶ *Id.*

¹⁵⁷ *Id.* (“We don’t wanna slow down smaller startups. We don’t wanna slow down open source efforts.” (quoting Sam Altman, CEO of OpenAI)).

¹⁵⁸ *Id.*

¹⁵⁹ *Id.* (quoting Emad Mostaque, founder and CEO of Stability AI).

¹⁶⁰ *Id.* (quoting Clem Delangue, CEO of Hugging Face).

¹⁶¹ Wei et al., *supra* note 150 (manuscript at 1).

¹⁶² *Id.*

¹⁶³ *Id.*

capture can be addressed through giving non-industry stakeholders greater access to the policy process.”¹⁶⁴ To this end, some have suggested “participatory processes such as notice-and-comment” or “consumer empowerment programs [that] could help enable civic participation in AI policy.”¹⁶⁵ This Chapter doesn’t seek to provide an exhaustive account of what AI co-governance might look like. But it’s safe to say it could include the same mechanisms that have popped up in other contexts: citizen assemblies (that are commonly used in the participatory-budgeting context¹⁶⁶), governmental liaisons connected to local communities (exemplified in the RPN and elder care case studies above¹⁶⁷), “scrutinizing assembl[ies],”¹⁶⁸ and something akin to notice-and-comment.

Regardless of which of these practices are ultimately adopted, co-governance seeks to place these mechanisms within a broader and more comprehensive framework, thus assuring that participation occurs not only at “the very early stage[s]”¹⁶⁹ of the policy process but also throughout the regulatory lifecycle. Perhaps it is the case that people generally want AI to be subject to regulations that deter further innovation, just as current skeptics fear will happen if today’s “tech giants” get to call the shots.¹⁷⁰ But co-governance produces that decision through a process that better reflects the interests of the diverse stakeholders involved.

In this sense, co-governance is directly reflective of at least one segment of the AI industry — that is, open source. Recall that proponents of open-source AI celebrate it for its accessibility and, thus, its ability to “promote a diverse AI ecosystem.”¹⁷¹ It is at least imagined to offer a more democratic future for AI.¹⁷² Co-governance pulls this thread — what is possibly the best part of open source — and weaves it into a broader and more robust framework that aims to strengthen the foundations for AI development, generally. That is, co-governance seeks to make possible an iterative relationship between the technology and its positive impacts; as the technology develops, so, too, may the values it

¹⁶⁴ *Id.* (manuscript at 13).

¹⁶⁵ *Id.*

¹⁶⁶ See *supra* notes 105–17 and accompanying text.

¹⁶⁷ See *supra* notes 118–37 and accompanying text.

¹⁶⁸ These allow citizens to “reflect on the content and consequences of the ballot options and to inform and enlighten the wider public on the matter before voting takes place.” Hendriks & Wagenaar, *supra* note 89, at 576.

¹⁶⁹ Wei et al., *supra* note 150 (manuscript at 13).

¹⁷⁰ Vincent, *supra* note 151.

¹⁷¹ Brooks, *supra* note 16; see *supra* p. 1613.

¹⁷² See Divya Siddarth et al., *A Vision of Democratic AI*, DIGITALIST PAPERS, <https://www.digitalistpapers.com/essays/a-vision-of-democratic-ai> [<https://perma.cc/M4E2-4TQZ>] (explaining that “[t]he open-source community can play a large role” in “[t]uning AI with [c]ollective [i]nput” because “[t]ransparency and open innovation is beneficial for enabling trust in the results and wider participation, while the opportunity that open source poses for comparatively rapid and dispersed experimentation and iteration can increase the rate of learning”).

embodies become more deeply entrenched.¹⁷³ In the context of open-source AI, this means committing to a regulatory framework that allows accessibility and transparency to flourish; the same values that underpin the technology motivate its governance. But the benefits of this approach need not be limited to one side of the open–closed spectrum. To the contrary, adopting co-governance makes possible a culture of AI that better supports the democratic values underlying open-source AI and our political community, more broadly.

Conclusion

All in all, what co-governance offers is a much-needed alternative to one-size-fits-all, top-down, centralized regulation. In the AI era, embracing this kind of institutional innovation is more important than ever. The effects of AI are predicted to be ubiquitous — to impact each and every one of us in fundamental ways. If that’s the case, then it makes sense for this technology to be governed not by only a few members of our political community but by as diverse and representative a group as possible. This group should not be composed of only legislators and those whose voices are already most familiar to our existing regulatory system: the most powerful companies and most influential experts. It should also include — and, in fact, be *most* responsive to — the people.

Implementing co-governance is an opportunity to achieve this end. Co-governance practices have already been pressure tested and have shown success in other contexts where policy issues have a direct impact on people’s lives. It is an approach that is well aligned with this country’s long-held democratic values; it gives people a meaningful opportunity to have their voices heard. And — despite the global nature of AI’s projected effects — it is well suited for addressing the challenges posed by the technology. Already, the same democratic values that inspired the adoption of co-governance in other contexts have manifested in the conversation about AI, most prominently in the debate about open- versus closed-source systems. This debate is a reaffirmation of the collective commitment to democratic values of accessibility, transparency, and public participation in the AI era. Co-governance is an opportunity to make those values salient; it ensures that they — and the people for whom they matter — shape the future of AI, and not vice versa.

¹⁷³ AI is already being used to facilitate public deliberation and “promote inclusive community-level engagement within wider democratic processes.” *About Us*, COMMON GOOD AI, <https://commongoodai.org/mission> [<https://perma.cc/8JVV-47VC>]; see also, e.g., ANDY PAICE & MARTIN RAUSCH, CREATING AN ONLINE CONVERSATION BETWEEN A NATION AND A MINI-PUBLIC: A CASE STUDY ON POLIS & THE AUSTRIAN CITIZENS’ CLIMATE COUNCIL 35, <https://www.co-intelligence.institute/case-study-polis-climate-council> [<https://perma.cc/H9W8-J8YG>]; Chris Horton, *The Simple but Ingenious System Taiwan Uses to Crowdfund Its Laws*, MIT TECH. REV. (Aug. 21, 2018), <https://www.technologyreview.com/2018/08/21/240284/the-simple-but-ingenious-system-taiwan-uses-to-crowdfund-its-laws> [<https://perma.cc/WE3K-JXES>].