

# Contents

<b>The Technological Republic, Reconsidered</b>	<b>1</b>
Preface . . . . .	1
I. The diagnosis . . . . .	2
II. The slide . . . . .	3
III. The historical claim . . . . .	5
IV. The empirical record . . . . .	8
V. China . . . . .	9
VI. The moral question . . . . .	11
VII. The reconstruction . . . . .	13
A coda on tone . . . . .	15

## The Technological Republic, Reconsidered

*A full Cube Commons response to Alexander C. Karp and Nicholas W. Zamiska’s The Technological Republic: Hard Power, Soft Belief, and the Future of the West (Crown Currency, February 2025)*

**Cube Commons Technical Report CC-TR-2026-005-D** Barton Nicholls · April 2026

---

### Preface

Alexander Karp and Nicholas Zamiska have written a serious book, and it deserves a serious response. *The Technological Republic* is not, despite some of the more dismissive reviews, a Palantir promotional document with a philosophical veneer. It is an argument — coherent, ambitious, well-written, and substantially wrong — about the relationship between software, power, and democratic legitimacy. It diagnoses a real problem in contemporary American technology culture and prescribes a remedy that, if adopted at scale, would foreclose the institutional architecture this report exists to defend.

The shorter response to Karp and Zamiska is the position paper *The Other Answer* (CC-TR-2026-005), which sets the broad alternative argument without engaging the book paragraph by paragraph. The shorter response sufficed for the audiences that needed it: legislators, grant officers, civic-tech peers. This longer response is for the audience that has actually read the book or is considering doing so, and that wants to understand both what is right about it and what is wrong about it without being asked to choose between treating it as gospel or treating it as nonsense. Neither response is correct. The book is a sophisticated wrong answer to an important question.

This essay proceeds in seven parts. Part I takes the diagnosis seriously: Silicon Valley *has* lost something real, and Karp and Zamiska name the loss with more precision than most of their critics give them credit for. Part II identifies the move where the argument goes wrong — the slide from “Silicon Valley should rebuild a relationship with its host society” to “concentrated state-aligned firms should be that relationship.” Part III addresses the historical claim that the actually-foundational vision of computing was state-defense fusion, and shows that the actually-foundational vision was something else. Part IV walks through the empirical record on the architectural model the book defends, which is losing in real time in the layer where it operates. Part V engages the China question as the book frames it and shows why the framing produces the wrong policy. Part VI takes up the moral question — the one Karp keeps circling and Zamiska keeps formalizing — about engineering competence and civic obligation. Part VII proposes what the book’s argument would look like if it were rebuilt around the institutions that actually exist in democratic societies, rather than around a fusion of two institutions that, jointly, encompass only a small fraction of democratic life.

This is the longest single Cube Commons document so far. It exists because the alternative — letting *The Technological Republic* shape the AI infrastructure conversation by default — would foreclose work that needs to happen in the next several years. The book is now influential enough, in Washington and London and Brussels, that civic technologists who don’t engage it find themselves operating inside its assumptions whether they meant to or not. This is the engagement.

---

## **I. The diagnosis**

The book opens with a claim that, on the surface, is hard to disagree with. Silicon Valley, Karp and Zamiska argue, has lost something. The early American software industry was built in partnership with the United States government — through Vannevar Bush’s Office of Scientific Research and Development, through the Manhattan Project, through DARPA, through the SAGE air defense system and the ARPANET. That partnership produced the foundational infrastructure of modern computing. By the 2010s, the partnership had hollowed out. Google publicly withdrew from Project Maven in 2018 after employee protest. Major Silicon Valley firms declined defense contracts. The talent flow into civic and military problems narrowed. The talent flow into consumer applications, advertising optimization, and short-form video widened. A generation of engineers who could have been working on the problems Vannevar Bush organized in 1940 — public health, national defense, civic infrastructure — chose instead to work on persuasion at scale.

The diagnosis is not new. Eric Schmidt and Robert Work made versions of it through

the National Security Commission on Artificial Intelligence (2018–2021). The American Affairs / Tablet right has been making versions of it for a decade. Jaron Lanier has been making versions of it for longer. What Karp and Zamiska contribute is the rhetorical force of insiders saying it: the founders and senior executives of one of the few Silicon Valley firms that *did* maintain a serious relationship with national-security work, telling the rest of the industry that the relationship is constitutive and that abandoning it has been a moral failure.

Several elements of their diagnosis are correct. There has been a real cultural shift in Silicon Valley away from what one might call civilizational seriousness. The center of gravity of the industry, measured by market capitalization and engineering attention, has moved from problems-of-state to problems-of-attention. The metaphysical commitments of the technology class have, on average, narrowed. Large numbers of capable engineers do work that is, by any reasonable measure, frivolous. The compensation structure rewards the frivolity. The cultural prestige structure inside the industry, until recently, also rewarded the frivolity. Project Maven was a bellwether: the question was whether engineers would help build automated targeting systems for the United States military, and the answer that emerged from significant Google employee organizing was *no*. Whatever one thinks of that specific project, the meta-question — whether American software engineers as a class consider themselves part of the American political community in any thick sense — was answered uncomfortably.

Karp and Zamiska are also right that the response of the industry to the events of the early 2020s was, in many cases, embarrassing. The combination of frantic ESG performance with a refusal to engage on hard infrastructural questions produced a culture of moral preening unconnected to the actual stakes of either the climate crisis or the geopolitical environment. The book is sharper than its critics on this point. There is a real failure mode in which a class of people in their twenties and thirties earning four hundred thousand dollars a year writing tracking pixels persuades itself that its work is morally neutral as long as the firm has good DEI metrics and the cafeteria food is sustainably sourced. That is a genuine cultural pathology, and the book diagnoses it accurately.

The diagnosis as far as it goes, then, is largely correct. Silicon Valley has experienced a thinning of civilizational commitment, a hollowing of the relationship between technical capability and public purpose, and an inflation of moralism over substance. The question is what follows from this diagnosis.

## **II. The slide**

Here is where the argument fails. Karp and Zamiska move, over the course of the book, from a defensible diagnostic claim (“Silicon Valley has lost connection to civic

purpose”) to a much stronger architectural claim (“the right form of that connection is partnership between Silicon Valley firms and the United States national-security state”). The slide happens gradually, across multiple chapters, and is rarely stated as a single proposition. But it is the load-bearing move of the entire book, and once you see it, the rest of the argument loses its grip.

Consider the structure carefully. The diagnostic claim is that the technology industry has detached from the political community it operates within. The architectural prescription is that the industry should rebuild a relationship with that community. So far, this is unobjectionable. The slide happens when “the political community” gets quietly redefined as “the United States national-security state and its allied institutions.” This is not the same thing. The political community in which any civic technology actually operates includes — inescapably — hospitals, schools, religious institutions, municipal governments, housing authorities, libraries, universities, trade unions, mutual aid networks, neighborhood associations, voluntary firefighting departments, and the ten thousand other intermediate institutions that make a democratic society more than an aggregation of individuals confronting a state. The national-security state is one institution among many. It is an institution of unusual size and budget, but it is not coextensive with the political community. Karp and Zamiska’s argument requires treating it as if it were.

This is the move that *The Other Answer* engaged most directly. *Distributed institutional sovereignty* is the framing for the alternative: the durable architecture of democratic civic life is the polycentric, nested, rule-governed network of intermediate institutions, and the right relationship between technology and that network is one in which each institution can hold its own data, run its own computation, and govern its own coordination. This is not a refusal of public purpose. It is a different account of what public purpose looks like at the level of architecture.

The Karp-Zamiska argument cannot accommodate this account. Its rhetorical structure requires a single Silicon Valley facing a single state, with the question being whether they are working together. There is no place in this picture for ten thousand civic institutions each operating their own infrastructure under their own governance, coordinating through open protocols, accountable to their own members. The book does not argue against distributed institutional sovereignty; it simply does not see it. The polycentric civic landscape is invisible from the vantage point the book takes up. A reader who comes to the book already inside that landscape — running a small civic technology project, working in a municipal government, building open-source civic infrastructure — will recognize the omission immediately. A reader who comes to the book from inside a venture-backed firm or a Pentagon contractor will not.

There is a further problem with the slide that is worth naming. The Karp-Zamiska framing makes the choice binary: engineering elites either serve the national-security state, or they retreat into commercial frivolity. This false binary erases the entire

third option that has, in practice, produced most of the world's actually-functional civic infrastructure: the open-source commons. Linux is not a Silicon Valley consumer product. Apache is not a national-security project. The Internet protocols were not built by a state-firm fusion. The Web is not the work of either a hyperscaler or a defense contractor. These institutions and the infrastructures they have built represent neither the frivolous Silicon Valley Karp and Zamiska criticize nor the state-aligned Silicon Valley they advocate. They represent a mode of organized engineering work, sustained over decades, that is structurally different from both — and that has, in the relevant scoreboard, outperformed both.

The book does not engage this third tradition. It is mentioned in passing, as background; it is not treated as a serious civilizational alternative. This is the book's most significant intellectual omission. A book about the relationship between technology and democratic purpose that does not seriously engage the open-source commons is missing the largest body of evidence on that exact question.

### **III. The historical claim**

Part of what makes the slide persuasive is the historical claim that supports it. Karp and Zamiska argue that the actually-foundational vision of computing was the state-defense partnership of the 1940s and 1950s — Vannevar Bush, the Manhattan Project, SAGE, ARPANET. The argument is that contemporary Silicon Valley has drifted from this foundation, and that returning to it is therefore a return to the discipline's authentic purpose.

The historical claim is roughly a third of the truth. There were state-defense partnerships in the early decades of computing, and they did matter. Bush's *Science: The Endless Frontier* (1945) is a genuinely important document. The Manhattan Project did produce significant computational infrastructure. SAGE did anchor a generation of systems engineering. The ARPANET, funded through DARPA, did become the Internet.

What the historical claim leaves out is that the *intellectual* foundation of computing — the body of ideas about what computers should do, who they should serve, and how human beings should relate to them — was substantially the work of figures who were either explicitly opposed to state-defense fusion or were thinking about computing as a tool for distributed augmentation rather than centralized command. The omission of these figures from the book's historical narrative is consequential because their actual ideas are the strongest available counter-argument to the book's prescription.

Norbert Wiener, the founder of cybernetics, wrote an *Atlantic Monthly* essay in January 1947 titled "A Scientist Rebels," refusing to share his work with the military. In August 1949, Wiener wrote to Walter Reuther, the president of the United Auto Workers, offering to help organized labor prepare for the labor displacement that

automated production would produce. Wiener's *The Human Use of Human Beings* (1950) is one of the foundational texts of computing as a discipline. It is also a sustained argument against exactly the kind of state-defense-industrial fusion Karp and Zamiska now propose as the discipline's authentic purpose. To call Wiener foundational and to omit his actual political commitments is to invent a different lineage.

J.C.R. Licklider's "Man-Computer Symbiosis" (1960) and his 1968 follow-up with Robert Taylor, "The Computer as a Communication Device," frame computing as a network of *autonomous* nodes augmenting human judgment. The end-to-end principle that Licklider's vision shaped — the principle that intelligence should live at the edges of the network, not in its center — is the foundational architectural commitment of the Internet. It is also the precise opposite of the architectural prescription Karp and Zamiska's book produces, which puts intelligence at the center and treats the edges as data sources.

Douglas Engelbart's 1962 framework "Augmenting Human Intellect" analyzed computing as a socio-technical ensemble in which the unit of analysis was not the individual user, the firm, or the state, but the recursive co-evolution of human practice and computational capacity. Engelbart's "Mother of All Demos" in December 1968 demonstrated, in real time, the working prototype of a vision that Silicon Valley would spend the next half-century failing to fully realize. The vision was not a state-defense vision. It was a vision of computing as a tool for collaborative human work — work that could happen in any institutional context that wanted to do it.

Ted Nelson's Xanadu (1960-) encoded at the protocol layer features the contemporary Web still lacks: two-way links, transclusion, fragment-level provenance, versioning. Nelson's vision was explicitly anti-centralization. The "failure" of Xanadu — its inability to become the dominant hypertext system — is one of the most consequential counterfactuals in computing history. The world we live in, with its centralized platform-mediated content and its broken-link ecology, is the world that emerged when the centralizing vision won.

Stafford Beer's Project Cybersyn in Allende's Chile (1971-1973) was nested-recursive computational coordination respecting subsidiarity by design. It was not a Western-aligned project; it was an explicitly socialist one, and it was destroyed by a military coup that the United States actively supported. Cybersyn is uncomfortable history for the Karp-Zamiska narrative, because it represents a serious engineering attempt to build computational coordination infrastructure that was neither corporate nor state-aligned-firm but something else — and that was destroyed by exactly the kind of state-firm fusion the book now prescribes as authentic.

Bill Joy, the co-founder of Sun Microsystems and the principal author of much of BSD Unix, published "Why the Future Doesn't Need Us" in *Wired* in April 2000 — an argument for restraint in the development of computational, biological, and nanotech-

nological capability. Joy is one of the great founding engineers of the post-ARPANET era. His political commitment, on the record, was to caution rather than acceleration.

Donald Knuth — the most distinguished computer scientist of the late twentieth century — has been clear throughout his career that his work is in the service of mathematical understanding, not state power. The TeX project, which Knuth gave to the world under one of the earliest open-source licenses, is a case study in how individual engineering excellence can produce durable civic infrastructure without any state-defense partnership.

Richard Stallman's launch of the GNU Project in September 1983, and the publication of the GNU Manifesto in *Dr. Dobbs's* in March 1985, was an explicit refusal of the proprietary-and-state-aligned model that Karp and Zamiska now defend. The GPL — the General Public License — is the legal instrument that uses copyright law to keep software permanently open. The fact that Linux, the operating system that runs most of the Internet, is licensed under the GPL is not incidental. It is constitutive. Without Stallman's legal innovation, the Internet would be running on proprietary infrastructure today.

Tim Berners-Lee invented the World Wide Web at CERN in 1989 and, with CERN's formal release of the Web into the public domain in April 1993, made one of the most consequential sovereignty-preserving infrastructure decisions in computing history. The counterfactual world — in which the Web was proprietary, owned by AOL or Microsoft or some other firm — is one that Karp and Zamiska's argument cannot easily condemn, because the Web they actually use was made by the opposite of the state-firm fusion they prescribe.

This is not a fringe alternative tradition. This is the actual intellectual mainstream of computing as a discipline. Wiener, Licklider, Engelbart, Nelson, Beer, Joy, Knuth, Stallman, Berners-Lee — these are not minor figures. They are the people whose work the contemporary technology industry is still living off. Their political commitments, where they made them explicit, were generally toward distributed augmentation, toward institutional autonomy, toward refusal of state-defense fusion, toward commons-preserving legal architecture, and toward the long-term survival of human practice over the short-term capture of human attention.

The book's historical claim treats this tradition as either invisible or peripheral. The result is a narrative in which Silicon Valley's drift from state-defense partnership is portrayed as a drift from its foundational purpose. The actual foundational purpose was something else, and the contemporary drift, where it has gone wrong, has gone wrong by abandoning that other foundation — not by abandoning the state-defense partnership.

## IV. The empirical record

The Karp-Zamiska prescription would be, in the kindest reading, a defensible bet on a particular architectural model if there were no live empirical record on whether that model produces durable infrastructure. There is a live empirical record. It is not kind to the Karp-Zamiska prescription.

Between 2018 and 2025, the layer of infrastructure where Cube Commons most directly operates — commercial open-source software — was the site of a series of attempted enclosures by well-resourced firms. MongoDB relicensed to the SSPL in October 2018. Elastic relicensed to SSPL/Elastic License v2 in January 2021. HashiCorp relicensed Terraform to the Business Source License in August 2023. Redis relicensed to the RSAL/SSPL in March 2024. Each of these moves was a textbook exercise of the architectural philosophy Karp and Zamiska’s book defends: a concentrated firm, with a strong engineering culture, attempting to capture rents from infrastructure it had helped build.

Each enclosure attempt produced a community fork. AWS forked Elasticsearch into OpenSearch in April 2021, with the project subsequently moving to the Linux Foundation. The Terraform community forked into OpenTofu in August 2023, reaching general availability under Linux Foundation governance in January 2024. OpenBao followed for HashiCorp’s Vault. The Redis community forked into Valkey thirty days after Redis’s relicensing. Within a year of the Redis fork, 83% of large enterprise Redis users were on or testing Valkey instead.

Then came the reversals. Elastic re-adopted the AGPL in August 2024. Redis re-adopted the AGPL in May 2025. HashiCorp was acquired by IBM for \$6.4 billion in February 2025, after the Terraform community had already moved on. A November 2024 academic study by Brockmeier and colleagues at CHAOSS and the OpenForum Academy found no evidence that any of the relicensings improved revenue trajectories for the firms that attempted them.

The lesson is direct. The polycentric commons-governance model — neutral hosting through the Linux Foundation, legal instruments like the AGPL designed specifically to prevent enclosure, fast-coalescing community forks, enterprise procurement pressure — defeated concentrated firm power in real time, in the exact infrastructure layer Cube Commons operates in. The firms that tried to enclose the commons failed at it. They did not even achieve the financial outcome they were trying to achieve.

There is no comparable empirical record on the side of state-aligned-firm fusion as a producer of durable civic infrastructure. The book’s central case study — Palantir itself — is an empirical record of a different kind: a firm whose primary contracts are with intelligence agencies, immigration enforcement, the Israeli Ministry of Defense, the German government (where the Federal Constitutional Court has partially banned its product), and the National Health Service of the United Kingdom (where its £330

million Federated Data Platform contract was opposed by Foxglove, medConfidential, Big Brother Watch, the Doctors' Association UK, the Good Law Project, and Amnesty UK). Whether the Palantir model produces durable civic infrastructure is, at minimum, contested. Whether it produces an infrastructure compatible with the institutional sovereignty of the institutions it serves is, after a decade of evidence, not contested. It does not.

The Karp-Zamiska book offers no live empirical record demonstrating that the architectural model it prescribes produces better civic infrastructure than the polycentric commons model. It offers instead a series of historical analogies (the Manhattan Project, SAGE, ARPANET) and a moral argument about civic obligation. The historical analogies are, as Part III argued, drawn from the wrong third of the historical record. The moral argument is the topic of Part VI.

For now, the empirical observation: in the layer where the question is being settled in real time, the commons-governance model is winning. The book's prescription would relocate civic infrastructure from the model that is winning to the model that is, where it is being tested, losing.

## **V. China**

The book's most rhetorically powerful move is its framing of the China question. Karp and Zamiska argue that the United States is in a civilizational competition with China, that this competition is decisive for the survival of democratic norms, that artificial intelligence is the decisive technology of this competition, and that the engineering elite has a moral obligation to build the infrastructure that allows the West to win.

The framing is rhetorically powerful because it taps into a real concern. The Chinese government has, during the same decade in which it has built increasingly sophisticated digital infrastructure, deepened practices that the United Nations human rights system considers possible crimes against humanity. The August 2022 OHCHR Assessment on Xinjiang is a UN document. The December 2025 sentencing of Jimmy Lai to twenty years in Hong Kong is a fact. The 2023 fourteen-year sentence of legal scholar Xu Zhiyong is a fact. The cumulative pattern of restriction on civil society, independent media, religious practice, and academic freedom in contemporary China is well-documented and serious.

The framing is rhetorically powerful, and it is also wrong. It is wrong in three specific ways that, together, produce a policy prescription that would harm the institutions Cube Commons exists to serve.

First, the framing treats "China" as monolithic. Contemporary Chinese intellectual life is not monolithic. The work of Wang Hui at Tsinghua, of Xiang Biao at Max Planck, of Yuen Yuen Ang at Johns Hopkins, of Cui Zhiyuan, of Yao Yang, of the imprisoned and silenced liberal critics like Xu Zhangrun and Xu Zhiyong, is not reducible to

state doctrine. The Chinese open-source ecosystem includes both state-stewarded projects under OpenAtom and genuinely commons-governed projects in the Apache Software Foundation, the Cloud Native Computing Foundation, and RISC-V International. Huawei is among the top contributors to the Linux kernel. The same firms that the United States now treats as security threats are simultaneously deeply embedded in the global open-source commons that the world depends on. Treating “China” as a monolithic adversary produces a policy that erases these distinctions, and the institutions that are most distinctive — the genuinely-commons projects, the dissident intellectuals, the cross-border collaborations — are the ones the policy most damages.

Second, the framing treats “the West” as coherent. It is not. The European Union has been actively building a digital sovereignty framework — through the GDPR, the Data Governance Act, the Data Act, the Digital Markets Act, the Digital Services Act, the AI Act, Gaia-X, the European Data Spaces — that is in substantial tension with the American hyperscaler model the Karp-Zamiska book defends. The Schrems II ruling in 2020 invalidated the Privacy Shield because of American surveillance practices, and the post-Schrems-II Data Privacy Framework remains under legal challenge. From the European perspective, “the West” looks less like a coherent civilizational project and more like an asymmetric arrangement in which the United States enjoys extraterritorial reach into European data while not extending comparable protections to European citizens. The Karp-Zamiska framing assumes a coherence that, on the actual data protection record, does not exist.

Third, the framing treats Global South countries as either with-the-West or against-it. This is, even more clearly, not the case. Indonesia, Brazil, India, South Africa, Saudi Arabia, the United Arab Emirates, Mexico, Turkey, Vietnam, and many other states have explicitly articulated strategies of *strategic autonomy* — the maintenance of working relationships with multiple poles without subordinating to any. The Non-Aligned Movement’s 2024 Kampala summit was the largest in decades. From the Global South perspective, the framing of a binary civilizational competition between the United States and China looks like exactly the kind of bloc-imposing pressure that strategic autonomy is designed to resist. The actual behavior of most of the world’s governments, in the actual choices they are making about digital infrastructure, is multipolar hedging — and the book’s framing makes this hedging look like betrayal rather than rational sovereignty preservation.

The combination of these three errors produces a policy prescription that is, in its actual operation, anti-democratic. It treats the natural diversity of democratic civic life — within China, within Europe, within the Global South — as either invisible or as a problem to be managed. It treats institutional sovereignty as something only nation-states can have, and only certain nation-states are entitled to. The result is that the prescription, if adopted, would consolidate power in a small number of American firms, justified by an emergency narrative that has been declared permanent. This is

not a defense of democratic norms. It is the use of democratic norms as warrant for technocratic concentration.

There is a different way to think about the China question, and Cube Commons proposes it. The question is not whether the United States can defeat China by building the largest concentrated AI capability the fastest. The question is whether civic institutions everywhere — in Boston and São Paulo and Lagos and Berlin and yes, in Chongqing and Wuhan if the Chinese system ever permits it — can build the infrastructure they need to govern themselves under their own constitutional traditions, their own legal frameworks, their own democratic rituals. The answer to that question is institutional sovereignty, not bloc concentration. And institutional sovereignty is structurally incompatible with the Karp-Zamiska prescription, regardless of which bloc one enlists in.

The genuinely concerning practices of the Chinese state — the Xinjiang practices, the Hong Kong NSL, the silencing of liberal critics — are not arguments for American AI concentration. They are arguments for international human rights enforcement, for cross-border solidarity with Chinese civil society, for sanctions where appropriate, for asylum where necessary, for the long patient work of supporting the Chinese liberal intellectual tradition that the state has tried to silence. They are not, on any honest reading, arguments for Palantir.

## **VI. The moral question**

The shortest version of the Karp-Zamiska argument is moral, not architectural. The technological elite, they argue, has an obligation to its host society. That obligation is not optional. It is constitutive of the elite's standing as a class. The frivolousness of contemporary Silicon Valley is, in this telling, a moral failure — a betrayal of the civic compact that produced the conditions for the technological class to exist in the first place.

The moral argument deserves a serious engagement, because it is the part of the book that most rewards close reading. The argument is at its sharpest in Karp's discussion of the Project Maven episode and at its deepest in Zamiska's framing of "soft belief" as the precondition for hard power. The point Zamiska wants to make — and it is a real point — is that civilizations rest on shared metaphysical commitments, and that the hollowing of those commitments produces, in time, the inability to defend the institutions that depend on them. A society that no longer believes in anything in particular cannot organize collective action of the kind that strategic competence requires.

The diagnosis is partly correct. There is a real way in which contemporary American culture has hollowed its metaphysical commitments, and there is a real way in which the technological class has been particularly affected by the hollowing. The cultural

symptoms — the proliferation of brand-positioning where conviction used to be, the substitution of identity performance for substantive belief, the inability to articulate what a society is for beyond the maximization of individual autonomy — are real. To the extent the book names these symptoms, it is doing something useful.

The diagnosis becomes incorrect at the point where it identifies the remedy. The remedy Karp and Zamiska propose is essentially: re-attach the technological class to the national-security state. The shared metaphysics they want to recover is the metaphysics of civilizational competition, of Western survival as a project, of engineering competence as a vocation in service of national power. Adopt this metaphysics, the argument goes, and the hollowness gets refilled.

The proposed remedy is a category error. The hollowness that Karp and Zamiska diagnose is real, but it is not a hollowness that civilizational-competition metaphysics can fill. It is a hollowness in the relationship between individuals and the intermediate institutions that make democratic life thick — families, neighborhoods, parishes, schools, voluntary associations, civic projects, mutual aid networks. The thinning of these institutions is what produces the metaphysical thinness that Zamiska is, partially, observing. The remedy, on any serious account of why intermediate institutions matter, is to rebuild them — not to short-circuit them by attaching individuals directly to the national-security state.

This is the substantive disagreement. The Karp-Zamiska argument treats civic obligation as a relationship between individuals (engineers) and the state (specifically, its national-security apparatus). The civic-republican tradition, the Catholic social tradition, the communitarian tradition, the polycentric-governance tradition, and the commons tradition all treat civic obligation as a relationship between individuals and the *intermediate institutions* that mediate between persons and the state. *Rerum Novarum* (1891) articulates this. *Quadragesimo Anno* (1931) articulates this with particular relevance, written explicitly against the corporatist fusion of state and economy that was being prescribed in Mussolini's Italy. Tocqueville articulates this. Robert Nisbet articulates this in *The Quest for Community* (1953), in language that anticipates the contemporary Silicon Valley pathology with uncomfortable precision: when intermediate institutions weaken, individuals become atomized, and atomized individuals invite concentrated power to fill the vacuum the lost institutions used to fill. The Karp-Zamiska remedy, in this older and richer framework, is exactly the failure mode Nisbet warned about. Individuals who have lost their intermediate institutions reach for direct attachment to the state. The state, accommodating, offers itself as the new locus of civic meaning. The result is the closing of the institutional space in which democratic life actually happens.

There is a further problem with the moral argument. Karp and Zamiska write as if the technological class is choosing between civic obligation and frivolousness. The actual choice is between several different forms of civic obligation. An engineer who works

on Linux is exercising civic obligation. An engineer who works on a municipal civic-tech project in Boston or São Paulo is exercising civic obligation. An engineer who works on a hospital information system that respects patient autonomy is exercising civic obligation. An engineer who works on the Wikipedia infrastructure or the IETF protocols or the World Wide Web Consortium standards is exercising civic obligation. The book's framing is that civic obligation means working at Palantir or its analog. This is, at minimum, an unusually narrow construction of the term. At worst, it is a category replacement: the substitution of "service to the national-security state" for "service to the polity" with the hope that the reader will not notice the move.

The moral argument, properly engaged, becomes its own refutation of the book's prescription. If civic obligation is real, and if it requires the rebuilding of intermediate institutions, then the architectural prescription that follows is the one Cube Commons articulates: build the technological substrate that allows intermediate institutions to govern themselves. Do not build the substrate that allows the state-aligned firm to govern them. The moral case for institutional sovereignty is at least as strong as, and probably stronger than, the moral case for state-firm fusion. Karp and Zamiska's argument, on its own terms, points away from its own conclusion.

## **VII. The reconstruction**

What would *The Technological Republic* look like if its diagnosis were preserved and its prescription reconstructed? This is a question worth engaging seriously, because the answer is not "the book minus its conclusions." It is a different book, with a different architecture, that retains the parts of the original that are sound and replaces the parts that are not.

A reconstructed version would keep the diagnosis. Silicon Valley has lost something. The technological class has detached from its host society. Project Maven was a bellwether. The substitution of moralism for substance is a real cultural failure. The hollowing of metaphysical commitment in the technological class is observable. These are real things, and a serious book about them is welcome.

A reconstructed version would replace the architectural prescription. Instead of "the technological class should re-attach to the national-security state," the prescription would be "the technological class should re-attach to the intermediate institutions that constitute democratic civic life." The implications follow. Build infrastructure that lets hospitals govern their own data. Build infrastructure that lets municipal governments operate their own services. Build infrastructure that lets schools, religious institutions, libraries, and voluntary associations coordinate without ceding sovereignty to a hyperscaler or a national champion. Build the protocols, the legal architecture, the open-source substrate, and the institutional governance practices that make this possible at scale.

A reconstructed version would replace the historical narrative. Instead of treating Vannevar Bush, the Manhattan Project, and DARPA as the foundational lineage, treat the actual broader lineage — Wiener, Licklider, Engelbart, Nelson, Beer, Joy, Knuth, Stallman, Berners-Lee — as the lineage to be recovered. This lineage produced most of the actually-functioning infrastructure of contemporary computing. Its civilizational seriousness is real. Its political commitments are diverse but, where they are explicit, they generally point toward distributed augmentation, institutional autonomy, and commons preservation rather than toward state-defense fusion.

A reconstructed version would replace the China framing. Instead of civilizational competition with a monolithic adversary, treat the international scene as multipolar with multiple legitimate poles — the United States, China, the European Union, the Global South strategic-autonomy bloc — each containing internal pluralism, each with civic institutions that deserve infrastructure that lets them govern themselves. The competition with the actual repressive practices of the Chinese state happens through human rights enforcement, asylum, support for Chinese liberal intellectuals, sanctions where appropriate, and patient solidarity with Chinese civil society. It does not happen through American AI concentration.

A reconstructed version would replace the moral argument. Instead of “civic obligation requires service to the national-security state,” articulate civic obligation as service to the polycentric civic landscape — which includes, among many other institutions, the actually-existing democratic state, the intermediate institutions of civil society, the international institutions of the multilateral order, and the commons-governance institutions that have produced most of contemporary infrastructure.

A reconstructed version would, in other words, be a different book. It would share the diagnostic frame and the moral seriousness of the Karp-Zamiska original. It would articulate a different architecture and a different politics. It would be, structurally, the book Cube Commons is now writing across multiple documents — the position paper *The Other Answer*, the multipolar position paper *The Multipolar Commons*, the primers, the research foundations, and this long-form response.

There is something genuinely sad about *The Technological Republic* in this light. The book is the work of serious people who have correctly identified a real problem and have, by force of habit and of vested interest, prescribed exactly the remedy that will deepen the problem. Concentration of technological capability in state-aligned firms does not rebuild the connection between the technological class and democratic civic life. It substitutes one form of detachment (frivolous consumer software) for another (proprietary infrastructure for the state). The technological class, in this remedy, gets to feel important again, but it does not actually re-attach to the polity. It attaches to the state, which is a different and narrower thing.

The remedy Cube Commons is building does the harder work. It rebuilds the con-

nection between technological capability and civic life one institution at a time — a hospital here, a housing authority there, a municipal government somewhere else. It does this through the substrate that has, historically, produced the most durable civic infrastructure: open-source software under copyleft licensing, polycentric governance under neutral foundations, federated architectures with institutional data sovereignty, and protocol-based coordination rather than platform-based capture. It does this slowly. It does this without the rhetorical theater of civilizational competition. It does this by building the infrastructure that lets ten thousand intermediate institutions hold their own data, run their own computation, and govern their own coordination.

This is the alternative to *The Technological Republic*. It is not a refusal of the diagnosis. It is the diagnosis taken seriously enough to demand a different remedy.

---

## **A coda on tone**

Karp and Zamiska's book is at times rhetorically aggressive, at times rhetorically grandiose, and at times rhetorically careless in ways that have invited dismissive responses. The dismissive responses have been a mistake. The book is serious enough that its specific argument deserves a specific response, and the dismissive critics have given Karp and Zamiska the gift of seeming, by comparison, the ones engaging in good faith. The opposite is the case. The dismissive critics are, in many cases, telling the truth about the book in compressed form. The book is, in many cases, defending a position that does not survive the close engagement they declined to give it.

This response has tried to give the close engagement. Where the book is right — in its diagnosis of Silicon Valley's hollowing, in its naming of cultural pathologies, in its insistence that the technological class has obligations to its host society — the response has agreed. Where the book is wrong — in its slide from diagnosis to architecture, in its incomplete historical narrative, in its empirical record being against it, in its monolithic treatment of China and the West, in its category-replacement of state-attachment for civic obligation — the response has said so directly.

The intended reader of this response is a person who has read *The Technological Republic* and is uncertain what to do with it. The intended outcome is that this reader leaves the engagement understanding that the book's diagnosis is partly correct, that its prescription is wrong in specific recoverable ways, and that there exists a constructive alternative that takes the diagnosis seriously without adopting the prescription. The constructive alternative is the body of work this response sits inside: *The Other Answer*, *The Multipolar Commons*, the primers, the research foundations, and the actual technical work Cube Commons is doing in the Cube Data Plane, the CUBEdesk

substrate, the agent-identity protocol drafts at the IETF, the condo-association portal, the Massachusetts law API, and the rest.

The book is influential. The conversation it has shaped is now substantial. Engaging it well is part of the civic-technology work of the next several years. This response is one engagement. There will need to be others. The forthcoming position paper *The Multipolar Commons* (CC-TR-2026-005-C) will continue the work from a different angle. The technical reports on the Cube Data Plane, the CUBEdesk diagram language, and the bias translation engine each carry pieces of the alternative architecture. Future work will continue the engagement.

The point is not to win the argument with Karp and Zamiska. The point is to make sure that the architecture civic institutions actually need — distributed institutional sovereignty, polycentric commons governance, protocol-based interoperability, institutional data custody — remains a live option in the policy environment that *The Technological Republic* is trying to shape. The book is currently winning the policy attention. The work of the next few years is to make sure that, when civic institutions across the United States, Europe, the Global South, and yes, eventually within China itself, look around for the infrastructure they need, the alternative is built, documented, defensible, and ready.

---

*Cube Commons, Inc. is a Massachusetts Public Benefit Corporation building local-first, open-source multi-agent coordination infrastructure for civic institutions. This response accompanies “The Other Answer” (CC-TR-2026-005), “The Multipolar Commons” (CC-TR-2026-005-C, forthcoming), and the associated primers and research foundations. It is deposited on Zenodo under CC BY 4.0.*