

The Ideologies of Silicon Valley

A Crooked Timber Seminar

November 2023

Contents

Introduction	3
Dave Karpf - Silicon Valley is the Church of Moore's Law	4
Sherry Turkle - Silicon Valley Fairy Dust	9
Neil Malhotra - Silicon Valley Liberal-tarians	12
Louis Hyman - Silicon Valley is the Detroit of the Future	14
Tamara Kneese - Observers Observed: The Ethnographer in Silicon Valley	16
Finn Brunton - The Agar Plate of Crypto Currency	20
Henry Farrell - The Religion of the Engineers; And Hayek Its True Prophet	26
Lana Swartz - "If it's a Ponzi, get in early": The Ideology of Scam Futures	29
Maria Farrell - Silicon Valley's Worldview Is Not Just an Ideology; It's a Personality Disorder	32
Shazeda Ahmed - From Algorithmic Monoculture to Epistemic Monoculture? Understanding the Rise of AI Safety	39

Introduction

In April 2023, Johns Hopkins' Center for Economy and Society and Stanford' s Center for Advanced Study in the Behavioral Sciences held a workshop on the political ideologies of Silicon Valley. It was a great event, in large part because it brought together a somewhat disconnected community.

People had been thinking about Silicon Valley in history, in sociology, in cultural studies and political science among other disciplines. They had read across disciplines out of necessity, to keep up with the ideas - but they often hadn't had a chance to meet the people they were reading. So this gave them an opportunity to talk. And as an unanticipated by-product of the meeting, we invited people who attended the meeting, and a couple of others, to write short pieces about their understanding of the ideology or ideologies of Silicon Valley.

The seminar is available under a CC BY-NC 4.0 DEED Attribution-NonCommercial 4.0 International License.

The participants in the seminar are as follows:

- Shazeda Ahmed is a post-doctoral fellow in the Transatlantic Digital Debates at the Global Public Policy Institute at Princeton University.
- Finn Brunton is a historian of science at the University of California, Davis, and the author of *Digital Cash* (Princeton, 2019).
- Henry Farrell blogs at Crooked Timber.
- Maria Farrell blogs at Crooked Timber.
- Louis Hyman is the Maurice and Hinda Neufeld Founders Professor at Cornell University's School of Industrial & Labor Relations.
- Dave Karpf is an associate professor of media and public affairs at George Washington University.
- Tamara Kneese is a Senior Researcher and Project Director of Data & Society Research Institute's AIMLab, and a Visiting Scholar at UC Berkeley's Center for Science, Technology, Medicine & Society.
- Neil Malhotra is the Edith M. Cornell Professor of Political Economy at Stanford Graduate School of Business.
- Lana Swartz is an associate professor of Media Studies at the University of Virginia.
- Sherry Turkle is the Abby Rockefeller Mauzé Professor of the Social Studies of Science and Technology at MIT.

Dave Karpf - Silicon Valley is the Church of Moore's Law

I have come to the conclusion that the most essential element of the Silicon Valley ideology is its collective faith in technological acceleration. More than the mix of libertarianism and tech determinism that Richard Barbrook and Andy Cameron described as the “Californian Ideology” in 1995, more than the breakdown of donations between Democratic and Republican candidates, the driving ideological conviction among the denizens of Silicon Valley is that the rate of technological innovation is accelerating.

Silicon Valley is the Church of Moore's Law.

Tech acceleration is at the center of the entire Silicon Valley worldview. This dates back at least to 1993, when WIRED magazine declared in its opening pages that the digital revolution had put us on the cusp of “social changes so profound their only parallel is probably the discovery of fire.” From this central assumption, it follows that:

- The entrepreneurs, engineers, and investors who travel in Silicon Valley circles (regardless of their actual zip code) are the driving force of history. The decisions that will reverberate through the decades are made on Sand Hill Road, not Capitol Hill or Brussels.
- Existing institutions – businesses, media, government, etc – will all soon be *disrupted*. The old rules no longer apply. New rules are being created.
- You judge the success of a business based on its *growth*, not its profit. The future is in the process of being created. The companies that monopolize an emerging sector are the companies that will define that future.
- And, finally, the future will be abundant. Everyone will reap the benefits brought about by Silicon Valley. So it is really not worth worrying about inequality, or about the distribution of those social benefits.

I have spent the past few years reading the entire back catalog of WIRED magazine, as part of a project tracing the history of the digital future. The surprising part is how little has changed in the intervening decades. Today's tech barons speak and behave much like the libertarian tech-optimist ideologues of 30 years ago. And yet there is ample evidence that the pace of technological innovation has, if anything, *slowed down* in recent years.

Belief in the immutable nature of Moore's Law is a matter of identity in Silicon Valley circles. But it is also a political project. Faith in Moore's Law implies a distrust of government – government is too slow and cumbersome. Legislators and federal agencies cannot hope to keep pace with innovation, so we are better off demanding that they stay out of the way.

So I suppose it is unsurprising that, even as the pace of “internet time” seems to have stalled over the past decade or so, our tech titans and Silicon Valley evangelists have remained insistently committed to the idea of Moore's Law.

The digerati have been propping up the corpse of Moore's Law, *Weekend at Bernie's*-style, for years now.

Meanwhile, the only thing that has exponentially increased is wealth inequality. We now have centi-billionaires like Elon Musk, and a host of social problems that are the consequence of extreme wealth inequality. We ought to distrust the ideology that tells us this is nothing to worth worrying about.

* * * * *

I imagine most Crooked Timber readers are already familiar with Moore's Law. But, as a refresher, let me explain where it originated. Writing in 1965, Gordon Moore predicted that the number of transistors you could fit on a silicon chip would double roughly every two years, while the price of those chips would be cut in half at the same rate. This prediction turned out to be durable for decades to come. The tech booms of the past fifty years – from mainframe computing to personal computers to the Internet, World Wide Web, laptops, cell phones, smartphones, and tablets – all have Moore's Law in common. Chips kept getting faster and cheaper, allowing for successive generations of digital technologies.

In *The Long Arm of Moore's Law*, Cyrus C. Mody examines Moore's Law as a social fact. Within the developing computer industry, Moore's Law functioned as a coordination point and an industrial target ("More an engine than a camera"). Developing new computer chips is a capital-intensive process. For decades, Moore's Law held true because axs vast array of actors *made* it true. Various parts of the emerging industry (one that, for much of its history, could rely on a steady stream of public funding) worked to meet the 18-24 month development cycle that Gordon Moore had effectively set as a goal. The investment in new chip production facilities was worthwhile, because chip manufacturers had a clear sense of the timeline and price points that computer manufacturers were working towards.

Mody's work focuses on the social fact of Moore's Law *within* the computer industry, though. The ideological claim made by tech acceleration is that a similar exponential curve is underway at a civilizational level. Drawing from the work of futurists like Alvin Toffler and George Gilder, Silicon Valley in the 1990s developed a heady, brash certainty that a new day had dawned. Gilder, for instance, insisted that there would be parallels to Moore's Law in telecommunications infrastructure. (He was wrong, but he was confident. And that confidence was well-rewarded).

In his 2005 book, *What the Dormouse Said*, John Markoff described Moore's Law as "Silicon Valley's defining principle. ... [Moore's Law] dictated that nothing stays the same for more than a moment; no technology is safe from its successor; costs fall and computing power increases not at a constant rate but exponentially: If you're not running on what became known as "Internet time," you're falling behind." It is, in other words, both a technical prediction about processor capacity and a deeply-held cultural belief about the pace of the digital revolution.

I have written about this elsewhere, but the pace of digital change was really quite noteworthy back then. The internet of 1993 bore little resemblance to the internet of 1998 or 2003. When technologists insisted that Congress and federal agencies could not regulate the emerging digital media environment that they did not understand it because it was evolving too fast, they were self-serving, but they also kind of had a point.

The social evidence of Moore's Law was also much more visible as a mass cultural phenomenon back then. I recently spent some time combing through the first five years of product reviews in WIRED magazine. The magazine celebrated the latest in high-end consumer goods -- \$6,000 laptops, \$8,000 flat screen televisions, \$2,500 video cameras. If you had struck gold in the dotcom gold rush, you could afford those items today. If not, you could wait a couple years and buy a superior product for half the price. "Last year's unattainable conspicuous consumption is the next year's Christmas gift."

This is just one of the ways that the mythology surrounding Moore's Law has faded in recent years. The desktop computer I brought to campus in September 1997 could barely run Windows XP when it was released in 2001. By comparison, the iPhone 15 was just announced last week. It costs about the same as the iPhone 14. It replaces the iPhone 14's lightning charging plug with a USB-C charging plug. If "Moore's Law" was once synonymous with consumer technology getting dramatically more powerful and cheaper at a constant rate, then that version of Moore's Law is now little more than a faded memory.

Faith in Moore's Law as an inevitable social force has remained strong within Silicon Valley circles, though. The victors of that era -- the young entrepreneurs who made their first millions in the first dotcom boom -- became the investors and thought-leaders of the decades that followed. There is a self-justifying logic to today's Silicon Valley ideology: the early entrepreneurs must have been right all along; Their insights being the essential ingredient of their much-deserved success. It's survivor bias, plus the influence that comes from decades of unchecked wealth accumulation.

* * * * *

A decade ago, I was myself a believer in tech acceleration and Moore's Law as a broader sociocultural phenomenon. I came to that perspective honestly: I was a graduate student in the '00s, and the pace of digital change was so fast that the research literature could barely keep up.

I became disabused of that notion in 2018, when I had my first opportunity to read the entire WIRED magazine back catalog. It became clear then that the Internet of 2018 was really quite similar to the Internet of 2013. Internet time had, effectively, slowed down.

The Internet of 2023 has many strong similarities to the Internet of 2013. It's, for the most part, the same large companies, offering the same (enshittified) services. The intervening decade has featured plenty of grand pronouncements

that a new era has just arrived (3D printing, robotics, self-driving cars, Google Glass, cryptocurrency, the Internet of Things, wearable electronics, Big Data, cryptocurrency again, the metaverse, etc etc etc). And while it is may still be too early to declare any one of these digital futures a permanent failure, we can confidently conclude at this point that the *pace* of their arrival is no longer speeding up as their promoters insist.

And yet 2023 has been something of a banner year for the acolytes of technological acceleration. Generative AI tools like Dall-E2 and ChatGPT seem solid in a way that “Web3” and the Metaverse never quite did. *This time* may be *different*. And even if ChatGPT starts to look like a cheap parlor trick after a few weeks of use, *just think how much better the next version will be*. Perhaps the past decade was just an interlude. Maybe the future is about to arrive.

Sam Altman, OpenAI’s President, wrote a sort of manifesto two years ago. It was titled “Moore’s Law for Everything.” It is quite a revealing document. If Silicon Valley is the church of Moore’s Law, then Altman, it seems, is a Cardinal.

Altman proclaims that the coming Artificial Intelligence revolution is “unstoppable,” part and parcel of the long arc of tech acceleration that we are living through. He writes:

“On a zoomed-out time scale, technological progress follows an exponential curve. Compare how the world looked 15 years ago (no smartphones, really), 150 years ago (no combustion engine, no home electricity), 1,500 years ago (no industrial machines), and 15,000 years ago (no agriculture).”

Fifteen years is an awkward anchor point. If he had instead chosen ten years, then he would have had little to point to. And it’s worth noting that, on a more reasonable timescale, the pace of technological progress doesn’t seem to be increasing at all. Back in 1997, a pair of futurists wrote a gaudy scenario for WIRED titled “The Long Boom.” They declared that 1980-2020 would feature a cluster of history-making scientific breakthroughs that would set us all on the path to global prosperity. (the authors have since taken stock of their predictions and decided “yep! Nailed it”.)

1980-2020 encompasses the rise of the personal computer, the World Wide Web, cell phones, and smartphones. Those are certainly nontrivial breakthroughs in the lives of mass consumers. But is that actually an increase in the rate of technological innovation? 1940-1980 was the television age and the nuclear bomb and The Pill. 1900-1940 was the automobile, the radio, the earliest commercial airlines, and the telephone. 1860-1900 included mass electrification and the end of slavery. 1820-1860 is basically the start of the industrial revolution. Moore’s Law only seems special because it’s the era we’ve been living through.

Altman insists that artificial general intelligence is inevitable. He also predicts that the first trillionaire will be someone involved in the AI revolution. And he speaks as though this will be a good thing. He urges his readers to preoccupy

themselves not with how we regulate and shape his new industry, but how we prepare and remold society to adapt to the natural next-step in technological evolution.

Altman has also insisted that he welcomes regulation. OpenAI has invested heavily in educating Congress about the promise and perils of this new, world-changing, general purpose technology. But OpenAI opposes regulation by existing agencies. The company only supports regulation if it gets to select the regulator and influence the rules.

* * * * *

The most important difference between the 90s tech boom and today is the relative power of Silicon Valley. “Big Tech” in the 1990s simply wasn’t all that *big* yet. And the evidence of Moore’s Law – not as a prediction about the power and price of computer chips, but as a broader social phenomenon – seemed to be everywhere. The technologists were gaining influence, the technologies were moving to the center of daily life. It stood to reason that the future would be *different* and, hopefully, better.

We care today about the ideology of Silicon Valley both because of how digital technologies have been integrated into the social, economic, and political realms, and also because of the financial power of the tech titans themselves. Elon Musk is a centi-billionaire, and he is not the only one. In the months before FTX was revealed to be an elaborate ponzi scheme, Sam Bankman-Fried had enough wealth that he openly considered buying the island nation of Nauru. Those levels of wealth are deleterious to a functional democracy. Ignoring wealth inequality made the problem much worse.

We will need to address wealth inequality in the years to come. And we should do so knowing that Silicon Valley is ideologically unwilling to help.

Sherry Turkle - Silicon Valley Fairy Dust

Silicon Valley companies began life with the Fairy dust of 1960s dreams sprinkled on them. The revolution that 1960s activists dreamed of had failed, but the personal computer movement carried that dream onto the early personal computer industry. Hobbyist fairs, a communitarian language, and the very place of their birth encouraged this fantasy. Nevertheless, it soon became clear that, like all companies, what these companies wanted most of all, was to make money. Not to foster democracy, not to foster community and new thinking, but to make money.

Making money with digital tools in neoliberal capitalism led to four practices that constituted a baseline ideology-in-practice.

1. The scraping and selling of user data. As users became accustomed to this, accepting it as the cost of online participation, the idea of privacy changed its meaning. The idea of living in a state of continual surveillance became normalized. As Foucault taught us, with this kind of change, the idea of personhood changed as well: intimacy, privacy, and democracy are woven together in an intricate connection.
2. The normalization of lying to the public while wearing a public face of moral high-mindedness. In 2021, when Facebook wanted to start an Instagram for under-thirteen-year-olds, it hid its internal research about how teenage girls felt after beginning to use Instagram. The girls said they had more suicidal thoughts, their eating disorders got worse, and they felt worse about their bodies. But Facebook was going to proceed with its under-thirteen Instagram until a whistleblower brought all this to light. The lack of commitment to truth in Silicon Valley companies is politically crucial because they are in a unique position to routinely dispense disinformation as information.
3. Silicon Valley companies that have user-facing platforms want, most of all, to keep people at their screens. Early on, they discovered that a good formula was to make users angry and then keep them with their own kind. When they were siloed, people could be stirred up into being even angrier at those with whom they disagreed. Predictably, this formula undermined the conversational attitudes that nurture democracy — above all, attitudes of tolerant listening. Digital manipulation undermines and then destroys the very possibility of conversation.
4. Avatars have politics. Online conversations make people feel less vulnerable than the face-to-face kind. As engagement at a remove has become a social norm, it has become more acceptable to stop taking the time to meet in person, even in professions where conversations was highly valued, such as teaching, consulting, and psychotherapy. In remote classrooms and meetings, in conversations-by-text, it's easy to lose skills of listening, especially listening to people who don't share your opinions. Democracy

works best if you can talk across differences. It works best if you slow down to hear someone else's point of view. We need these skills to reclaim our communities, our democracy, and our shared common purpose. In today's political climate, we most need the political skills that screen objects erode.

Silicon Valley suggests that technology will cure social problems, but it exacerbates the social problems it claims its connectivity will cure. Facebook claims to be our cure for loneliness, but online, we became alone together, less able to find a common cause.

Fewer and fewer Americans know even one other person they could call in time of emergency. We suffer an epidemic of loneliness, even as we live immersed in our technologies of hyperconnection. This loneliness is at the heart of growing rates of depression, drug abuse, and suicide. If Americans could learn to turn toward each other in the real and to act together to save their communities, that would help us save ourselves. Silicon Valley ideology insists that virtual organizing will translate into real-world connection, but online life puts you into habits of mind that don't make this translation easy. Ultimately, this has to do with tolerating and even embracing friction. In this context, Silicon Valley proposes a meta-object, the metaverse, that threatens to put us into precisely all the wrong habits of mind.

In real life, things go awry. We need to tolerate each other's differences. Virtual reality is friction-free. The dissidents are removed from the system. People get used to that, and real life seems intimidating. Maybe that's why so many internet pioneers are tempted by going to space or the metaverse. That sense of a clean slate. In real life, there is history.

Silicon Valley ideology wants to deny history because we might object to what is being done to our own. When we are online, our lives are bought and sold in bits and pieces.

From early on, pointing out this harm was most often met with a shrug. It was the cost of having social media "for free," then of having Gmail "for free." In the early years of Facebook, one young woman told me she wasn't much concerned that Facebook was looking at her data. She said: "Who would care about me and my little life?"

"Who would care about me and my little life?"— Well, Facebook did. Social media evolved to sell our private information in ways that fractured both our intimacy and our democracy.

But even after so many people knew this, conversations about this, such as conversations about climate change, tried to *not* talk about its reality. Here is how Lana, who just graduated from college, talked about how she organizes herself to *not* think about the realities of online privacy:

On Facebook, I try to keep it light. So I don't use Facebook for conversations of real consequence. And I'm glad not to have anything

controversial on my mind because I can't think of any online place where it would be safe to have controversial conversations.

Now, in fact, Lana had no lack of controversial opinions. But we can hear her convincing herself that they are not worth expressing because her medium would be online, and there is no way to talk "safely" there. This is Foucault brought down to earth. The politics of Facebook is a politics of tutelage in forgetting. Lana is learning to be a citizen in an authoritarian regime.

Lana says she'll worry about online privacy "if something bad happens." But something bad has already happened. She has learned to self-censor. She does not see herself as someone with a voice. In this small example, we see how our narrowed sense of privacy undermines the habits of thought that nurture democracy.

The former chairman of Google once said that if you're worried about privacy, don't be a Luddite, "Just be good." In a democracy, we all need to begin with the assumption that everyone has something to hide, a zone of private action and reflection, a zone that needs to be protected no matter what your techno-enthusiasms. You need space for real dissent. A mental space and a technical space. It's a private space where people are free to "not be good."

This conversation about technology, privacy, and democracy is not Luddite, and it is never too late to remember to have it.

Neil Malhotra - Silicon Valley Liberal-tarians

To analyze the ideology of Silicon Valley, one can take two approaches. One is to start “from the bottom” and qualitatively examine the writing and influence of key intellectual figures in the community. This method will yield a host of arcane and idiosyncratic ideologies and worldviews, which may or may not be reflected in political competition and policymaking, either directly or indirectly. The second approach is to consider the current structure of political cleavages and see where Silicon Valley elites – people who play a key role in the Silicon Valley business community – fit onto this existing mapping.

Both approaches have value, but I have focused on the second methodology in my own research. I’ve done this because it is useful for understanding how the elite actors in an outside industry or interest group might influence politics in the short term. Other social scientists have taken a similar approach in the past. For example, the Democratic Party transformed was transformed from an organization with segregationist leadership to a racially liberal organization over the course of the Twentieth Century. In *Racial Realignment: The Transformation of American Liberalism, 1932-1965*, political scientist Eric Schickler explains that this transition didn’t happen because of strategic decisions made by party leaders but because of bottom-up pressure from interest groups. In particular, the incorporation of African Americans into the racially progressive Council of Industrial Organizations (in contrast to the more racially conservative American Federation of Labor) played a crucial role in both the civil rights movement and the realignment of American politics.

So how do outside elites exert influence on a party? As the economist Albert Hirschman has suggested, they can attempt to influence a party from within its coalition (*voice*) or leave the party and engage in external competition (*exit*). By examining Silicon Valley elites’ views on *existing* political debates, we can attempt to predict whether they will exhibit voice or exit.

That is the basic idea behind a February 2017 survey I conducted with David Broockman, a political scientist at U.C. Berkeley, and Greg Ferenstein, a technology journalist. We constructed a sampling frame of Silicon Valley elites based on a database from Crunchbase, which compiled a list of technology company founders based on SEC filings of Series A venture capital investments. We also conducted surveys of the mass public and partisan political donors at the same time as points of comparison.

Our findings led us to greatly rethink what we think of as “the Silicon Valley ideology,” which many pundits equate with libertarianism. In fact, our survey found that over 75% of technology founders explicitly rejected libertarian ideology. Instead, we found that they exhibit a constellation of political beliefs unique among any population we studied. We call this ideology “liberal-tarianism.” Technology elites are liberal on almost all issues^¾including taxation and redistribution^¾but extremely conservative when it comes to government regulation,

particularly of the labor market. Amazingly, their preferences toward regulation resemble Republican donors.

These findings help explain why the technology industry was a core part of the Democratic coalition and why the Obama Administration was fairly lax when it came to the regulation of the technology industry. Silicon Valley worked from *within* the Democratic coalition to move regulatory policy to the right, while supporting the party's positions on social issues, economic redistribution, and globalization. Perhaps it is not so surprising that the revolving door between the Obama Administration and Silicon Valley companies seemed to always spin people in or out.

However, things might be changing. There is increased hostility between the technology industry and the Biden Administration. And this is not just with respect to Elon Musk. The Biden Justice Department has recently filed an antitrust lawsuit against Google (on the heels of legal action against Microsoft), and the Federal Trade Commission as generally been more supportive of regulation of the industry. Prior to 2016, the technology industry was seen by Democrats as a "clean" industry (unlike fossil fuels and manufacturing) that was acceptable to engage with and take donations from. The three main fundraising stops for major Democratic politicians were usually New York (finance), Los Angeles (media), and San Francisco (technology). However, the left soon realized that companies can be associated with social externalities besides air and water pollution, such as misinformation and foreign influence in elections.

At the same time, technology elites may be becoming more conservative on issues besides regulation. I don't have survey evidence on this, but my conjecture is that the industry has taken right-wing positions on some new issues that have emerged on the political agenda. For example, many technology elites spoke out against Covid restrictions such as lockdowns and vaccine mandates. After the murder of George Floyd in 2020, some pushed back against renewed emphases on racial equity and the re-definition of merit. More broadly, many technology elites have chafed at the left's pushback against unbridled free speech. Technology elites have always seen themselves as disruptive, countercultural and iconoclastic, stretching back to the computer industry's origins in the 1960s. Now that the left dominates culture, the countercultural rebels seek to fight against it.

This all leaves the technology industry between a rock and a hard place, not fitting in well with either party. It will therefore be hard for them to follow the lead of the Council of Industrial Organizations (which shared the party's views on labor issues but not civil rights) and change the party on select issues from within. Their influence on politics must come from outside party structures. There isn't a great roadmap for that, potentially suggesting that Silicon Valley ideology may become less politically influential.

Louis Hyman - Silicon Valley is the Detroit of the Future

In the U.S., there is a city where industrial visionaries, state leaders, and financial titans all clamor to go. They want to see the future being made today. Revolutionary new ways of working are being combined with technology to create a better standard of living for the employees, cheaper products for the consumers, and unimaginable returns for the investors. Driving around, you can see the boom everywhere you look. Factories. Warehouses. Mansions. This town is going to be the center of a new world. This town is Detroit in 1920.

Today Detroit is synonymous with capital flight and urban decay, but there was a time when it was synonymous with the future. Henry Ford made a new kind of corporation that, in turn, made Detroit. Ford's methods were so influential that they became known simply as Fordism (something that Google and Facebook still haven't achieved.) A poll in the 1920s found that college students rated him the third most important person in human history after Jesus and Napoleon. Polls today might not celebrate Zuckerberg or Musk as saints in quite the same way, but there is certainly a sense that they aren't far behind. If we think Silicon Valley's apotheosis is unique, we are mistaken.

Yet Detroit fell. Detroit followed the Ford and Fordism in losing ground. It is hard to imagine that as the economic moment that buoyed Silicon Valley passes, it too will not suffer the same fate. And here, I mean literally the people, communities, and built environment circling the Bay.

Facebook, now Meta, built so many buildings in the pre-pandemic years, spanning from Palo Alto up to Redwood City, that they could be measured in miles, not square feet. As we all know, office vacancies are at unprecedented levels as employees have begun to demand to work from home. This summer, Meta, like many other Silicon Valley giants, took \$2 billion in losses as it tried to offload its inventory. While the shift to work-from-home has reduced, and will reduce demand for office space, the real question is when we shift from Facebook/Meta, or in the longer term, from the Silicon Valley economy.

As we think about Silicon Valley's rise, we must also think about its inevitable fall. In doing so, Silicon Valley seems less like a discontinuity than just another case in the history of capitalism. While its fall might be inevitable, the consequences of that fall might not be, especially if we plan ahead. I have great faith, disappointing as it might be for some, that capitalism will continue to reinvent itself, but little faith that it will do so in the same place over and over again. Cities rise and fall. And given our own short history, we have seen many quick rises and even faster falls. Places, like Detroit, but also Buffalo, St. Louis, which were at one time engines of wealth creation, fall apart. The real estate crisis of the office might be a signal of Silicon Valley's end or not, but those buildings created for a particular moment will last long after the companies are gone.

As we think about the future of Silicon Valley, we should always think about the history of Detroit, as a useful analogy, considering what will be after the heyday, and how best to create a path forward after the saints are gone, in a near future era when their sanctuaries are crumbling into desuetude and ruin.

Tamara Kneese - Observers Observed: The Ethnographer in Silicon Valley

It is undeniably powerful to hear workers' stories in their own words. Movements can emerge from the unlikeliest sources. The oral histories of ordinary workers are often seen as distinct from the memoirs of outsiders in tech, many of them women, who have written about their experiences. The latter range from Ellen Ullman's 1990s memoir from the perspective of a woman software engineer to Anna Wiener's viral essay and then monograph-length account, *Uncanny Valley*.

Elsewhere, I have written about the politics of collecting stories from the margins of Silicon Valley. I argue that the femme tech memoir, as an iteration of the personal essay genre, can be read alongside workers' inquiry as a way of finding solidarity across job descriptions and positionalities. Workers' inquiry combines research with organizing, allowing workers themselves to produce knowledge about their own circumstances and use it in their labor organizing. This seems especially vital in an industry where even short-term history is hard to access and most workers stay with specific companies or in particular roles for short stints.

I came to this project as someone who was first trained in cultural anthropology and then entered the tech workplace myself as a researcher and, to my mind, an organizer-observer. I was specifically recruited into my research position *because* of my ethnographer background and my labor organizing skills.

In this brief essay, I call attention to another form of Silicon Valley personal narrative: the stories of the ethnographers who became tech workers, those whose "soft skills" and human-computer interaction expertise are often discarded during times of economic upheaval and who sometimes treat their own workplaces as yet another field site for ethnographic investigation. What do the stories of the many generations of ethnographic researchers who joined and sometimes left the tech industry have to tell us about how Silicon Valley ideologies are taken up, embedded, and contested in workflows and products? How do the collected personal stories, or oral histories, of UX researchers interface with those of tech campus janitors and engineers? And is there something valuable that can be learned from their varied experiences about the sometimes ambivalent relationships between research, work, and collective action?

Since Xerox PARC's early, experimental practice of incorporating social scientists into the product development process, ethnographers have entered the tech workforce and contributed to R&D efforts. Today, Responsible AI and other DEI (diversity, equity, and inclusion), sustainability, and ethics teams outside of or adjacent to product teams—corners of tech where social scientists have an outsized presence—are often marginalized within corporations. They are not necessarily empowered to make decisions about which products and features are developed and when they are shipped. Tech companies have the opportunity to use these teams to rubber stamp their products and policies, enfolding

humanistic research and internal employee resource group-related organizing work in marketing campaigns.

At the same time, there is a long history of ethnographers going into the “field” of Big Tech, positioning themselves as participant observers in order to provide insights into the ideologies underlying Silicon Valley production. They tend to view themselves as perpetual outsiders while making new markets visible to corporations. Following the methodology of cultural anthropology, through the notion of “fieldwork,” or gathering data while embedding themselves in the collective lives and social rituals of those they study, through participant observation and interviews, many social scientists within tech don’t see themselves as full members of the enterprise. They always exist slightly outside of it and their observations of their everyday work experiences constitute a form of fieldnote documentation.

Xerox’s Palo Alto Research Center (PARC) launched in 1976. The STS scholar Lucy Suchman writes about her role at Xerox PARC in her book chapter “Consuming Anthropology,” tying her personal experience to the larger history of anthropology. The discipline of anthropology made itself legitimate and legible by partnering with colonial governments and other institutions, which valued ethnographic information on the people in their colonies. Anthropology has a rich history of observers observing themselves and others in the field, a form of self-reflexivity or auto-ethnography which at times teeters on the edge of memoir and that occasionally attempts to do penance for anthropology’s undeniably murky past.

Many anthropologists in tech have similarly examined their own roles within corporate hierarchies. Suchman, along with several other anthropology PhD students at UC Berkeley, were placed in Xerox sales offices to track and analyze the company’s workflows and social relations in their sales and customer service work. Suchman writes that she is attempting to think “about the embedding of anthropological research within corporate enterprises in relation to the turn to markets as a research object in the social sciences.” By entering tech corporations, ethnographers created new field sites for themselves.

Suchman goes on to describe this method as being connected to, but a departure from, a pro-labor Scandinavian approach:

In particular, we drew the model for our own practice from colleagues in Denmark, Norway and Sweden; academic computer scientists collaborating with Scandinavian trade unions to develop union-sponsored demonstration systems informed by values of quality of working life and workplace democracy. In our representations of the value of participatory design to the corporation, however, political values were minimised in favour of potentially superior design outcomes, producing information systems better suited to working practices.

The North American justification for ethnographic research, typically internally

referred to as UX, focuses on optimizing work through technology. Anthropologists anticipate the future needs of users, which can be used to guide company-wide strategy and R&D.

Anthropologists have long been employed at places like Microsoft Research, IBM, Intel, and the Institute for the Future. They are often explicitly charged with making the speculative future of work more visible and tangible. Lonny Brooks and Geof Bowker were resident anthropologists at IFTF, the consulting firm, as corporations attempted to navigate the growth of digital media in the 1990s. Brooks conducted ethnographic research on IFTF clients and internal research practices;

A symmetry of the research site was that the observer was himself being observed. If IFTF works (as we shall see) to develop a playground to contain the Other, the researcher represented for the company just the sort of Other they were seeking to incorporate – a marginalized graduate student from a relatively radical department.

For IFTF, Brooks was attractive precisely because of the outsider perspective, the role of the Other and the ethnographer, that he could play.

In recent years, there have been more examples of UX researchers looking back on their work in the field. John Sherry worked at Microsoft and then at Intel starting in 1997, when he was hired as their first anthropologist. The following year Genevieve Bell, with a PhD in anthropology from Stanford, was recruited into a research role at Intel, where she started the company's first User Experience group in 2005. In a 2019 interview, Sherry argues that anthropologists should not be limited to UX teams. Rather, they provide a point of view and change the trajectory of technologies that are developed. He says:

New technologies are changing the world along so many dimensions, not just in terms of personal experiences but affecting our economies, our political systems, our whole social fabric. . . We need smart anthropologists thinking about impacts and helping us identify more equitable and humane futures. Luckily, a lot of companies are starting to wake up to this as well."

But as happens with the ebbs and flows of industry, it is unclear if such nuanced roles will survive the growth of generative AI alongside mass layoffs.

As Mel Gregg says in her book chapter about her experience as a gender studies expert at Intel, there is often a disconnect between how the outsider-researcher is imagined by the company and their individual experiences. As a rare "technical female" in the enterprise, Gregg learned to carve out time to do the kinds of research, reflection, and writing that might actually produce a more humanistic form of technology. She didn't have an office of her own and she booked meetings with herself to escape the forced innovation and collaboration of the open office plan.

Gregg, along with many others like her, myself included, always felt like they

were slightly outside of tech because of their gender, sexuality, race, and/or class position. “Technical Female” is part of the same milieu of discussion as works by Ellen Ullman, Anna Wiener, Wendy Liu, and Joanne McNeil. But there is something especially rich about the stories of those who came to tech as observers, whose research was intended to deliver value to the corporation. What do their stories tell us about the changing values and labor conditions of the tech enterprise?

The participant observer position is not unique to anthropologists in tech. Many outsiders enter tech and never quite feel a part of things. The UX designer and organizer Yindi Pei’s *The Valley of Misfit Tech Workers* zine is emblematic of the organizing potential and camaraderie contained in the shared experience of being a misfit. On a quite different note, there are many hand-wringing think pieces about the death of UX research because of its failure to deliver business value. During a time of increased precarity, it is clear that many corporations see ethnographic insights as expendable. Having left the tech industry, I feel compelled to comb through the stories of generations of fellow ethnographer-outsiders in tech in order to find commonalities and offer a way forward.

Finn Brunton - The Agar Plate of Crypto Currency

One useful way to think about Silicon Valley ideologies is through looking at cryptocurrency. Payment systems like Paypal were as much a part of the Silicon Valley story as platforms like Facebook. They have always been entangled with political aspirations. According to Peter Thiel, Neal Stephenson's *Cryptonomicon* was required reading for Paypal's leaders. They wanted to displace the US dollar with their own private currency, building a new global economy. Then came Bitcoin, and Ethereum, and a whole host of privately issued cryptocurrencies that entangled various political aspirations with the desire to make a whole heap of money on the foundations of blockchain.

Cryptocurrencies, as a category of technology, are extraordinarily well-suited to hosting ideologies. Extraordinary relative to what? Consider email. Email can certainly be a subject of ideological concern – I wrote a whole book about how it became a space for negotiations about speech, markets, and political and private power. But at the level of developing the Simple Mail Transfer Protocol (SMTP) to copy text files from one box to another, at the level of fiddling with address and transport schemes for messages, it partakes of the implicit ideological background radiation of a network technology developed in modern society, but without attempting to deliberately embody a coherent political-economic system of theories, goals, and policies. You might call email an ordinarily ideological technology. Crypto's very foundations are a set of explicitly ideological projects whose ongoing thrash and debate inspires forks (when one protocol splits into two or more), fights, and further development at the level of protocol and platform.

Put it this way: can you imagine someone identifying politically, in seriousness, as an “email maximalist”? Or someone saying that email is the postnational “flag of technology”? Back in the early years of digital cash, Tim May – an actual no-kidding ideologue who wrote manifestos and coined -isms for his mutant strain of militant libertarian political economy – described this problem. For purposes of developing the technology, “it's fairly clear what ‘encrypting’ or ‘remailing’ means”: making email secure and private has a stable and shared set of concepts and questions. But “the situation becomes much murkier for things like digital money . . . just what is a ‘digital bank’?”

The question is this: what is money, and what has the authority to mint it? That eminently political query lies behind the bewildering myriad of terms used to describe different aspects of crypto- and traditional currencies: centralized, decentralized, distributed; inflationary, deflationary, stable, pegged, speculative; currency, commodity, ticket, asset; private, transparent, pseudonymous, anonymous; reversible, irrefutable, airdropped, contractual; institutional settlement, microtransaction, collateral, on-chain, off-chain; proof of work, of burn, of space, of stake – and more, some presupposing each other, others mutually exclusive.

What constitutes identity, ownership, and the rules of exchange? National or extranational, and if so what's the interface – is there one? – to territorial, central bank currencies? What is mandatory, easy, encouraged, difficult, or forbidden? And, above all, who has the capacity to make these decisions and on what grounds? To develop the technology – to actually write and commit code – you have to start with answers to ideological questions.

This fact has made crypto into an ideological agar plate, a growth medium for political-economic ideas and theories. Some are preexisting, some are widely accepted, some are novel to crypto and its sister technologies, and some are deeply weird. What follows is a tentative, brief, and somewhat snarky typology aimed at identifying and grouping these ideological formations and subformations, and at reflecting on some of the larger ideological changes they reflect.

Hard Money Libertarians

There are lots and lots of libertarians in all their ideological and theoretical variety in crypto, but the main current of old-school libertarianism is the *hard money* crowd. They believe in money and other financial tools secured by something other than the authority of a central bank to issue and redeem it. Usually this means precious metals, sometimes land, or, in a more avant-garde mode, energy, human labor, or computational work. This conviction is the most salient part of a larger complex of beliefs, from disputing the right of the state to mint money, to debates about humans as rational market actors. Subfactions include:

Digital Gold Currency Entrepreneurs

DGCs were the original internet payment platforms, created at a confluence of market opportunity and ideology – a still deeper history that threads back through gold- and silverbug commodity money activists in American history for more than a century. DGC proponents and architects are still around, and their ideas and examples continue to play a role in modern crypto, albeit diminishing and rather unfashionable.

Heterodox + Austrian Economics Nerds

People who read doctoral dissertations from the 1930s and wrestle with brontosaurus tomes of econo-metaphysics to explain the need for free and private digital banks that can issue their own currencies. *End-the-Fed Activists*. Entirely devoted to the destruction of the creature from Jekyll Island, for which crypto provides a useful ally. The End the Fed crew stand like the Archdeacon Frolo in Hugo's *Notre-Dame de Paris*, pointing from the book on his desk to the

cathedral outside and explaining, “This will kill that.” They would rather not talk about the whole Eustace Mullins thing.

Accelerationists

Those for whom crypto in form or another is a means to the realization of a potential future human condition (as opposed to the return of an idealized political-economic past with the hard money crew).

Extropians

Early transhumanists, the Extropians combined a philosophy of human potential with a sci-fi sensibility and a reading of libertarianism focused on emergent order and the potential of free markets to drive technological innovation. This heady cocktail is a cosmic ideology in which new digital currencies are fuel for a cascade of breakthroughs heralding the end of the human condition; it led to much foundational work in cryptocurrencies. Their philosophy, aesthetic, and membership has had a massively outsized and under-recognized cultural influence, not least on the crypto space.

Singularity-Longtermist-Effective Altruism Ethical Wealth Capturers

Slightly tangential to crypto as such but exerting a heavy pull: the utilitarian-utopian ideology of capturing wealth, primarily through production of and speculation on crypto. This captured wealth is supposed to be deployed in worthy endeavors that tend to tip over from near-term priorities like mosquito nets and carbon capture into more futuristic and theoretically groovy projects like friendly artificial general intelligence, space exploration, and existential risk scenario management. Crypto is a means to underwrite this philosophical project, but both the means and the ends share a similar attitude of seeing past the limits of contemporary irrational ideological institutions to a rational, data-driven future order.

NRX + Dark Enlightenment Fascists / Monarchists

The ultimate in tech CEO worship identifies crypto as a vanguard component of an emergent post-national tech-finance elite who deserve to be put in charge. Prolix and verbose in style, characterized by multi-hundred post threads, novella-length documents, and droning podcasts. There are various degrees of overtness about the “race science,” antidemocratic, and eugenicist dimensions of the program. Crypto here acts as an ideological store of value, a membership card or flag of allegiance, and sometimes an outright technical version of a

philosophical claim (as with Urbit, which began as a neoreactionary digital feudalist authority-and-property scheme for a computer network).

Smart Contractors + DAOists + NFT Metaversists

Most notably Ethereum. I'm putting this in the accelerationist bucket because many of the technologies and implementations involved reflect a diverse ideological landscape of ways that society and property could be organized, with the technology as a means of getting there. Crypto as voting rights, crypto as a component of consensus or constitutional processes, crypto as a tool for the ownership of intangible property; providing access to savings, credit, and financial tools for the unbanked; self-sovereign identity, polycentricity, 21st century versions of the Hanseatic League: attempts to move towards political processes and institutions that do not yet exist, united by the family of crypto-blockchain tools and imaginaries they use. This subcategory could easily be broken out into several distinct sub-subcategories.

Cryptoanarchists

The third leg of the historical stool of cryptocurrency, along with Extropians and hard-money libertarianism. The cryptoanarchist current identifies the protection of transactional data as the most important part of the technology. Their reasons reflect some interesting ideological distinctions.

Cypherpunks

Another small but remarkably influential movement, identifying the components of cryptography that could be used to protect computer users from surveillance. Their shared belief was that the mathematical breakthroughs that constituted modern cryptography should not be classified military secrets, but their reasons were varied – ranging from protecting classic civil liberties and personal privacy from the power of a computerized state, to the promise of the collapse of governments into “crypto-anarchy” due to the proliferation of untraceable, untaxable, unseizable digital cash. These distinctions play out in their current day heirs.

Agorists

Ideologically, agorism is the esoteric postpunk to libertarianism's pop-punk: intellectual, challenging to get into, more European than American, and with better graphic design. Founded by a couple of ultra-obscur theorists and science fiction writers, agorism is premised on the creation of “counter-economics,” anonymous covert grey and black markets whose operation outside the state diminish its capacity to govern and produce non-state spaces for people to

conduct more and more of their lives. Crypto was well suited to adopt into this ideological project, and the Silk Road contraband marketplace was a deliberately agorist project that ran, of course, on Bitcoin and provided its first substantive use case. Numerous copycat marketplaces (several of them honeypots built by law enforcement) may not share the same sense of mission, but they share the operational need: the currency should be anonymous.

Privacy Oriented Civil Rights Defenders

People who recognize the threat of transactional surveillance and seek to avert or mitigate it with currency, among other tools. Privacy is the main benefit of these technologies, and others (portability, ease of transaction) are secondary. David Chaum epitomizes this current. They are part of the deeper history of the analysis of the dangers of governmental overreach; when Paul Armer was warning about the threats of electronic funds transfer surveillance in the 1970s, his reading list included Nixon administration memos on domestic intelligence gathering and Lucy Davidowicz's *The War Against the Jews*.

Basic Beige Neoliberalism

Operating in the context of a default finance-capitalist political-economic model of the world. Powerful, entrenched, ubiquitous, boring – the index fund of ideology. The interesting parts relevant to crypto and Silicon Valley clump into three clusters of mutant activity that express, in my opinion, three failures on the part of beige neoliberal ideology to address its crises and internal contradictions. I don't think any of these qualify as ideologies on their own, exactly. Instead, they act as decadent, cultish, performative ideological formations.

FOMO VCs

Suffering from 0% interest rate psychosis, investors poured money into venture funds, and venture fund bets went beyond wild into some other dimension. This produced a fascinating evolution of the greater fool into a more comprehensive *fear of missing out*. Crypto, blockchain, and NFTs all got a chance to be the subject of frenzy at the institutional and retail levels – not because people held ideological convictions about them, quite the contrary: investors followed one another like an ant mill. This changed the shape of the crypto space by emphasizing stuff that was fundable rather than functional, feasible, or useful.

Grindset Lifestyle Entrepreneurism

Hustle culture bros on TikTok leaning on what is clearly a rented Lamborghini and telling you about how they get up at 3AM and read five books a day, then

segueing to their crypto trading course or pitching you on investing in some variety of crypto you've never heard of. A performance of the ideal subject of beige neolib ideology, for which crypto is the ideal product: pure branding and confidence and hustle, free of the burden of actually making or producing anything.

YOLOism

The decadent stage of investment nihilism: buying into, speculating on, and memeing about cryptos for the sake of drama and clout. The wild swing, the ruinous bet, the diamond hand bravado of “Lambos or food stamps.” Think ugly, low-effort NFTs and cringe meme cryptos like Dogecoin, asinine jokes that lurch into actual money. This is a retail investor's financial capitalist ideology when there is no plausible future on which to base an adult portfolio strategy – you only live once. This overlaps with the meme stocks scene, parts of which have now entered a completely conspiratorial unreality.

Henry Farrell - The Religion of the Engineers; And Hayek Its True Prophet

Marc Andreessen's recent "tech optimist manifesto" is one of the most significant statements of Silicon Valley ideology. As I've written elsewhere, it's actually less a political manifesto than an apostolic credo for the Religion of Progress. The words "we believe" appear no less than 113 times in the text, not counting synonyms.

The core precept of this secular religion is faith in technology. From Andreessen's opening section: "We believe growth is progress . . . the only perpetual source of growth is technology . . . this is why we are not still living in mud huts . . . this is why our descendents [sic] will live in the stars."

Andreessen invokes the right wing economist, Friedrich von Hayek, as one of the "patron saints" of this dogma. That might seem like a surprising assertion. Hayek was ferociously critical of what he described as the "religion of the engineers" – the efforts of Saint-Simon's followers to create a quasi-messianic faith applying engineering insights to society. Their fervid belief in the inevitable benefits of progress purportedly justified the efforts of an elite to remake society along better and more rational lines.

Hayek quotes an early Saint-Simonian journal as describing a program to "develop and expand the principles of a philosophy of human nature based on the recognition that the destiny of our race is to exploit and modify external nature to its greatest advantage." Compare to Andreessen: "We believe in nature, but we also believe in overcoming nature. We are not primitives, cowering in fear of the lightning bolt. We are the apex predator; the lightning works for us."

The obvious difference is that the earlier religion of the engineers glorified the state, while the new one glorifies markets (that's why Hayek is one of its patron saints). But the similarities are *at least* as important. Both the old time religion and the new one invoke grand visions to wave away the mess, disagreements and complexities of the present. They depict those who oppose the actions of a tiny self-elected elite as champions of ignorance and enemies of progress. If we *only just* let the engineers run things, we could be sure that our descendants will have the universe for their inheritance.

* * * * *

I've been trying to work out my thoughts about the relationship between the old and the new religions of the engineers for years. Hayek plays an interesting and complicated role, as erstwhile CT contributor, Corey Robin has pointed out. His suggestion that rich elites will and should play a crucial role in guiding the progress of an apparently decentralized and pluralistic system helps justify the world-shaping ambitions of founders. So too, does Schumpeter's theory of the entrepreneur and of the general benefits of monopoly. But my sense of what

is going on was really crystallized by Daron Acemoglu's and Simon Johnson's recent book, *Power and Progress*.

This book gets Andreessen's shtick down cold, in a book that was published well before the manifesto (Andreessen is expressing the collective wisdom of those around him as much as his own thoughts). Acemoglu and Johnson describe a standard optimistic mythology, according to which we are "heading relentlessly toward a better world, thanks to unprecedented advances in technology." Whatever problems we experience are the birth pangs of a better world that is just around the corner. In their description, "[p]eople understand that not everything promised by Bill Gates, Elon Musk, or even Steve Jobs will likely come to pass. But, as a world, we have become infused by their techno-optimism. Everyone everywhere should innovate as much as they can, figure out what works, and iron out the rough edges later."

More specifically, the book explains exactly how claims about the awesome freedoms of the markets are interwoven with practical restrictions on people's liberties. It emphasizes the importance of Jeremy Bentham's ideas about the general benefits of surveillance for economy and politics: "before the panopticon was a prison, it was a factory." These ideas paved the way for factories that turned workers into "mere cogs" and the later notions of Frederick Taylor and others who looked to use new technologies of surveillance to squeeze as much productivity out of workers. The standard response is that everyone benefits from this in the long run, but Acemoglu and Johnson stress that this is hardly a given. How the benefits are distributed depends on politics, and specifically on whether those who are on the receiving end are able to organize and ally with others, to create "countervailing power" that ensures that the benefits of new technologies are evenly distributed, and to avoid technological trajectories that maximize on exploitation rather than general benefits.

These historical lessons have relevance today. I've heard it said (correctly or incorrectly) that Andreessen's tirade was largely motivated by his anger at AI skeptics. Certainly, one of his proposed articles of faith is that "We believe any deceleration of AI will cost lives. Deaths that were preventable by the AI that was prevented from existing is a form of murder." Acemoglu and Johnson point out that AI is regularly being used to replace workers or to surveil them. They stress that this is a political decision, rather than an inevitable consequence of technology. We can choose differently, and we ought to.

* * * * *

Like other religions - like Marxism too for that matter - the religion of the engineers is centered on a myth about the world to come. A lot of people talk about the influence of science fiction on Silicon Valley, and how people like Peter Thiel and the Paypal Mafia were inspired by Neal Stephenson's ideas about money. Stephenson is an important part of the story that Silicon Valley tells itself about its present - the Metaverse, Google Earth and so on. But I can't help wondering if the Culture novels of Iain M. Banks (cited for example by

Jeff Bezos and Elon Musk as core texts) are more important to the stories that Silicon Valley tells itself about the future.

Banks' future is one where humanity (I simplify here – the Culture's relationship to actual Earth-humans is complicated, and much happens in our past, Elsewhere in the Galaxy) has figured out how to produce universal abundance. Within very broad reason, the people of the Culture can have whatever they want, traveling the universe in massive starships, constructing vast Orbitals, glanding drugs, having lots of sex, changing gender at a whim (Musk may have changed his mind on that bit) and throwing wild parties, all overseen by more-or-less benign AIs. It's a very attractive future, where socialism and libertarianism blur into each other.

I can't say whether Andreessen's manifesto is directly influenced by Banks' novels, but its imagined trajectory at the least adjacent, with AI as "our alchemy, our Philosopher's Stone" and a future in which:

We believe the global population can quite easily expand to 50 billion people or more, and then far beyond that as we ultimately settle other planets. We believe that out of all of these people will come scientists, technologists, artists, and visionaries beyond our wildest dreams. We believe the ultimate mission of technology is to advance life both on Earth and in the stars.

In contrast, I am reasonably sure that Banks would have *absolutely fucking hated* the tech optimist manifesto and the project behind it. His books have plenty to say about people who promise paradise tomorrow to justify purgatory and hell today. None of it is complimentary. His books are all about the complexities and the tragedies of politics.

There isn't any room for complexity in Andreessen's vision. The politics are all stripped out. There is only a struggle between the Good who embrace technological progress, and the Enemies of Progress. The religion of the engineers is the hopium of Silicon Valley elites. It's less a complex theology than an eschatological soporific, a prosperity gospel for venture capitalists, founders and wannabes. It tells its votaries that profits and progress point in exactly the same direction, and that by doing well they will most certainly do good. It should barely need pointing out that the actual problems and promise of technology lie in the current political struggles that this vision of the future waves away.

Lana Swartz - “If it’s a Ponzi, get in early”: The Ideology of Scam Futures

When he opened the seminar that prompted these essays, Fred Turner said that Silicon Valley built more than semiconductors or search engines or smart phones or sharing platforms. Indeed, he suggested that Silicon Valley’s true product is ideology. In my notes, I wrote and underlined, “Silicon Valley creates and retails visions of the future.”

This resonated with my own research. Money—the main technology I study—is one way to do futurity, as I (and many scholars including Finn Brunton, another seminar participant) have argued. We only accept money from other people today because we think that someone will accept it from us tomorrow, and so on, into multiple tomorrows. When we invest, we are laying bets on particular visions of the future.

Retail investing, whether in crypto or meme stocks or more traditional tech industry IPOs (companies like Tesla or Coinbase are making the distinction ever blurrier), is one way that Silicon Valley quite literally retails a future. If you invest in an asset, you are both supposing and summoning a future in which that asset is worth more than you paid for. Retail investors can be active market participants in the future sold by Silicon Valley, not merely passive feminized “consumers” of that future.

Studying retail investing is one way to explore how Silicon Valley ideologies move from centers of power, such as the actual physical place called “Silicon Valley,” and diffuse to the rest of the world. Retail investing resembles Althusser’s notion of the classic Ideological State Apparatus. It is a vector of ideology, a way of mediating it. I have been told, by probably about four different interviewees in crypto, that they (or “someone they know”) became more of an ideological believer in the politics of crypto as they watched the line go up and the potential cash-out value of their investment grow. When the line goes down, they don’t abandon those beliefs. Instead they revise them, and qualify them to rationalize either selling at a loss or “hodling” on.

Some of these futures, of course, are scam futures. I have accidentally become an expert on scams. That’s what happens when you’re a researcher immersed in the arcane worlds of e-commerce and cryptocurrency: you get to know all the ways that scams are produced and prevented. You watch the processes of how the people who build and maintain economic infrastructures decide what’s a scam and what is just business as usual, buyer beware.

As I argued in my article on the 2017 crypto ICO (Initial Coin Offering) bubble (which I’ll draw from quite a bit in the rest of this piece), many of today’s Silicon Valley retail investing opportunities are a frenzied effort by a collective to bring about a future, even though the collective doesn’t necessarily believe in it.

Just as we saw with ICOs in 2017 and NFTs (Non-Fungible Tokens) more recently,

investors aren't just buying a crypto asset that they hope they hope to sell at a profit. They are also buying a vision of a future in which these crypto or Web3 technologies have profoundly disrupted some or all of society. This kind of retail investing regularly involves literal scams, notably rug pulls and pump and dumps, both of which reveal a kind of arbitrage on belief in the future. Those who are true believers are focused on the long term promised future. Those who make money are focused on the short term. As one interviewee told me, "Everyone tells everyone else to hodl, but no wants to be a bagholder."

This way of doing futurity is fundamentally characterized by ambiguity and asymmetry. Among those buying in, there is an uneven (but perhaps knowable) likelihood of benefit from the scam and an uneven (and *unknowable*) belief in the likelihood of its promised future.

It's clear that those *selling* ICOs (or NFTs or meme stocks) are more likely to get material benefits now and in a range of possible futures than those *buying* and *speculating* in them. It's less clear, and really impossible to know with any certainty, who really believes in the dream and who is trying to turn a quick buck buying low and selling high.

Those who stand to benefit and who are cynical about promises offered by today's new economic formations could be called scammers, but the scam is only possible because of the effervescence resulting from the collective enthusiasm of everyone involved. In my article, I call this a "network scam."

Were the VCs who were (maybe) duped by Theranos similarly participating in a network scam? Does the scamminess that happens when we use Robinhood on our phone following the same logic as venture capital? Can we trace a ideological flow from one site to another?

I do think it's important to note that retail investing doesn't just copy-paste ideologies from founders and VCs and technocrats into the heads of Robinhood users. We need to understand better how the ideologies, ramified via markets and gamified apps, changes as it is circulated and inhabited. As the entire project of media and cultural studies has shown us, it's important not to be content with just-so stories.

Nevertheless, as crypto impresario Dave Portnoy put it as he was launching his own cryptocurrency, "If it's a ponzi, get in early." Scammers don't usually invite their marks to participate in scams, promising that they will at least have a high relative position in the scam pecking order. In this milieu—and the one that brought riches to countless VCs through vaporware if not fraud-- traditional notions of progressive consumer protection fall flat.

I grew up in Florida, a land quite literally terraformed by real estate scams. There is an osmotic threshold where scam reality just becomes a reality. Even if the promised future doesn't come to be, some future inevitably does. What kind of future happens in the aftermath of scams? The key question on my mind these days is: how do you keep living in a future that was never meant to

actually exist because it was supposed to be a scam?

Maria Farrell - Silicon Valley's Worldview Is Not Just an Ideology; It's a Personality Disorder

1. Ideology

Silicon Valley's ideology is this: Libertarianism for me. Feudalism for thee. In more detail:

- Surveillance, manipulation and coercion; at first, just for profit, later by necessity, and ultimately for the hell of it.
- Disruption and capture, not competition; monopoly or at least duopoly in each industry it envelops.
- Oligarchy to begin with, creeping autocracy for the win. Overseas autocrats the best of friends.
- Pick me or China wins.
- Ever-increasing inequality and the concentration of capital within a small, interconnected group who back each other's companies and public moves.
- There is no such thing as human rights. There is only identity politics and culture war, which are profit centres.
- Far right white supremacism; libertarianism for white men, forced birth for white women. Eugenics for everyone else.
- A series of bullshit dark utopias designed to drive the hype and private equity cycles, distract and dazzle gullible politicians and policymakers, and convince everyone else that there is no alternative. E.g. crypto-currencies, Facebook's Metaverse, AI and, of course, Mars.
- Systematic racism and misogyny in the workplace, the destruction of organised labour, the ever-worsening of working conditions, extreme inequality.
- Denigration of human agency and creativity, beginning with writers, artists and musicians. Systematic destruction of their ability to earn a living and suggest alternatives.
- Obsessive optimisation along narrow spectrums; externalisation of risks and costs to others, i.e. workers, 'data subjects', the public sector.
- Gutting of independent media, hatred of journalism in particular and accountability in general. Buying out or shutting down all opposition.
- State subsidies and tax dodging. Hollowing out the state. Making private – both in terms of ownership and secrecy – what used to be accountable and universal public services.
- The spoils to the strong, the costs to the weak. Might is right. Winner takes all. The state is an enforcer, not a support. Let the long tail starve.

Silicon Valley ideology is a master-slave mentality, a hierarchical worldview that we all exist in extractive relation to someone stronger, and exploit and despise anyone weaker. Its only relations to other humans are supplication in one direction and subjugation in the other, hence its poster-boys' constant yoyoing between grandiosity and victimhood. Tech bros like Thiel, Musk and Andreessen are the fluffers in the global authoritarian circle jerk. Putin is the bro they'd be tickled to receive calls from, making them feel they're on the geopolitical insider's inside track. MBS is the bro they envy but tell each other scary stories about. Like most of them, MBS inherited his head start in life. He has all the money, all the power, a nice bit of geo-engineering on the side, and he dismembers uppity journalists without consequence. A mere billionaire like Thiel can only secretively litigate them out of business.

Silicon Valley ideology is organising economic, political and social relations into a zero-sum hierarchical chain in which democratic accountability is irrelevant, where beta politicians suck up to the alpha tech-oligarchs, offering their citizens as tribute.* To wit, the thoroughly interchangeable Matt Hancock, Rishi Sunaks, Wes Streetings; all selling out UK citizens' data and life chances for pennies on the pound and a glint of northern California's reflected glory. (Grant Shapps is unusual in having had the initiative to craft and run his very own internet-based get-rich-quick scheme long before he became a government minister.) Most politicians just beg for scraps from big tech's table, enacting the same alpha-beta hierarchy of the incels, but in power poses and slightly better suits. (The early, motivating animus of the Tories' Online Safety Act was envious rage at Nick Clegg's escape to become a tanned and wealthy Facebook lobbyist.)

Silicon Valley ideology is using private equity to buy a new marketplace, flood it with capital to flush out competitors, and use economic dominance to eviscerate working conditions and the cost of labour before jacking up the prices again, this time with the surplus all going to investors. It's hyping specific technologies as universal, structural game-changers in accelerating hype cycles designed to fleece their marks quickly enough to drive growth and cash out before most people realise the technology simply doesn't work as they were told. Bonus points for damaging trusted institutions (crypto) or labour (AI) along the way.

Silicon Valley ideology valorises disruption in the board-room but crushes it in the increasingly digitised and surveilled classroom, and grinds its face into the concrete floor of the Amazon warehouse. Disruption is for CEO's and funders, not for people who protest pipelines or strike to limit labour exploitation. Disruption is something that is done to us. It is not something open to us to do.

Silicon Valley ideology is robbing states of tax and taking over the wrecked public services that result. (I write this from a country whose public tender for health data was bespoke tailored for Palantir. Soon my most private information will sit on Peter Thiel's servers and there is nothing I can do to stop it. Meanwhile, I can't get an appointment to see a GP.)

Silicon Valley ideology blames others for its harms. Its titans built the machines

currently dismantling democracies. So, to absolve themselves of responsibility, they've come to see democracy itself as flawed and weak. Silicon Valley ideology quietly admits (its) freedom is not compatible with (our) democracy. So it wrecks it, destroying our information systems, gutting our infrastructure and essential services, and gathering digital lynch mobs to hound women and people of colour out of public life. Then, like the violent abuser who stands back, momentarily awed at what he has wrought, it says in a moment of startled vulnerability; 'Look what you made me do.' Its tools are beloved of autocrats and authoritarians, and as these tend to be the kind of men it most admires, Silicon Valley ideology has come to more publicly align itself with fascism. It claims the only fix for the violent disorder it foments is more surveillance, more control, at a significant mark-up.

Silicon Valley ideology worships 'intelligence', defined narrowly as mathematical and engineering capability, with all its IQ-related ties to racism, misogyny and eugenics. The worship of 'intelligence' drives ideologues' obsession with billionaire fecundity and longevity, white natalism, space colonisation and the alleged existential threat of AI. The dark futures that set human mathematical intelligence against machine intelligence work not only to recast the grubby scrap for economic dominance as an epic battle against species extinction; they expand the horizon of Silicon Valley's harm from the present, where its built-in biases increase inequality to cause harm and death every day, to the far future, when planetary super-brains might turn people into paper-clips. Their technology reflects this displacement. So, the most important harms are not harms, but risks, and the people best placed to address those are those building the machines that create them.

Silicon Valley ideology is split on this point, however. Its more radical cult, long-termism, centres the omnipotence phantasy of future AI risks, but rather than use these drummed up extinction scenarios to lock in control and economic dominance, the true believers speculate about which AI geniuses to assassinate, to avoid Armageddon. Other devotees fight theological battles similar to the number of angels who could fit on a pin-head, but about just how few survivors will be needed to re-seed humanity, after the nuclear war they believe necessary to forestall artificial generalised intelligence. It's really something to see Silicon Valley's more zealous children turn the trolley problem back on the founders.

Silicon Valley ideology says safeguarding intelligence in the future is more important than its systems systematically crushing and killing black and brown people right now. Long-termism grabs attention back from people being harmed, who were beginning to make too much noise. When confronted with his silence about AI and inequality researcher, Timnit Gebru, who was sacked when she criticised the built-in racism and misogyny of Google's AI systems, British AI star Geoffrey Hinton told Rolling Stone that Gebru's ideas "aren't as existentially serious as the idea of these things getting more intelligent than us and taking over." This tracks with the only critique of technology that Silicon Valley ideology permits – and amply funds via organisations like the Center for Humane

Technology – the concern that algorithmic distraction hijacks smart people’s attention and time. Silicon Valley’s extractive systems are only a real problem when they come after what the tech bros most value, their own brain function and autonomy. Racism, for them, is not ‘existential’. Misogyny is a matter of indifference when your goal is to ‘extend the light of consciousness’ across the solar system.

It’s only when you look straight at Silicon Valley’s leaders you realise its core beliefs aren’t an ideology. They’re a personality disorder.

2. Personality disorder

Silicon Valley’s most famous funders and CEOs veer strongly into sociopathy, narcissism, and abiding Daddy issues. They are thin-skinned, vicious, gormless. Now middle-aged men, their emotional development ended when they made their first hundred million.

They talk about ‘community’ but cultivate slavish fandom and sic their fanboys and lawyers on anyone who disagrees, especially if they’re female, nonbinary, queer, black or brown.

Friendless because they’re incapable of acting as anyone’s true friend, they make a doctrine of misanthropy.

They fear laughter because they’re unfunny.

They fear contempt because they’re roiled by self-loathing.

They prize displays of masculinity because they are cowards.

Optimisation culture masks the abusive and bizarre relations with the human body, primarily their own. Dorsey’s eating issues. The penis electric shock guy. The one who transfuses the blood of his teenaged son. Musk’s compulsive reproduction with multiple women via scientised surrogate.

Partly, they just want to extend their own lives to match the two-century death-grip they plan for their companies. But the body-hacking and transcendence-drive express their horrified bafflement that mortality applies to them, too. Most of us have no choice but to come to terms with physical frailty and with what can be done to us because of the kind of body we’re in. Sick, disabled, fat, queer, black and brown people don’t get to opt out of their bodies, because society doesn’t let them forget for a moment. You might expect such people to wish most to be a brain in a jar or a ghost in a machine orbiting Mars, because of the micro and macro hurts they experience daily. But no, it’s those who won the lottery behind the Rawlsian veil - the wealthy and well, the white and male - who most fervently wish to discard their winning ticket. This tiny group of men self-isolate and self-optimize to excise the messiness of humanity from themselves, and now work diligently to exit from human concerns entirely. There is something fundamental about the human condition that they cannot, will not come to terms with. They require planetary levels of wealth to fuel their death drive to escape velocity.

These are terrible, terrible people. Damaged and cruel, vain and venal. I get why they came to power – by lying, cheating, bullying and stealing, same as any robber baron that ever existed. What’s curious and indefensible is how tightly packed in they are by concentric circles of mostly men who worship them; men who are powerful in their own right, but who accept a subordinate position in the implied pecking order of power. What’s going on with that?

3. Silence of the Sensibles

dividually, the Silicon Valley brats are nothing special. I’ve worked in tech for over two decades, and each arbitrarily made billionaire is as interchangeable an extrusion of post-war, northern Californian high-tech capitalism as was each identical food delivery app it produced. Brats will be brats, and capital is always gonna capital, but it’s the legions of merely moderately powerful enablers who maintain the tech billionaires’ status and ensure their untouchability.

The fawning regard of sovereign wealth fund managers, newspaper editors and senior journalists, industry-funded professors and think-tankers, tame NGO leaders, politicians and SPADs, and the whole masturbatory ouroboros of technology’s own influencer economy, insulates the tech billionaires from accountability. I’ll never get what these men see in Silicon Valley’s boy-kings. I don’t mean that rhetorically. There’s clearly an itch the tech oligarchs scratch for those who brush up against them, but looking at the exact same person, my brain clocks ‘predator’ at a thousand paces, and theirs seem to switch into a purring, excited mode that’s wholly unavailable to me.

(The silent and often determining role played by corporate PR should never be under-estimated. Let’s pause for a moment and think of how many individual PRs used to be journalists, before tech ate media, too.)

The sensibles identify with the aggressor, align themselves with money, flutter like fangirls in the face of power. They never say ‘far right’ or ‘fascist’. They pat themselves on the back for occasionally calling Silicon Valley’s titans ‘controversial’. They refuse even to acknowledge the intentional far right zone-flooding tech companies spray all over their prized values of civility and reason. Their craven servility seems to be based on a misapprehension that the fawning regard for the boy emperors of tech is reciprocal, that the sensibles represent, to those they serve, anything more than somewhat above averagely useful tools.

I read a paper copy of the Financial Times six days a week, and through all the business reporting of Twitter’s (“X”) travails this year, I’ve seen hardly anything about Musk’s insidious but growing antisemitism and overt white suprematism, the NDAs of assaulted women and his hounding of female public figures, the systemic racism in his factories, his illicit drug usage and emotional volatility, compulsive lying, personal amplification of Nazi content, and the conspiracy theories he now feels at total liberty to spread. (And no, FT, taking Musk’s tame biographer out to lunch isn’t journalism.) The silence of the sensibles lets business reporting live in a magical universe where politics doesn’t intrude, at least when the politics is on the right. There’s a consequence-free sweet spot

where billionaire bootlicking and centrist civility co-exist. The merely influential are rarely the first to say the Emperor has no clothes (or is a Nazi in plain sight).

4. We are all very, very tired.

I, quite frankly, am tired. I find myself yet again in a conversation dominated by beneficiaries of a dirty system while the conscience, critique and force of collective action for alternatives are provided by women, and women of colour, predominantly.

Observer columnist John Naughton wrote two pieces a while back, pointing out that Silicon Valley's most able critics are nearly all women or nonbinary. Of course, not all men, etc. etc. and, for damn sure, not all women either. But we – I speak as one of Naughton's "formidable female tech critics" – have our own ideas and dreams, too. We didn't grow up yearning to pen exquisite critiques of shitty ideologies. We'd like to build our own things, too, you know? When one moderately powerful person steps up it emboldens others to act. It would signal to Musk's shoulder-shrugging supporters inside US government – and especially the DoD – that you cannot run critical communications and defence infrastructure while being a far-right stooge sympathetic to foreign powers.

So, to the moderately powerful men who prop up and benefit from Silicon Valley's ideology, and leave us the real work of fixing what it breaks, here is some of what we've learned and hope you might use:

Silicon Valley's ideology and personality carry within them the seeds of its destruction. Its thin-skin means it can often be rattled and made to back down. (The trick is watching for the same bad act to be done again later, just more quietly.) Policymakers are rattled, too. And so, but only a little, are university VCs who still accept donations and suborn research. The social licence of Silicon Valley is now just a provisional permit, at risk of cancellation each time a new outrage is revealed or new legislation produced.

Silicon Valley's totalising concentration makes its companies run increasingly afloat of competition authorities, who've belatedly come round to broader notions of economic harm and the possibilities for transatlantic cooperation.

Silicon Valley's titans are so far removed from reality they can no longer distinguish between whims and crimes. The criminal lens is resisted by regulators, prosecutors and commentators alike, but it increasingly pulls systematic harms into focus.

The security lens is stricter again, as certain billionaires' behaviour clearly flouts the requirements for security clearances they need for all those defence contracts. And their companies are far more brittle than they appear.

It's not a lot, but it's a start. Get to it.

** About the politicians sucking up to tech oligarchs, I just want it out there that I wrote this six weeks before Rishi Sunak's toe-curling AI Summit "fireside chat"*

with Elon Musk, where Sunak looked like some rich guy's kid on work experience, mortifying the few people left who saw the UK as a serious country.

Shazeda Ahmed - From Algorithmic Monoculture to Epistemic Monoculture? Understanding the Rise of AI Safety

From November 1-2, the UK government hosted its inaugural AI Safety Summit, a gathering of international government officials, AI business leaders, researchers, and civil society advocates to discuss the potential for creating an international body to govern AI, akin to the IPCC for climate change. On its surface, ‘safety’ appears to be an unobjectionable concern after years of instances in which AI systems have caused errors that have denied people state benefits and cast them into financial turmoil, produced hate speech, and denied refugees asylum due to mis-translations of verbal testimony.

Yet the conception of safety that motivated the Summit is unconcerned with this category of harms, instead looking to a future hundreds of years from now where advanced AI systems could pose an ‘existential risk’ (x-risk) to the continued existence of humanity. The ideas behind the emerging field of ‘AI safety,’ a subset of which operate on the assumption that it is possible to prevent AI x-risks and to ‘align’ AI systems with human interests, have rapidly shifted from a hobbyist interest of a few small communities into becoming a globally influential, well-resourced, and galvanizing force behind media coverage and policymakers’ actions on preventing AI harms.

Where did these ideas originate, what material outcomes are they producing in the world, and what might they herald for the future of how we regulate and live with AI systems?

In spring 2022, my colleagues at Princeton University and I began to track AI safety’s growing influence and sought to map its intellectual origins. How are the ideologies that underpin this new field moving people, money, research, community-building, and career advising—in sum, the activities that people within AI safety refer to as ‘field-building’—towards a utopic vision of living with AI?

We began to see a broad-strokes argument unfolding that as large AI systems such as large language models (LLMs) including ChatGPT scale up, they could develop advanced capabilities beyond those their original creators had anticipated, and cause widespread harm to humanity if left unchecked. Related to this fear of rogue AI systems, AI safety proponents worry about whether a human bad actor could amass vast quantities of computational power (“compute”) to build a bioweapon.

To understand the genesis of these fears, we read texts foundational to x-risk studies, and traced their relationship to both the concept of long-termism—concerns with the future of human existence hundreds of years into the future (and longer yet)—and the effective altruism movement. Effective altruists (EAs) are concerned with how to “do the most good” given finite resources. Drawing

from utilitarian philosophy, they seek to optimize returns on expected value in a range of proposed interventions for cause areas including the prevention of pandemics, nuclear wars, and AI x-risks.

Early in our research, I attended a talk by Sam Bowman, an NYU professor of computer science who had taken leave to work at Anthropic, a company founded by former OpenAI staff with the purported aim of building safe AI systems. Bowman highlighted a point that later came to be central to our research: a variety of communities have sprung up around ideas such as EA, x-risk, and long-termism, and people within these communities are in the majority of those advocating for AI safety.

Yet a growing number of people coming to the field of AI safety have no affiliation with these ideas. The latter group nonetheless must pass through institutions that have been shaped by what my colleagues and I refer to as AI safety’s “epistemic culture” (from sociologist Karin Knorr-Cetina): the cultural practices of how knowledge is constituted and disseminated within the sub-communities that unite to work towards a shared idea of AI safety.

To better grasp what is included in this emerging field’s conception of “AI safety,” we started by identifying what research, community cultivation, fellowships, and institution-building in the field were being funded. Then we descriptively analyzed the picture that these funding stream’s outputs formed. What emerged was a clear image of a tightly networked series of communities with at least four distinct features. The first is online community-building— both through EA web forums such as LessWrong and EA career advising hubs such as the nonprofit 80,000 Hours. Someone who may have read Oxford philosopher Nick Bostrom’s book *Superintelligence*, for instance, may first encounter the idea that future AI systems can one day ‘surpass’ humans at performing every task and attain artificial general intelligence (AGI).

Then, seeking a community of like-minded others debating the consequences of such a future, they could find a lively set of discussions about this topic on the Effective Altruism Forum, or the more recent Alignment Forum. Career advisories like 80,000 Hours create a pipeline where young people in these networks can find jobs in EA cause areas, for instance working on AI alignment at OpenAI.

The second feature is AI forecasting, which takes two forms. One involves hiring professional forecasters to cast and defend predictions about specific outcomes, such as the date by which a particular model may attain a specific benchmark— if a model outperforms the forecaster’s guess, some in the field see it as a sign that AI systems are developing at a clip that humans attempting to ensure safe deployment cannot keep up with (therefore, the logic goes, this would justify the need to invest in more AI safety research.)

The third feature constitutes this research itself, which focuses on issues such as monitoring emergent capabilities as AI systems scale up and develop unanticipated new functions, developing methods for alignment of AI systems with

pre-defined human values to avert x-risk, defending the robustness of these systems against highly improbable but massively destructive events such as crashes of automated financial trading systems, and finally the catchall of “systemic safety,” referring to deployment of systems in context—including issues such as privacy, cybersecurity, and algorithmic bias.

Many of these ideas are disseminated via the fourth feature, prize competitions where entrants can submit code or papers in response to technical challenges in the field. Given the deep coffers of AI safety funders, and the urgency with which they believe AI safety must be addressed, prize competitions in the field carry massive prize pools ranging up to \$1 million split across winning participants. While some of these competitions are hosted by AI safety nonprofits, we saw a trend toward competitions being embedded in academic computer science conferences as part of a broader effort to bring AI safety into mainstream academic computer science.

What happens when this field has influence beyond its epistemic community? Media coverage has tended toward replicating AI x-risk narratives, platforming the small handful of companies and figures at the center of this epistemic culture and positioning them as speaking for a larger constituency than they represent. For instance, companies such as OpenAI and DeepMind have long made it their mission to attain AGI— and OpenAI CEO Sam Altman not only presents this as possible, but as a path to a future utopia. The flipside of this premise is often presented within AI safety as one where we fail to rise to the challenge of ‘aligning’ AI with human interests.

Laying bare the ideological scaffolding of AI safety raises the question of whether these ideas will continue to be central to the field going forward, and how essential they are to pursuing the field’s definition of safety in the first place. Rishi Bommasani et al describe the algorithmic monoculture that arises when a small handful of companies like OpenAI or Anthropic produce ‘foundation models’ such as ChatGPT that a plurality of third party actors across the platform economy come to rely on, from Spotify to Zoom.

In parallel to these same companies that are the leading industry actors in the field, do the ideas underpinning AI safety have the potential to become an epistemic monoculture, crowding out other points of view? Furthermore, does buying into one part of this framing of AI safety essentially amount to being indefinitely locked in to a utilitarian, x-risk centered approach? If so, what might the consequences be?

Eight years ago, Oxford philosopher Amia Srinivasan wrote about the effective altruism movement as dodging specificity in favor of general approaches to mitigating x-risks. This is echoed in how organizations in the epistemic community are presumed, based on names like the Future of Life Institute or the Future of Humanity Institute, to see protection of the entirety of humankind as an achievable mission.

A recent coalition of Chinese and US organizations and individuals proposed

in light of the UK’s AI Safety Summit that companies spend up to 30% of their research budgets on AI safety. If this happened, what opportunities might prioritizing this generalist approach foreclose? One possibility is that finding solutions to well-documented problems— discriminatory uses of AI tools in hiring, or the wide range of shortcomings in AI functionality — would, in this framework, become downgraded in urgency.

Srinivasan notes that part of the wide appeal of effective altruism is that it ultimately does not challenge, but rather reinforces, the status quo. As with other ideologies that have incubated in Silicon Valley, while many ideas coming out of the AI safety epistemic community can appear at first to run counter to the mainstream, they nonetheless embed neatly into what Henry Farrell and Abraham Newman term the ‘weaponized interdependence’ model of the United States.

For instance, leading figures in the AI safety community have long treated compute” resources as akin to a wartime resource that the United States government must monitor to prevent individual bad actors from amassing them to build bioweapons— a position that has become solidified in the mainstream through the recent White House Executive Order’s inclusion of references to governing foundation models using the Defense Production Act. While some commended the EO for acknowledging the need to protect workers and alleviate social harms of AI systems, the language regarding these known harms that stem from historical injustices was far vaguer, calling for “standards” and “guidelines” whose enforceability remains an open question.

Understanding the epistemic culture of AI safety can help us anticipate what industry leaders and other influential academic and philosophical voices in the epistemic community will advocate for in the future. In a recent interview with Foreign Policy magazine’s editor-in-chief, White House Office of Science and Technology Policy representative, sociologist, and leading architect behind the Blueprint for an AI Bill of Rights Alondra Nelson noted that the government has opened up the conversation about regulating AI to public debate, encouraging citizens to appeal to their representatives for advancing legislation.

What do people advocate for when a prevailing narrative is that AI may one day end humanity, and that we must put faith in a small group of technocrats who claim to speak for all of humanity to prevent that outcome? And how do we avoid a doubling down on the status quo of self-regulation and voluntary agreements that AI safety appears to be amenable to? As we head towards what many are hoping will be a banner year for more concrete regulation of AI, the AI safety epistemic community should be more receptive of external critiques, and aim to be accountable— both for the knowledge they produce that has gained major global influence in a very short period of time, and for the specific, contextual harms that arise in the AI systems they aim to improve.