SECOND AND THIRD ORDER EFFECTS OF A.I.

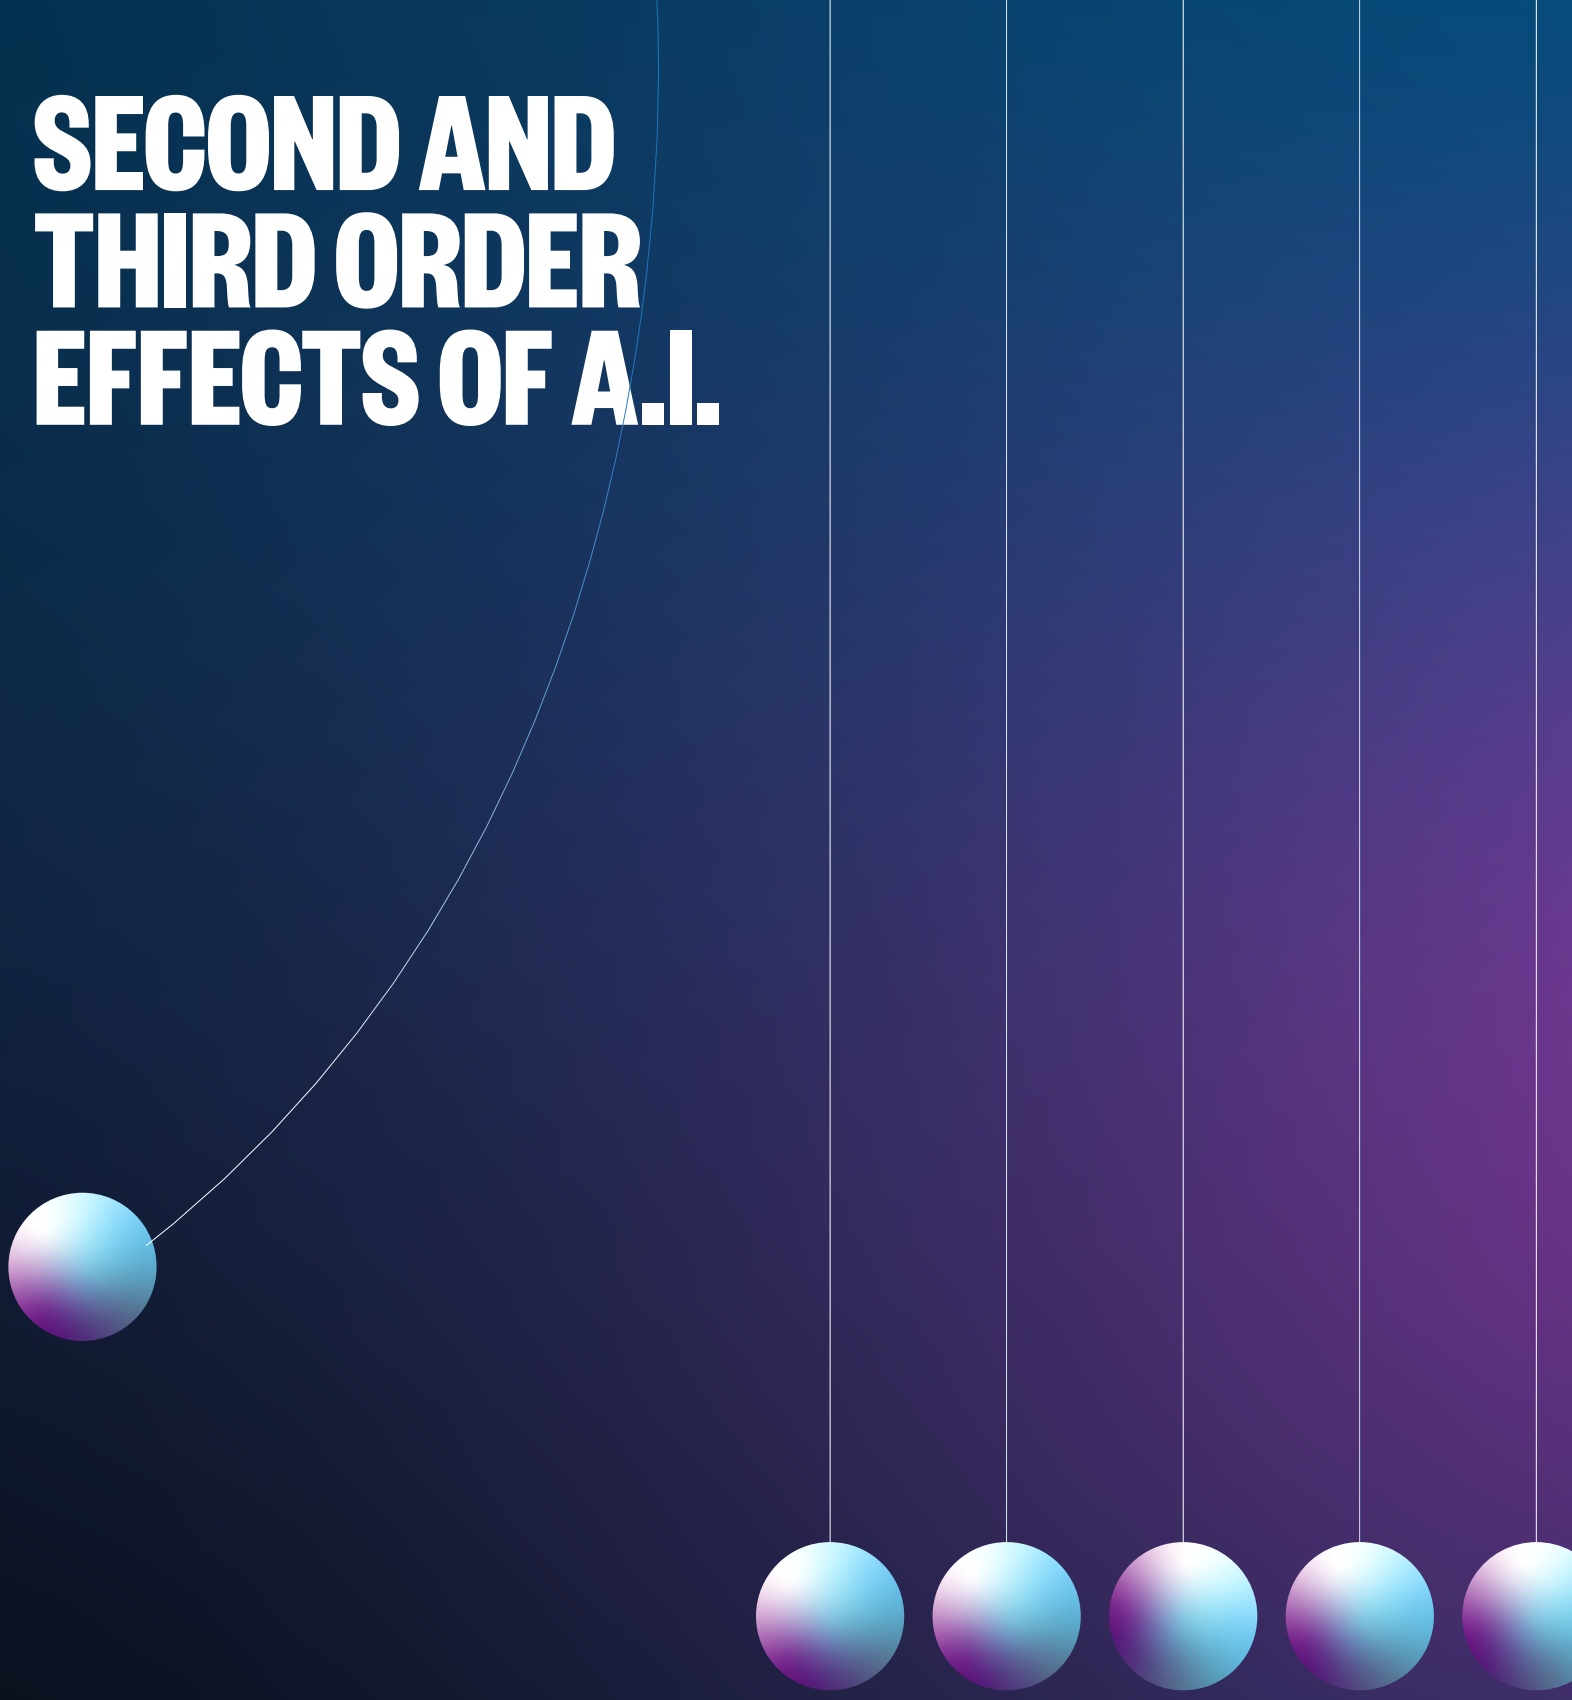

ASPEN DIGITAL
aspen institute

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This document provides a summary and an analysis of the most significant themes and key insights from the foresight exercise delivered by Aspen Digital in August 2024 at the Doris Duke Foundation's Shangri La estate. The gathering was focused on the second- and third-order effects of artificial intelligence (AI). We used a series of foresight methods to imagine a series of possible futures and their implications. Participants from diverse fields (including technology, academia, philanthropy, and the arts) engaged in discussions, offered provocations, and constructed scenarios under the Chatham House Rule, to explore the potential societal impacts of rapidly evolving AI technologies. The gathering was intended to be a kickoff for a broader, more distributed set of discussions on the theme, with an overall aim to mitigate projected harms and amplify potential goods.

The Aspen Institute worked in close collaboration with Omidyar Network to identify a list of individuals who would uniquely contribute to a series of wide-ranging, future-focused discussions on AI. The facility at Shangri La, on Oahu in Hawaii, provided a serene and secluded location for deeper engagement over a two-and-a-half-day event that included full-group sessions, breakouts, and social engagements to bring the participants together.

Early on in the planned agenda, participants engaged in an exercise that asked them to identify transformative technologies of the past 75 years. They were then prompted to identify the second- and third-order effects of those technical, social, and cultural revolutions. Considering innovations ranging from the introduction of residential air conditioning to the washing machine, and from GPS to national weather networks, they identified a myriad of technologies that have led to widespread societal change. In those retrospective discussions, and with the benefit of full hindsight and the context and perspective of history, they were able to connect the dots, not only from A to B, but also from A to C, D and E. This hindsight exercise acted as a warm-up for subsequent discussions of how AI may lead to transformative social change in a multitude of ways. Admittedly, rooted in the present, it is challenging to see into the future. Moreover, AI presents a particular conundrum as it touches on just about every field of human endeavor. The landscape of possible futures, and

possible downstream impacts associated with them, is vast. Yet the diverse group of people gathered, including talented thinkers, leaders, advocates, and builders, allowed for dynamic, insightful, and productive conversation.

Participants brought with them their varied backgrounds, talents, gifts, and key competencies. With time and trust, this diversity of experience can often push and pull people to explore new ideas. Yet it can be hard to pull focus, especially when harms are closely felt, and intense feelings rise. There is also an intense fatigue among those who have been active in trying to build more fair, just, and equitable technology futures. Among this group in particular, AI's arrival feels like rocket fuel poured on to an already raging inferno of unresolved issues, known harms, and aggrieved policy arguments.

With this in mind, it's no surprise that the group felt the gravitational pull of the first-order effects of AI on economics, democracy, media, and society and culture. They often focused on near-term measures to both mitigate the potential harms and lift up potential benefits. Put more plainly, they were still processing today's reality, and while that helped them imagine what was to come, it took time and was not fully realized. In retrospect, this is likely a necessary and needed part of the journey that will help us find our way to imagining more distant futures. The findings of the group should be considered a first leg of a longer journey, charting a course towards unknown futures.

Despite these challenges, the group rose to the occasion, and engaged deeply in the complexity of the topic and embraced the task assigned to them. As you will see in the summary below, their discussions took them in a myriad of directions, with a series of discussions covering nearly every aspect of modern life—from the economy, to climate, to social cohesion, to education and the very nature of knowledge.

Everyone acknowledged the uniqueness of the moment: while AI is obviously more than a new technology, it is one unlike any that has come before, presenting challenges and disruptions at a speed and scale we have not seen. One participant noted that the story of a technology is not just that of its invention, but also of its adoption. While the AI tools that captured the public's imagination in 2023 were unprecedented, the way in which they have been released and adopted is even more surprising. One

participant lamented the degree to which incomplete (or even non-functional) applications have not only been released to the public, but aggressively sold to businesses and consumers. Today's state-of-the-art AI tools routinely make glaring errors, and fail at common tasks and queries. ChatGPT 4.0 had a 'hallucination rate' of over 25%. By comparison, in 1994, Intel recalled $475M worth of Pentium processors because they occasionally produced mathematical errors under certain rare conditions (1 in 9 billion operations).

The urgency and pressure to compete is also driving unprecedented capital expenditures, unsustainable energy consumption, and unrestricted drive to competition with other nations. As a result, the adoption and widespread use of AI has urgent downstream implications. At one end, the demands of investors to see returns are accelerated, because the cost of development and operation is so astronomical. At the other, the raw power required to provide these services with as-of-yet-unoptimized AI models is creating an energy demand that cannot be met and sustained by today's power grid. Finally, the emergence of AI has unleashed a new arms race, with companies speeding to develop and release models and establish market leadership, especially with China, and a corollary pressure from those companies to avoid regulation or guardrails, which could reduce harms, but would also slow development and increase innovation costs. In the short term, it seems the decision is to spend more, consume more, and regulate less.

EMERGING NARRATIVES AND THEMES

The following outlines the emerging narratives and major themes of our large and small group discussions. It is by no means a complete list of every topic mentioned, but shines a spotlight on the most significant topics and ideas. Most importantly, while the ideas below may be presented with definition and forceful clarity, our discussions often highlighted the uncertainty, complexity, and diversity of thoughts, approaches, and opinions. Our assembled group included enthusiasts and skeptics, practitioners and philosophers, whose varied expertise took us in many directions. Finally, we must acknowledge that we don't know what we don't know: in many ways, today's AI tools have yet to fully reveal themselves—the heights of their potential, and also the depth of their failings.

With these considerations in mind, our major themes included:

- **Development, Governance, and Competition**
- **Labor and Augmentation**
- **Education, Epistemology, and Trust**
- **Democracy, Social Cohesion, and Human Connection**
- **Art, Culture, and Creativity**

DEVELOPMENT, GOVERNANCE AND COMPETITION

"Is this just a toolkit for authoritarian regimes? Or is it a way to equalize computational power?"

MOMENTUM AND FOMO

People are simultaneously scared and excited about AI's development and deployment. Some experience feeling both at the same time, though the focus of conversation was largely on the negative. With a few exceptions, participants expressed frustration and anger at the realities they see emerging in the AI field. They feel like Cassandras, as the rest of society and industry hurl themselves into a new technological era, all while repeating the mistakes of social media, blockchain/crypto, and beyond, fueled by a combination of ambition, curiosity, and FOMO.

EMERGENT TOOLS AND OVERWHELMING HYPE

Despite enormous progress, many AI tools are still in their infancy. Companies have released half-baked products and turned mainstream consumers into their beta testers, often for products that can't perform basic functions reliably, or simply don't work at all. This represents a notable and dramatic externalization of the risk and costs of product development. One participant talked about quality standards and the acceptance of risk of innovation, and the move to externalize the harms. One cited Intel's recall of faulty Pentium chips and LLM tools being deployed despite major failure points.

We are continuing to experience a massive hype and investment bubble. A small number of (mostly) American companies have dominated the conversation regarding AI, including its various applications in the forms of prompted content, algorithmic labor, and other potential. A former OpenAI employee raised $1B at a $5B valuation for a spin-off 'company' that had fewer than 10 employees and no working code. There are early signs that Wall Street may pull on the reins of this expansive phase as AI requires

a capital investment that is unparalleled in previous technology cycles; one that will demand returns or at a minimum, the ability to cover the debts incurred for energy, compute, hardware, chips, and cloud.

OVERSIMPLIFICATION OF A.I. IN MULTIPLE IMPLEMENTATIONS

AI is not one thing—its application in transformer-based LLMs is not the same as how it is being developed in other fields or uses, like drug discovery (or how it has been developed and deployed for the last 20 years in consumer and industrial applications). There is no single thing that is "AI." There are simply more and less complex algorithms. Those algorithms have different purposes and uses, and they come from differing lineages. Public discourse needs nuance here to be more useful and avoid generalizations.

We shouldn't let our view of today's technology be so myopic as to think that LLMs using transformers are the only innovation we need or will see; the flaws, liabilities, and limitations embedded in these technologies are already becoming more apparent. Participants referenced Yann LeCun, who has argued for new approaches, and an MIT study showing how LLM's are likely not capable of reason (despite being very good at appearing to reason).

MONOPOLY AND CONCENTRATION OF POWER

Concerns about monopoly and corporate capture of regulators and legislators are a serious consideration. As one participant noted, "We're living in a world where AI has been defined by a handful of tech companies and the people who run them, and we're reacting to that. We (our organizations and foundations) could develop an alternative narrative. Storytelling doesn't stand alone, but history shows the power of storytelling. FDR told a different story that changed America forever. I'd like us to think about the power of words and the power of the way we talk about these things. We are living in Silicon Valley's imagination right now. There's a moment here to shift away from someone else's language."

LABOR AND AUGMENTATION

"AI is primarily a knife-sharpener for the blade of capitalism. They're salivating at the prospect of what they can do with very sharp knives."

"We were promised the Singularity, but we got DoorDash."

DIGNITY OF HUMAN LABOR

AI has the capacity to impact all human endeavors—at least, all endeavors that engage a computer, which includes nearly all white-collar work, and substantial aspects of blue-collar jobs. That's an existential threat to the frameworks that sustain today's society, particularly for people who work for a living and don't want to be replaced by something that is cheap and tireless. Humans gain a huge amount of identity and personal meaning from their work—they can become alienated from it, experience moral injury when working in hostile environments, and grow demoralized when their expectations do not align with their environment. Conversely, they can also experience validation, agency, dignity, fellowship, and feelings of self-worth and self-respect while at work.

LOSS OF MEANINGFUL WORK

AI could contribute to a reality where fulfilling and creative work, one of the highest expressions of human endeavor, becomes a luxury available only to a select few, while most others are subjected to a kind of AI-driven exploitation for the benefit of technology elites. Participants talked about working "above and below the robots"—describing those whose work is assigned to them by AI bosses, and those who direct AI to serve them. In this vein, AI could exacerbate existing labor inequalities, concentrating benefits among a smaller number of elites while promoting greater immiseration among the most precarious workers.

If AI threatens to make more work even more meaningless for more people, it could lead to a crisis of meaning, value, and purpose for everyday Americans and others all around the world. Historically, these types of crises have not gone well: they end in war, in terrorism, in ethno-nationalism, and fascism.

LOSS OF SKILLS AND MASTERY

AI could result in the loss of expertise and mastery as AI tools automate tasks and replace skilled workers, leading to a workforce that is less capable and less knowledgeable. This in turn could bring about the "Last Masters" problem, where the next generation of experts is not adequately trained due to an increasing reliance on AI tools, coupled with the absence of training for junior-level tasks (now replaced by automation).

AUTOMATION AND AGENCY

AI could help to create a more equitable and just economy by automating tasks that are dangerous, repetitive, or low-paying, freeing up workers to pursue more fulfilling and meaningful work (assuming their jobs are not simply eliminated, and those savings taken as profits). The potential for an increased reliance on AI in the workplace raises important questions about the future of human agency and control. Conversely, AI could empower workers by giving them access to new tools and technologies that can increase their productivity and efficiency. AI could also level the playing field by providing new opportunities for learning and knowledge creation, especially benefitting those who are neurodivergent or who lack access to traditional education. Could AI make management better? Imagine an AI manager that does the weekly schedule—one you don't have to curry favor with, or who doesn't hold a grudge, denying you desirable shifts? Most shift workers would be thrilled to replace that boss with a robot.

It's possible that several of these scenarios will come true at the same time. Will AI augment human capabilities or be deployed in ways that replace human decision-making altogether? Will AI lead to greater worker autonomy or to a more automated and controlled workplace?

EDUCATION, EPISTEMOLOGY, AND TRUST

"Having to work through a problem yourself is itself knowledge production. Basically, this could make a more brittle knowledge system."

"I just want people to start majoring in history and humanities instead of computer fucking science."

COMPARISON REDUCES OUR ABILITY TO SEE A.I.'S POTENTIAL

AI resembles some transformations of the past (e.g., search, automation) yet it is fundamentally an entirely new thing. We struggle to wrap our minds around all that this entails. Applying comparisons always forces us to leave out crucial elements. Treating AI as unrelated to its preceding technologies denies us the ability to apply the lessons of past technologies (e.g., bias, exploitation). As Tobias Rees asks in a recent interview about AI,"... is there anything in the here and now that is new, so radically novel or different that we cannot think about it in terms of the assumptions that we have inherited from the past? … AI lies outside of the categories that thus far have stably organized our world. I want to understand this outside, to develop new vocabulary, to make it navigable."

We still see AI as a machine, a replacer, and an automator of things we know. Yet at the same time, we often dream about the aspirational 'Star Trek' version of AI — an embodied voice of a computer with superior analytical capacity to answer any question on the spot and to solve problems as a member of our crew. We picture life-like robots imbued with the agency to act on our behalf and to work in concert with us. While this still feels like the future, not the present, our desire for it (and its touchstones in modern popular culture) make it easy for us to be lured into solutions that promise such a reality. Still, attempts to commercialize agentic AI have been embarrassingly bad (Rabbit R1, the AI pin) and do not come close to the idealistic future we aspire to.

CHALLENGING ESTABLISHED CONCEPTS OF KNOWLEDGE AND REASON

AI, both as a technology and as something of a cultural moment, is challenging our ways of understanding knowledge, agency, and reason. This is so in terms of how these concepts can be represented in technology, and also how we think about our own human capabilities. Previously, the paradigm was simple: human beings possess the power of reason, while plants, animals, and objects do not. But now there are "objects" (AI) that appear to reason, though in fact they are responding to prompts by analyzing past behaviors and offering 'the most likely' response. Yet, the risk is that we misinterpret plausible AI responses as human, despite the fact they are not made from actual reasoning. Equally concerning is that these 'impersonations' may still be completely adequate for an entire category of uses—uses that, if deployed, could displace human labor.

PERSONALIZED EDUCATION: RESPONSIVE, ADAPTIVE, AND LONELY

The group discussed how AI-powered personalized education could lead to increased homeschooling, furthering societal tribalization and the decline of shared public experiences. Those who had direct experience of the difficulty their students had with advanced literacy, focus, nuance, research, fact-checking, or analog methods of recording and note-taking all worried that further automations would only limit their students' baseline capacities.

In an AI-powered world, some wondered why students would ever bother attending school or investing in their education, but learning anything at all, ever again. This could result in an erosion of public identity and shared narratives due to the loss of schools as geographical tribes, with a reinforcement of existing educational biases. Participants imagined how this could lead to a widening gap in educational quality, a rise in the importance of standardized tests as gatekeepers to knowledge and opportunity, and, ultimately, the potential for AI to exacerbate inequalities in access to quality education.

Overall, there was a feeling that the reliance on AI in education could lead to overdependence on technology and a decline in critical thinking and human interaction. AI tutors could personalize learning, but their long-term effects on motivation and the

potential for simulating parental affection raise concerns. That said, it was not all negative; one participant shared a compelling and moving story of their child's struggles with learning, and the ability of today's tools to improve their ability to understand the materials, and advance at their own pace.

A.I. IN ACADEMIA SHIFTS POWER AND METRICS

Participants made several observations about potential futures in academia. AI tools that enable non-technical researchers to produce analysis of large swaths of scientific or qualitative data (e.g., substantial collections of either a great volume of works, or many works in fields not well known by the researcher) could dramatically change who can take on research, and also how it is conducted, potentially black-boxing knowledge and altering the power dynamics between students and professors.

Separately, AI could shift the metrics used to evaluate the impact of research. Currently, the simple metric of citations is the prevailing determinant, but one that fails to capture many important aspects of a research product's innovation, influence, sustained relevance, and even how it may be disproven over time. New AI tools might be developed to measure the impact of scholarly work in real-time (and over time, constantly re-evaluating), potentially incentivizing certain types of knowledge over others and raising concerns about transparency and the "quantified self" in academia.

DEMOCRACY, SOCIAL COHESION, AND HUMAN CONNECTION

"What we tell our machines is important."

"Children need answers from people who love them."

IMPACT ON HUMAN CONNECTION

Participants worried about the potential for AI to replace genuine human connection with artificial interactions, leading to increased loneliness and isolation. In a world of ubiquitous, always-on chat companions, is there a risk that AI could replace small talk, potentially altering how humans interact with each other? Will increased reliance on AI lead to further social isolation and loneliness, particularly if it enables further "exit from common spaces"? Concerns were also voiced about AI fostering echo chambers and reinforcing existing biases, ultimately leading to increased polarization and tribalization.

Conversely, the group also explored opportunities for AI tools to help people avoid loneliness, including providing avenues for aging populations to build connections and supporting neurodivergent people in making more effective connections with others. AI might be developed and deployed to promote understanding or connections across different groups, perhaps through pro-social media platforms that connect people with shared interests, and analytical tools that help bridge different understandings and ways of speaking.

EROSION OF DEMOCRACY AND A RISE OF AUTHORITARIANISM

Participants expressed a fear that AI could further the concentration of power in the hands of a few, particularly large tech companies with profit-driven motives. This concentration, coupled with AI's potential to manipulate behavior and spread disinformation, could lead to the erosion of democratic systems.

The increasing use of AI in governance and decision-making processes, particularly in areas like law enforcement and the military, heightened apprehension about the lack of transparency and accountability. These concerns spanned both the nature of the

data and training that was used to develop those tools, and also the lack of public knowledge of when such tools are being deployed and for what purpose. That absence of knowledge could deny the public agency and the opportunity of defense against negative outcomes. AI's potential to amplify existing power structures, especially when trained on biased datasets, could exacerbate inequalities based on race, gender, class, and access to technology.

AMPLIFICATION AND TARGETING OF DISINFORMATION AT SCALE

AI could contribute to the spread of misinformation and the erosion of trust in institutions, particularly journalism, further undermining democratic processes. AI-powered tools for political manipulation and propaganda could exacerbate societal divisions and political polarization, potentially leading to the rise of authoritarianism. AI avatars deployed at scale might intimidate and persuade, while also inciting hatred, division, and violence. That said, the group felt that mis- and disinformation and propaganda were perpetual societal issues, and that AI was an accelerant more than a source. While AI is concerning, the risk continues to lie with bad actors, who remain the real issue.

Several participants raised concerns about AI and its impact on journalism, particularly regarding the erosion of trust and the distortion of perceptions and trust in the public sphere and institutions. They considered how generative AI might replace the written work of journalists, with compelling prose that lacks oversight of bias or veracity, flooding the infosphere, corroding the fourth estate so much so that journalists might become vilified and the targets of hostility.

ART, CULTURE, AND CREATIVITY

"Fulfilling and creative work is one of the highest expressions of human endeavor."

"Animators are becoming directors."

DEMOCRATIZATION OF CREATIVITY AND INNOVATION

AI tools could make creative tools more accessible, potentially allowing anyone to create art, and narrowing the gap between talent, skill, and high-quality creative outputs—anyone could become a "James Cameron". The group discussed both positive and negative effects of this increased access to artistic tools. AI tools could destabilize creative industries or lead to a reconsideration of the value of art in society. On one hand, a potential flood of derivative, low-quality work could devalue creative output, echoing the transition from skilled artisanship to "cheaply made mass-produced 'premium mediocre'" seen with industrialization. On the other, if anyone can create, then value might shift from the act of creation to inimitable factors like taste, originality, or access to stars.

AI could foster new models of creativity, emphasizing collaboration and remixing over individual originality. The capacity for AI tools to augment and accelerate the creativity of an artist were discussed at length. One participant described how today's emergent tools are allowing animators to create works of unparalleled detail, because the amount of time required has been diminished by orders of magnitude. This is a counter to the waves of derivative and imitative works that are seen coming from most diffusion image generators. Some participants were particularly concerned that displacing the creative arts and creative acts generally—the act of "solving a problem" in an unexpected way, regardless of the field or area of study—would diminish humanity's ability to adapt in the long term.

CONTROL AND OWNERSHIP

The group discussed control and ownership of AI technology, particularly its impact on intellectual property and the potential and patterns of exploitation and labor theft. There are substantial copyright questions that remain unresolved, and many court

challenges progressing through various jurisdictions. Even with new rulings, there likely remain many unresolvable, no-turning-back-level issues of copyright and appropriation that have created an existential challenge for those who create, and those who sell or share their art—from illustrators to actors. We likely can't undo what has been done to artists and other creators.

One exciting question is around collective ownership of community-developed models. Could workers band together to develop their own AI models and assert control over their creative output? It is notable how quickly this existential threat to creators has spurred legal challenges and negotiations, from court cases to collective bargaining. As these agreements play out, consumer protection or other regulations focusing on data ownership and usage could offer a way to address concerns about AI's impact on creative industries.

KEY INSIGHTS

- **The hype cycle:** There is a need to move beyond the hype surrounding generative AI and explore a broader spectrum of AI applications and their potential impacts.
- **Shifting power dynamics:** AI has the potential to shift power dynamics in society, both within and between nations.
- **The role of capital:** The current model of AI development is heavily driven by private capital, raising concerns about access, control, and the distribution of benefits.
- **The need for a public option:** The development of a public option for AI is crucial to ensure societal benefit and mitigate the risks of monopolization and control by private interests.
- **The importance of human agency:** Maintaining human agency in an AI-driven world is essential, with emphasis on fostering creativity, thinking critically, and navigating complex information landscapes.
- **The value of alternative perspectives:** Incorporating perspectives from diverse cultures and philosophical traditions is crucial to envisioning a future where AI benefits all of humanity.

RECOMMENDATIONS

The gathering explored both the risks and opportunities of AI in possible futures. With that in mind, we prepared a series of recommendations that provide methods to respond to both the promise and the perils of AI.

RECOMMENDATIONS TO TAKE ADVANTAGE OF THE PROMISE OF A.I.:

- **Focus on AI for augmentation, not just automation:** While automation grabs headlines, the real potential of AI might lie in its ability to augment human capabilities. We should think about ways that AI tools can help someone do their job better (e.g., solve problems more effectively, or access information more readily). One participant noted that animators are becoming directors, suggesting a shift towards higher-level creative work enabled by AI. How can surface AI tools that enhance skills and creativity rather than simply replacing human tasks?

- **Develop digital literacy and a critical eye:** As AI becomes more integrated into various aspects of life, it will become essential to cultivate digital literacy and a critical understanding of how AI systems are integrated into daily life, how they work, their limitations, and their potential biases. We should seek ways to enable continuous learning to stay informed about AI advancements and their implications, as they are likely to impact all aspects of modern life.

- **Advocate for transparency and accountability in AI development:** We should support initiatives that promote transparency and accountability in the development and deployment of AI systems. This includes advocating for regulations and ethical guidelines that ensure fairness, preventing discrimination, and holding developers responsible for the impacts of their creations. We should also embed human values into the AI tools we deploy in our society and workplaces. Participants emphasized the need for more open and accountable AI development processes.

- **Explore and experiment with AI tools and platforms:** As new tools are emerging, we don't know what they're capable of (and often their creators don't yet either). We must not be afraid to experiment with different AI tools and platforms to

discover how they work, where they fail, and what they can do. We should encourage engaging with AI in a learning capacity to explore its creative possibilities in their work, showcasing its potential in diverse fields.

RECOMMENDATIONS TO RESPOND TO AND AVOID THE ANTICIPATED HARMS OF A.I.:

- **Be aware of the potential for bias and discrimination:** AI systems are trained on data, and that data reflects existing societal biases, so the AI will perpetuate and even amplify those biases. We need to be vigilant about identifying and challenging bias in AI systems, particularly those used in decision-making processes that impact individuals' lives. Participants' statements throughout the workshop highlight the need to address the ways AI can exacerbate existing inequalities.

- **Support the development of a public option for AI:** Publicly-funded and publicly-controlled AI infrastructure can serve as a counter to the dominance of large tech companies and their profit-driven motives. Public tools can have public interests, values, and constraints that reshape how AI is deployed, trusted, and understood. We should advance a public option to ensure more equitable access to the benefits of AI and help steer its development in a more socially responsible direction.

- **Develop educational campaigns to increase understanding and promote discourse:** We should consider a campaign to demystify AI's jargon and inner workings, in order to bring more people into the discussion. Drawing parallels to the early days of the internet (how its inner workings were similarly abstruse) could help illustrate the potential for broader public understanding and engagement and allow those with relevant subject expertise to engage in this largely tech-focused space.

- **Promote human-centered design and ethical considerations:** Many participants discussed the need for the integration of human values, ethics, and social impact considerations into all stages of AI development and deployment. We should support organizations and initiatives that prioritize human well-being and social good alongside technological advancement.

- **Engage in critical conversations and resist hype:** There is a continued need for further thoughtful conversations about the potential risks and benefits of AI, and participants frequently discussed getting swept up in hype cycles driven by marketing, media, and Silicon Valley sensationalism. We should advance more workshops to advance critical conversations and challenge dominant narratives.

CONCLUSION AND NEXT STEPS

We're writing the conclusion to this report several months after the gathering in Hawaii. The main characteristic of this AI moment is its relentless velocity and capacity for change. Even as our insights from our discussions continue to hold up today, the landscape has already shifted. Conversations about bias, hallucinations, and adoption have given way to concerns about agentic services, inefficiency and energy consumption, and the surprise emergence of a $5M open weight LLM from China that can hold its own against OpenAI.

The main thing we know is that whatever we think we know today may not hold true for long. Anyone who says they know for certain what is to come is ignorant, lying, selling something, or all three. But in this moment, we have an opportunity to shift course—to fill in some of the current gaps with values, ideals, and a drive to protect the public good. If we seize this opening, we might be able to drive progress in an equitable direction.

It's no surprise that all our final recommendations call for more discussion. At times, we felt like the group could have spent another week, or even another month, deep in conversation. The problems are societal in scale, and the depth of what we don't yet know so vastly exceeds what we do at this nascent, experimental, janky, hype-filled, inefficient, expensive, and world-changing point in the arc of both invention and adoption.

What we all do next will determine the kind of future we'll share together.

ACKNOWLEDGEMENTS

PARTNERS

We are grateful to Omidyar Network for leading the charge and supporting this program, especially Tui Shaub. Many thanks to Sam Gill and the Doris Duke Foundation, who provided us use of their Shangri La global center in Hawaii, which served as an inspiration for our discussion. Additional support was provided by the Patrick J. McGovern Foundation, the Rockefeller Foundation, the Freedom Together Foundation, the John S. and James L. Knight Foundation, the Sloan Foundation, and the Mellon Foundation.

ASPEN DIGITAL

Ryan Merkley guided our discussion and was the principal author of this report. Special thanks, as well, to Isabella Sarmiento for leading the operations of the programming. Aspen Digital is led by Vivian Schiller, who also participated in the discussion and the crafting of the report.

ASPEN DIGITAL
aspen institute