

by: Robert Lempert

Robert Lempert¹

Artificial intelligence (AI) augurs changes in society at least as large as those of the industrial revolution. But much of the policy debate seems narrow - extrapolating current trends and asking how we might manage their rough edges. This essay instead explores how AI might be used to enable fundamentally different future worlds and how one such future might be enabled by AI algorithms with different goals and functions than those most common today. ²

AI machine learning algorithms now meet or surpass capabilities once regarded as uniquely human and will grow more capable over time. The algorithms draw their power from their ability to learn from vast amounts of data and their own experiences. For instance, autonomous vehicles, while still imperfect, roam our roads each learning not only from its own encounters, but also from the data gathered by other such vehicles. Amazon and other online services, learning from the choices of millions of customers, recommend books we might read and music we might enjoy. In 2016, the machine learning AI program AlphaGo beat Lee Sedol, the world's best human Go player, in a five match tournament, exhibiting strategies that displayed an astounding level of non-human creativity (Krieg, Proudfoot, and Rosen 2017). The AlphaGo version that beat Sodol learned its craft from records of previous human Go games. AlphaGo's even more powerful successor learned from millions of games it played against itself.

Such capabilities portend vast social, economic, and political transformations. Alluringly, we can imagine a future world of vast material wealth and convenience. Freed from human error and shortcomings, services such as transportation and medicine become much safer and more efficient. Localized, on demand, customized

by: Robert Lempert

3D manufacturing satisfies human wants and reduces environmental footprints. Freed from the drudgery of work that the machines can better handle, people embrace more meaningful tasks.

Yet AI also has dystopian portents. Most concretely, the technology threatens to destroy vast numbers of jobs. By some reckonings, 40 percent of the world's jobs could be replaced in the next 15 years (Roose 2019). AI may threaten privacy, as exemplified by the digital assistant devices that increasingly serve, and monitor, us in our homes. The power and wealth flowing to the individuals and firms that successfully commercialize AI may also exacerbate the income inequality straining our society.

The policy debate has begun to engage with these challenges (Alden and Taylor-Kale 2018). Some discussions focus on skill training to allow people to take new jobs. Others suggest a universal basic income (UBI) to reduce the adverse consequences of unemployment. Social insurance programs predicated on full-time employment might be expanded to include those in part-time jobs.

But such policy responses only nibble around the edges of an unfolding societal transformation. Extrapolating current trends, we might imagine a world in 2050 with AI offering unparalleled convenience and widespread material comfort while making most people economically irrelevant and concentrating power and wealth in the small number of firms and individuals that create and own the machines. No amount of job training or UBI would make such a world remotely similar to our own.

Envisioning scenarios offers one means to grapple with societal changes as fundamental as those augured by AI. Scenarios represent "focused descriptions of fundamentally different futures, often presented in a coherent script-like or narrative fashion" (Schoemaker 1993), often crafted to inform decision making. In

by: Robert Lempert

this spirit, we might envision a scenario called *Bezos World* in which current trends unfold into a world with all the wealth, power, and robots concentrated into a few hands. We can then posit other scenarios to help explore the extent to which *Bezos World* is foreordained, or whether we might imagine, understand, and influence the creating of alternative worlds that some people might find more to their liking.

The *Levelers* Scenario

Among its profound implications, the industrial revolution of the 19th and early 20th centuries reshaped the way citizens of liberal societies experienced agency and freedom. For many centuries, proponents of unfettered markets had also been natural advocates for popular sovereignty and economic and social equality (Anderson 2017). Markets helped to shatter aristocratic hierarchies and government chartered monopolies. In a feudal world, artisans labored for lords in a relationship of deference. A market economy made the two more equal, since an artisan could choose his customers with a freedom similar to that of a lord choosing his vendors.

As late as the mid-19th century, proponents of freedom such as Abraham Lincoln could envision a world of independent proprietors, whose few employees would be young apprentices on their way to running their own small firms.

The industrial revolution sundered this connection by creating economies of scale that required vast enterprises to exploit. These economies of scale generated immense material wealth, vastly expanding people's freedom as consumers. The choices of food to eat, clothes to wear, places to travel, and (with electric lighting) when to be awake and active expanded greatly. But the enterprises that swelled abundance and choice also generated new hierarchies of power and wealth. In the new managerial capitalism, assembly line workers and office clerks labored for their

by: Robert Lempert

bosses in a relationship of deference. Industrial age production also required large agglomerations of capital, further expanding inequality among citizens. In response, the U.S. and other industrialized countries transformed their governments and new non-governmental organizations emerged, such as labor unions, to more equitably share wealth and political power. The resulting social contract generated history's greatest expansion of freedom and material wealth.

But this social contract is now unraveling. As one important driver, new information technology threatens many jobs, and firms increasingly use technology to replace managerial capitalism with a new economy much less reliant on full-time workers.

For instance, organizations such as Uber foster a gig-economy in which firms contract with independent workers for short-term engagements, thereby capturing more of the profits and control of the workplace for themselves. Uber has roughly 10,000 full time employees and in the US 750,000 drivers operating as independent proprietors. While these drivers often enjoy the benefits of setting their own schedules, they exercise little control over their working conditions, have few benefits of employment, and exist at the bottom of an economic hierarchy which is investing heavily in automation to eliminate their jobs.

AI seems poised to exacerbate such trends. Today's firms actively pursue the *Bezos World* scenario, deploying AI towards a future in which machines do all the work and technology is used to maximize power and wealth in the hands of the small number of people who own the machines (O'Reilly 2017).

We might, however, envision a purposely very different scenario. Rather than using technology to automate away workers, this scenario envisions a world in which society uses technology to unwind the firm. We might call this scenario *Levelers*, based loosely on a political movement of that name during the English Civil War,

by: Robert Lempert

which was among the first to advocate markets as a means towards equality and popular sovereignty (Anderson 2017). In this *Levelers* scenario, a combination of populist uprising and appropriate technology would establish a 21st century version of what Lincoln envisioned for his time - a society of prosperous, independent proprietors without large concentrations of wealth - a world with many Uber drivers, but no Ubers. In this *Levelers* scenario, technology enables new kinds of work and is designed to spread power and wealth more equally across society.

Economists point to transaction costs to explain the necessity of firms (Coase 1960). Many economic arrangements, for instance running a railroad (Chandler 1977) or the design and production of complex goods such as automobiles and airplanes, require webs of connections among people and capital stock too complex to be organized and managed entirely by market forces. Prior to AI, these productive assets needed to be organized within managerial hierarchies associated with large stocks of capital. Today, firms also create vast wealth from the symbiotic relationship among their centrally managed, proprietary data (often gathered from customers), the services they sell that enables them to collect the data, and the AI learning from that data in order to make the services more efficient and effective.

AI technology can reduce many transaction costs, which enables a reimagining of the firm. Today, however, these capabilities are being used by firms to reimagine themselves for their own ends. For instance, ride-sharing services such as Uber and Lyft use ubiquitous web connectivity, databases, and route optimization algorithms to shift most of their workforce to independent proprietors and then to machines.

In contrast, the *Levelers* Scenario envisions a future in which widely deployed AI provides great wealth and convenience, but with the power relationships of the *Bezos World* scenario reversed. In the *Levelers* Scenario AI helps labor hire capital,

by: Robert Lempert

instead of the other way around. For instance, the firm Gigster currently uses AI to help its corporate clients efficiently identify ideal teams of temporary workers, thereby reducing the need to nurture well-known teams of full-time staff.³ In the *Levelers* Scenario, small teams of people might join together to make a car. Using a Gigster-like capability, they could find others with the skills they need and rent 3D manufacturing facilities to produce their design. In the service sector, AI could help drivers and passengers use competing databases to find each other without the need for an Uber to own the network. Achieving or maintaining the *Levelers* Scenario might require that a strong government use its authority to break up and prevent any large concentrations of economic power, such as those arising from data or network monopolies. But supported by AI created for the task, the *Levelers* Scenario envisions by 2050 a gig economy of radical social and economic equality, a world with material abundance, choice, and convenience but without managerial hierarchies and without large concentrations of wealth and power.

How AI algorithms might enhance human agency

This essay offers the *Levelers* as a normative scenario, which describes one way in which the transformative power of AI might lead to a future that enhances rather than reduces human agency. The scenario raises numerous questions, not least of which include the mix of government policies and social conditions necessary to bring such a future into being, the policies and conditions required to make such a future stable against the pressures that might undermine it, as well as the many potential unintended consequences that a world without many large commercial institutions might entail. But this essay explores one particular issue – how AI algorithms with different goals and functions than those most common today might

by: Robert Lempert

help enable the *Levelers* Scenario.

In thinking about normative future scenarios, a first step requires being explicit about what values constitute the good. We focus here on the capabilities framework, which Amartya Sen (2009) developed and then used to support a framework for measuring and helping to bring about a more just world. The capabilities framework offers an alternative to value systems based entirely on welfare economics. Outcomes matter, but so does the process by which the outcomes are achieved. For instance, choosing to spend a quiet day at home is different than spending a day under house arrest though the physical outcomes may appear similar. A just society, as emphasized by the capabilities framework, enables individuals to make reasoned and consequential choices about their own lives, to act on those choices, and to evaluate the results in terms of their own values and goals, respecting the diversity of such goals and values. In addition, a just society enables individuals to participate meaningfully in shaping their society, sharing in an informed and consequential way in choices made about its economic, social, and political attributes.

The *Levelers* scenario envisions widespread use of AI algorithms that *enhance*, rather than *reduce*, such human agency. But many of today's most touted algorithms fall in the latter category. They seek some best outcome and are less concerned with enriching the decision-making processes of the humans with whom they interact. In operating on the world, such algorithms start with a clear set of objectives to achieve and a set of actions they can take. They then assess a current situation, predict the consequences of various actions, and choose those actions that best achieve the objectives. For instance, AlphaGo seeks to win a Go game, can make moves consistent with the game's rules, and each turn makes the move that most increases its chances of victory. Autonomous vehicles seek to travel safely and

by: Robert Lempert

efficiently to a desired destination; create an understanding of the objects around them; and each moment decide whether their objectives are best met by turning, accelerating, or braking. The Amazon website aims to display those products to viewers that result in the highest probability of a sale to a satisfied customer. While humans may set the objectives (e.g. the destination of the AV or a high sales volume), the algorithm aims to reduce human involvement in the choices that lead to those outcomes. The algorithm's success is judged by the extent to which it enhances human welfare, based on the outcomes for at least some humans.

What would it look like for an algorithm to enhance human agency? The algorithm would certainly help achieve welfare-improving outcomes. But the algorithm would also help individuals gather the resources - including skills, capital, and material resources - relevant to their goals, team with appropriate human collaborators, help all the individuals involved make good choices on how best to coordinate their activities and deploy their resources in pursuit of multiple goals, explore how these goals might be expanded or modified in light of what is possible, examine all these steps from different vantages and points of view, and explain their choices to themselves and to others. For instance, a Gigster-like algorithm that helps people to form teams could be designed, not to promote just efficiency and profit by adding skills to an existing organization, but to help people assemble new teams and networks around a common purpose, help them to assemble and manage the resources needed to achieve their objectives, and help the participants' understanding of their common and individual purposes grow and deepen as they worked together.

Algorithms do exist that help to enhance humans' ability to make reasoned and consequential choices about their lives. This is most apparent in the field of decision support, in which algorithms, much less capable than today's AI, are used to help

by: Robert Lempert

groups of stakeholders deliberate about contentious policy challenges such as improving the resilience of a community to climate change and then seek consensus on actions to address the challenge (Marchau et al. 2019). Such algorithms are inherently multi-scenario and multi-objective, the former to reflect multiple ways to view and interpret the world, the latter to reflect alternative ways of judging outcomes due to different interests and different ethical frameworks (Lempert, Groves, and Fischbach 2013). The algorithms, often adapted from those in the classification and robust optimization literatures, are configured to support what is called *agree-on-decision* analysis, because the analytics aims to help people with different objectives and expectations about the future nonetheless reach consensus on near-term actions.⁴ In contrast, algorithms are more commonly developed and used to support *predict-then-act* analysis, which assumes that all the parties to a decision will accept a single, often computer-generated understanding of the future and then seek prescriptive recommendations from the computer on the best actions to take.

Such *agree-on-decision* algorithms are often embedded in a process called “deliberation with analysis” in which stakeholders deliberate on their objectives, options, and problem framings; algorithms generate decision-relevant information; and the parties revisit their objectives, options, and problem framings influenced by the algorithms information products (NRC 2009). The process envisions that participants’ understanding and views will evolve over time in response to interactions with each other and with scientific information. In brief, the algorithms aim to facilitate a democratic process of social choice, in which diverse parties agree on actions based on a Habermasian discourse in which parties recognize the inescapable plurality of competing views; facilitate clear explication of reasoning and logic; and accept the legitimacy of multiple views. The process, and the algorithms

by: Robert Lempert

that support them, are judged not only by the welfare outcomes they help achieve, but also by the extent to which they empower the people involved in the process to make what they regard as meaningful choices about their lives and society.

Human agency-enhancing AI algorithms are clearly not on their own sufficient to bring about a *Levelers* scenario. But it nonetheless remains useful to ask to what extent such algorithms could be designed to facilitate such a future? It remains hard to know, in part because the requisite technology landscape remains under-explored. Many people and institutions currently developing AI have incentives to reduce human agency, since excluding humans best serves the purposes of those most involved in designing the algorithms. Firms' goals generally do not focus on the agency of their customers or workers. Rather firms seek control over the latter and want their customers to make choices good for the corporation, not engage in self-reflection and enlightenment. Researchers seek objective truth as scientists and technological virtuosity as engineers, so seek algorithms that operate independently from human subjectivity and influence.

The experience with multi-scenario, multi-objective decision support may be, however, instructive. Explicitly changing their goals from *predict-then-act* to *agreement-decision* analysis enabled researchers and institutions focused on the latter to recraft existing classification and robust optimization algorithms, originally developed to operate independently of humans, and develop interlocking processes and algorithmic tools that enhance human agency. While at best a necessary condition, and certainly not a sufficient one, decision support processes using such agency-enhancing algorithms do seem to help generate improved social outcomes (Knopman and Lempert 2016).

by: Robert Lempert

Moving towards a *Levelers* scenario?

If those seeking alternatives to the *Bezos World* scenario wanted to take near-term actions that might steer society towards a *Levelers* scenario, what might they do?

As one step, they might establish research activities and institutions that have the incentives and resources to develop human agency enhancing AI algorithms. People might then launch pilot programs in a few sectors of the economy in which such algorithms might have the most success in helping people replace the firm rather than the firm replacing workers. The government could institute policies that create space in the economy for such experiments to thrive, perhaps akin to the renewable portfolio standards in the energy sector.

The industrial revolution and its enabling technologies created vast material abundance but exacerbated tensions among the several dimensions of human freedom. Today's AI augurs social transformations at least as profound. Much AI research, development, and deployment currently seeks to replace humans in pursuit of technological virtuosity and economic efficiency. The *Levelers*, and speculation regarding the algorithms that might support it, are offered as one of many potential scenarios intended to help people systematically explore whether and how AI might be configured to facilitate a future in which machines collaborate with humans to enhance the latter's capabilities, agency, and freedom.

REFERENCES

Alden, Edward, and Laura Taylor-Kale. 2018. "The Work Ahead: Machines, Skills, and U.S. Leadership in the Twenty-First Century." In, New York: Council on Foreign

by: Robert Lempert

Relations.

Anderson, Elizabeth. 2017. *Private Government: How Employers Rule Our Lives (and Why We Don't Talk about It)* (Princeton University Press).

Chandler, Alfred Dupont. 1977. *The Visible Hand: The Managerial Revolution in American Business* (Harvard University Press).

Coase, Ronald H. 1960. "The Problem of Social Cost." In, *Classic Papers in Natural Resource Economics* (Palgrave Macmillan: London).

Knopman, Debra, and Robert Lempert. 2016. Urban Responses to Climate Change: Framework for Decisionmaking and Supporting IndicatorsRAND. RR-1144-MCF

Krieg, Gary, Kevin Proudfoot, and Josh Rosen. 2017. "AlphaGo." In, RO*CO FILMS.

Lempert, Robert, David G. Groves, and Jordan Fischbach. 2013. "Is it Ethical to Use a Single Probability Density Function?" In, RAND.

Lempert, Robert J., Steven W. Popper, and Steven C. Bankes. 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-term Policy Analysis* (RAND Corporation).

Marchau, Vincent A. W. J., Warren E. Walker, Pieter J. T. M. Bloemen, and Steven W. Popper. 2019. *Decision Making Under Deep Uncertainty: From Theory to Practice* (Springer) (in press).

National Research Council. 2009. "Informing Decisions in a Changing Climate." In, Panel on Strategies and Methods for Climate-Related Decision Support, Committee on the Human Dimensions of Climate Change, Division of Behavioral and Social Sciences and Education (The National Academy Press).

by: Robert Lempert

O'Reilly, Tim. 2017. *WTF?: What's the Future and Why It's Up to Us*. (Random House).

Pendleton-Julian, Ann, and Robert Lempert. 2019. "World Building Workshop on Technology and Work Worth Doing (Not Jobs) Post AlphaGo." In, RAND CF-398 (forthcoming).

Roose, Kevin. 2019. 'The Hidden Automation Agenda of the Davos Elite', *New York Times*, Jan 29, 2019.

Schoemaker, Paul J.H. 1993. "Multiple Scenario Development: Its conceptual and Behavioral Foundation". In, *Strategic Management Journal*, 14: 193-213.

Sen, Amartya. 2009. *The Idea of Justice* (Belknap Press).

Shapiro, Ian. 2003. *Moral Foundation of Politics* (Yale University Press).

Wu, Tim. 2018. *The Curse of Bigness: Antitrust in the New Gilded Age* (Columbia Global Reports).

1. Director, RAND Pardee Center for Longer Range Global Policy and the Future Human Condition
2. This essay draws heavily on several sources, including the work of Anderson (2017), O'Reilly (2017), Sen (2009); a RAND workshop on related issues Pendleton-Julian and Lempert (2019); and the author's experience with multi-objective, multi-scenario algorithms designed to support deliberation with analysis among humans (Lempert, Popper, and Bankes 2003).
3. See Gigster. Accessed February 6, 2019. <https://gigster.com>.
4. This essay is not the place to speculate on the specific government policies required to enable a Levelers scenario. But an aggressive anti-trust policy might

by: Robert Lempert

prove part of the mix. It is thus worth noting that the original goals of anti-trust policy, for instance as expressed by Louis Brandeis, have strong resonance with the view of a just society that emerges from a capabilities framework. In past decades, the U.S. government has done little to combat growing concentrations of power, in part because its anti-trust policy has focused on a narrow goal of consume welfare (Wu 2018). In contrast, Brandeis favored breaking up monopolies and restoring an economy of competitive markets among many small entities in order to free people from the arbitrary will of the monopolist.