January 2022

Machines, Bureaucracies, and Markets as Artificial Intelligences

CSET Issue Brief

CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

AUTHOR

Richard Danzig

Executive Summary

Markets, bureaucracies, and machines are inventions designed to process information at speeds, in quantities, and with accuracies that surpass human capabilities. In all three systems this information processing is made possible by reducing reality to narrow inputs (e.g., bits, prices, entries on bureaucratic forms) and then detecting patterns and pattern conformance from these inputs. Recognizing this commonality, this paper treats these systems as members of a set of artificial intelligences and uses the experience with markets and bureaucracies to suggest descriptive, predictive, and prescriptive insights about machine intelligence.

The resulting insights complement and sometimes modify observations derived by comparing machine intelligence with the "natural" intelligence of the human mind. They normalize machine intelligence—it is not as unique as the human-machine contrasts alone suggests. Moreover, considering markets, bureaucracies, and machines together highlights that many concerns about machine intelligence are about how that intelligence amplifies the powers of bureaucracies and markets. Finally, this approach illuminates ways in which machines, bureaucracies, and markets function as an ecosystem, exchanging data, co-evolving, posing analogous challenges of control, and intensifying and complicating requirements for regulation.

Section 1 describes inadequacies in prevalent definitions of intelligence, "artificiality," and the composite concept of "artificial intelligence," then proposes that a different approach to these concepts can broaden and deepen understanding of machine intelligence. It describes how machines, bureaucracies, and markets can usefully be regarded as a set of artificial intelligences invented to complement the limited abilities of individual human minds to discern patterns in large amounts of data. This section identifies the foundational reductionist and correlative character of these three systems.

Section 2 applies this perspective to show that the "artificiality," and related characteristics often described as singularly disquieting and alienating attributes of machine intelligence, are understood

better when considered alongside the histories of bureaucracies and markets. When bureaucracies and markets were introduced, they too were seen as artificial and alienating. The fact that over time these concerns dissipated suggests that artificiality is predominantly a marker for unfamiliarity and not likely to endure as a concern about machine intelligence.

Section 3 addresses the widespread concerns about machine opacity and unpredictability. While noting that stock markets pursue transparency through methods such as standardized accounting principles and required disclosures, the discussion points to widespread acceptance of opacity and unpredictability in these markets. It observes that if markets were predictable, we would not need them. Moreover, in general, intelligence is an emergent attribute. If a system is not opaque, it is commonly described as "just computation." Drawing on human experiences with markets, bureaucracies, and machines, this section distinguishes four contexts in which opacity ranges from completely acceptable to unacceptable and suggests that the design of machine intelligence and policies to control that intelligence should be focused accordingly.

Section 4 examines dependency on data as a shared attribute of machines, markets, and bureaucracies. Though the shibboleth that "data is the new oil" is commonly recited with respect to machines, data is also the lifeblood of bureaucracies and markets. All three systems depend on obtaining, processing, and thwarting distortion (or "poisoning") of information. However, their most characteristic modes of obtaining information differ: bureaucracies have historically secured information by command, markets by seduction, and machine systems by simulation or scavenging. As a result, complex patterns have emerged in which these systems share and compete for data. These patterns will change as simulation, the construction of virtual worlds, and the flow of information from the Internet of Things become more significant as sources of data for machines. Understanding this and other likely developments is helped by seeing that machines, bureaucracies, and markets function as an ecosystem, each co-evolving with the others.

Section 5 explores the nature of this ecosystem and shows how, because of their commonalities, the commingling of markets, bureaucracies, and machines amplifies characteristics in each. In addition, this commingling gives rise to hybrids such as Uber that combine these intelligences in new ways. The discussion observes that, when introduced, all three systems were often described as reliable, value-free processes for efficient allocation and decision-making. However, experience with markets and bureaucracies shows that as failures occur and values underpinning these systems are revealed and challenged, efforts at regulation result. This section argues that this history is a likely harbinger of crisis and regulation of machine intelligence. However, while bureaucracies and markets had no competition as artificial intelligences at the time of their introduction, the future of machine intelligence will be affected by the pre-existing power of bureaucracies and markets. In fact, the most immediate and important concerns about machine intelligence will not be about machine intelligence, but rather about how bureaucracies and markets utilize machine intelligence and then about how society regulates the three systems together.

In sum, this paper argues that the presently dominant tendency to compare machine intelligence to the intelligence of the human mind narrows and distorts understanding. Present understanding should be complemented by consideration of bureaucracies and markets. Widening the aperture of comparison opens an array of possibilities for present insight and future investigation.

Table of Contents

Executive Summary 1
Introduction 5
Section 1: Machines, Bureaucracies and Markets as Reductionist, Correlating Intelligences 8
- The seductive path of defining intelligence by reference to an individual human mind 8
- Machines, markets, and bureaucracies as intelligent systems 10
- The reductionist and correlative characteristics of these artificial intelligences. 11

Section 2: Artificiality & Dehumanization 16
Section 3: Opacity 19
- Opacity as a Characteristic of the Intelligence that Emerges from Correlating Systems 19
- Mapping a Spectrum of Concern About Opacity in Bureaucracies, Markets and Machines 20
- Subjects 21
- Users 22
- Partners 23
- Overseers 25

Section 4: Bureaucracies, Markets, and Intelligent Machines as Data Dependent and Interdependent Systems 28
- Data as a Source of Power for Bureaucracies and Markets 28
- Machines Require Data and Amplify the Power of Bureaucracies and Markets to Manipulate Data 29
- The Challenge of Sharing Data 30

Section 5: The AI Ecosystem and Its Regulation 35
- The Two-Body System 35
- Adjustments caused by machine intelligence as a newcomer to this ecosystem 37
- Operational and Regulatory Controls 38

Conclusion 41
Author 44
Acknowledgments 44
Endnotes 47

Introduction

The power and proliferation of machine intelligence creates evident problems: How will humans assimilate and control this technology? This paper takes a fresh approach to these questions by viewing machine intelligence as a member of a set of artificial intelligences that includes bureaucracies and markets. This context prompts insights different from those that arise within a dominant paradigm that focuses on comparing machine intelligence (viewed as "artificial") with individual human intelligence (implicitly thought to be "natural").

To demonstrate the value of this approach, succeeding sections reassess present debates about concepts of artificial intelligence, machine opacity, and data dependence. This paper then analyzes the interactions between bureaucracies, markets, and machines, treating these systems as an ecosystem in which each affects the evolution of the other. It concludes that the challenge for societies is not to manage machine intelligence in isolation, but rather to manage it as a newcomer that will find a place in an extensive system of artificial intelligences on which we rely, in whole or in part, for consequential decisions.

As a starting point for this discussion, it should be recognized that intelligence is not precisely defined or understood. Rather, "intelligence" is a description often applied to entities that display a cluster of attributes including an opaque ability to receive information, process it, and provide "decisions" or information outputs of higher value.[1] Machines, bureaucracies, and markets can be described this way. Though bureaucracies and markets are largely absent from the literature on artificial intelligence, individual human intelligence should not be the sole standard for assessing machine intelligence: societally important manifestations of human intelligence are often collective.[2]

Viewing machine intelligence in conjunction with bureaucracies and markets normalizes machine intelligence. Several characteristics of concern—for example, opacity, brittleness, lack of common sense, a blinkered focus on some but not all relevant variables, and risks of loss of control—are not singular to new systems of machine intelligence, but rather familiar attributes of bureaucracies and markets. These problems arise in all three systems because members of this set of intelligences reduce complex data to standardized inputs and then identify patterns without comprehending contexts or distortions that shape these inputs. Markets, bureaucracies, and intelligent machines compensate for the limited data processing and computational abilities of a human mind operating alone. But, at the same time, their shared foundational characteristics render these systems rigid, vulnerable, and alienating to many humans they affect.

Differences in the methods these systems use for obtaining data further determine their powers and limitations. Intelligence is not synonymous with omniscience. In particular, valid inferences about human activities depend on constantly refreshed data. The three systems obtain that data in many ways, but they differ in their most characteristic means: markets seduce, bureaucracies command, and machines scavenge. As a result, the systems share data and this is an important means of empowering or constraining them.

Considering these artificial intelligences as a group yields prescriptive insights. Each of these systems is conceived (and commonly defended) as a neutral means of classification and allocational efficiency. However, the operations of markets and bureaucracies have frequently triggered challenges to the value premises, risk tolerances, and sensitivity to externalities embedded (or not embedded) within them. Concern about these variables repeatedly has provoked interventions by government agencies and the employment of auditors, inspectors, judges, and others whose duties include reviewing and constraining these systems. It is likely that efforts to control machine systems will follow a similar path.

Plans for controlling machine intelligence are naturally predisposed to employ procedures like those designed for controlling other machines. These emphasize careful design and then conditioning acceptance on test, evaluation, verification, and validation. The perspective advanced here, however, suggests that the problem of controlling evolving intelligent machines may be better understood as analogous to controlling lower-level bureaucrats. After bureaucratic actors are trained, tested, and certified, they are constantly reevaluated through systems of probation, supervision, audit, and promotion to widened responsibilities or removal. Analogous systems of continuous reassessment must be developed for machines whose intelligence will evolve and will be employed in situations that cannot be completely predicted.

The intrinsic difficulty of establishing a system for controlling machine intelligence will be amplified by problems of pace. Mechanisms for control of bureaucracies and markets co-evolved with the development of these systems over centuries. The speed of proliferation, the speed of operation, the diversity, and the power of machine intelligence will afford less time for reflection and reaction.

Simultaneously, information retrieval and processing power of intelligent machines will alter the procedures and powers of bureaucracies and markets. To manage the power and proliferation of machine intelligence, the triad of machine learning, bureaucracies, and markets must be viewed as an ecosystem of artificial intelligences with areas of competition, complementarity, hybrid combinations that bring together different forms of intelligence, and interactive effects that will transform each component.

Section 1: Machines, Bureaucracies and Markets as Reductionist, Correlating Intelligences

> [W]e see a complicated network of similarities overlapping and criss-crossing .... I can think of no better expression to characterize these similarities than 'family resemblances'; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc. etc. overlap and criss-cross in the same way. – And I shall say: "games" form a family."
>
> -- Ludwig Wittgenstein, *Philosophical Investigations* (1953)[3]

The seductive path of defining intelligence by reference to an individual human mind

The compound concept of "artificial intelligence" is difficult to operationalize or even comprehend. A preeminent analyst of intelligence in humans has decried "endless muddling about the definition" of the concept.[4] As described in the next section, characterizations of things we call "artificial" are not what they were some years ago or what they will be some years from now.

A foundational problem is that discussions of intelligence explore something only partially understood. This is not unusual—for example, concepts of "life" or "consciousness" have eluded widely accepted, precise definitions.[5] The ascription of intelligence to someone or something is not an assertion of an objectively determinable fact. Intelligence, like beauty, is in the eye of the beholder. Lakoff and Johnson call terms like these "ontological metaphors." They emphasize that how we conceive and construct a term will shape our perception. Chosen attributes "allow us to comprehend...[but] will necessarily hide other aspects of the concept."[6] Definitions, accordingly, may be multiple and may be employed or put aside according to their usefulness for improving understanding.[7]

What do prevalent definitions of artificial intelligence illuminate and what do they hide? The most common approach is defining intelligence as what human beings do when they reason cognitively. This conceptual strategy distinguishes between human (or human and animal) activities and the activities of machines—calling the latter "artificial." It then leans the concepts of artificiality and intelligence against one another, apparently in the hope that combining one confusing concept with another confusing concept will result in a clear concept. For example, the U.S. Department of Defense defines artificial intelligence as "the ability of machines to perform tasks that normally require human intelligence."[8] Similarly, a highly regarded professional publication on the "State of AI" defines AI as:

> A broad discipline with the goal of creating intelligent machines, as opposed to the natural intelligence that is demonstrated by humans and animals. It has become a somewhat catch all term that nonetheless captures the long-term ambition of the field to build machines that emulate and then exceed the full range of human cognition.[9]

The Turing test, which suggested assessing machines by asking whether they could produce outputs indistinguishable from those of humans, may have encouraged an inclination towards this way of defining artificial intelligence.[10] This approach was reinforced by theories that computer reasoning could apply networking techniques like those that integrated neurons in the human brain.

Assessing machine intelligence by comparison with an individual human mind has some advantages. Sometimes it is desirable to evaluate the pros and cons of machine intelligence as an alternative or complement to human reasoning—such as when a weapon may be employed with or without human intervention. Moreover, if we must come to grips with the fuzziness of the concept of intelligence when we grapple with the concept of "artificial intelligence," then it is attractive to at least define artificiality as an unvarying bright line distinguishing carbon-based systems as natural systems and silicon-based systems as artificial.

Machines, markets, and bureaucracies as intelligent systems

Despite its attractions, a binocular perspective that contrasts machine and individual human intelligence disposes us to overlook entities like bureaucracies and markets.[11] These are not individual intelligences, and they are not adequately described as carbon or silicon. But, both functionally and psychologically, bureaucracies and markets are routinely treated as intelligent machines. Functionally, humans value the decisions these systems autonomously recommend or implement[12] after processing complex or noisy information.[13] Psychologically, individuals who encounter these systems often feel that their operations are opaque. Their components and processes may be quite visible, but the reasons for particular outputs are obscure. In other words, these systems have attributes that are emergent.[14] This is important because a system that is readily comprehended and predicted is commonly regarded as "just computation," not intelligence.[15]

Thus, for example, the response to Deep Learning's triumph over a top-ranked human in the game of Go was not just a consequence of its triumph over the reigning human champion.[16] It resulted from a machine decision that humans could not have foreseen or explained. "Somehow a computer program knew something about the game that we didn't. Somehow it's [sic] intuition was both different and better than human intuition."[17] In the words of another commentator, "It was a move that demonstrated the mysterious power of modern artificial intelligence…"[18]

Similar emergent and unpredictable qualities are evident in markets. It is not too strong to say that we value markets as tools to produce emergent effects.[19] If the workings of markets were predictable, we wouldn't need them.[20] Probably no subject in the history of mankind has produced so much effort, by so many, with so little success as efforts to predict stock markets. Even retrospectively, daily newspaper columns bear witness to diverse and inconsistent efforts to explain the previous day's market movements.

This emergent quality is constrained in bureaucracies by modern regulations intended to force transparency (for example, in the United States by the Administrative Procedure Act). But many have observed and experienced bureaucracies (particularly in less regulated situations, like those responsible for security clearances or immigration approvals) as black holes whose operations are beyond comprehension. Franz Kafka offers classic presentations of this view.[21]

Societies' acceptance of bureaucracies and markets as mechanisms for consequential and often emergent decisions suggests that we may learn something about the new phenomenon of machine intelligence by considering it in light of our extended, substantial experience with markets and bureaucracies. Any such comparison immediately confronts the difficulty that markets, bureaucracies, and machine intelligences are diverse and numerous.[22] This paper's strategy for coping with this expansive space is to adopt a very specific exemplar of machine intelligence, Deep Learning,[23] as the anvil against which to hammer out observations. That choice can be contested, but many commentators recognize that Deep Learning, a form of machine learning, is now the most significant form of machine intelligence and accounts for most of the recent progress in the field.[24]

At the same time, this discussion is not confined to particular markets or bureaucracies, but uses different examples as seems most useful to increase insight about machines. The aim is not to increase insight about bureaucracies and markets. Though some incidental benefit may be gained in that respect, the discussion of these systems makes no effort to be comprehensive or representative. This paper suggests generalizations and points to examples as analogies, subject no doubt to qualification, but valuable to the degree they afford insight about machines.

The reductionist and correlative characteristics of these artificial intelligences

Close examination of markets, bureaucracies, and Deep Learning reveals two common foundational characteristics: they are correlative and reductionist. Correlating intelligence identifies patterns. For example, when the barometer goes down, rain follows in a calculably high percentage of cases. Similarly, when a disease appears less often in a vaccinated group than in a matched unvaccinated group, the vaccination may be inferred to be useful with a level of confidence linked to statistical power of the observed results. The correlation however reveals nothing about how or why the vaccine works or what causes rain and barometers to move synchronously. The often-recited proposition is correct: "Correlation is not causation."[25]

An intelligence that is "conceptual" claims insights as to causation. It propounds a theory about the linkage or lack of it between phenomena of interest. For example, an understanding that though barometric behavior is correlated with rainfall, the barometer does not cause rain. Human thought encompasses both correlative and conceptual capabilities.[26] To take a famous example, Johannes Kepler spent four years intensely trying to fit the orbits of planets to an identifiable form. Eventually, without comprehending why, he discovered that all such orbits could consistently be represented by an ellipse. Two-thirds of a century later Sir Isaac Newton developed a causal explanation that illuminated why and how this movement conformed to this geometry.

Markets are purely correlative entities. Assimilating many information inputs, they arrive at a price without any pretension of understanding what drives that price. Market mechanisms are indifferent to the character or consequence of the commodities they evaluate. They treat wages, automobiles, wheat, and Rembrandt's paintings with indifference. Markets, while immensely useful, are "mindless and purposeless."[27]

At their lower levels, bureaucracies, whether in corporations or government agencies, operate similarly. (The paper returns later to how higher level supervisors of these bureaucracies can add a layer of conceptual thought.) "[O]rganizational behavior, particularly decision-making, involves rule following more than the calculation of consequences."[28] Observing bureaucratic behavior in the formative years of the Prussian state, Max Weber summarized the situation: "The individual obeys the order, setting aside judgments either of its rationality or morality… ."[29]

Machine intelligence is similarly "rule-following." It abstains from "judgments of rationality or morality." It ignores consequences for which it is not programmed. And, as presently developed, it does not comprehend causation. Insight about causation remains an exclusively human power. In the words of one authority:

> In their 1955 proposal for the Dartmouth summer AI project, John McCarthy and colleagues wrote, "An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves." Now, nearly seven decades later, all of these research topics remain open and actively investigated in the AI community. While AI has made dramatic progress over the last decade in areas such as computer vision, natural language processing, and robotics, current AI systems almost entirely lack the ability to form humanlike concepts and abstractions.[30]

Two experts put the point unequivocally: "That's all that a deep-learning program can do: fit a function to data....Current machine-learning systems operate almost exclusively in a statistical, or model-blind, mode, which is analogous in many ways to fitting a function to a cloud of data points...."[31]

Or in the words of the authoritative textbook on the subject:

> Machine learning is essentially a form of applied statistics with increased emphasis on the use of computers to statistically estimate complicated functions and a decreased emphasis on proving confidence intervals around these functions....[32]

This is an extraordinarily powerful capability. Search engines, facial recognition, perception of changes in the environment, state surveillance and many other capabilities can be constructed from the ability to compare and correlate demographic, economic, physical, and other characteristics.[33] For example, deep learning acquired translation ability simply by calculating correlations between French, English, and German translations published by the European Union. “This approach dispenses with linguists: the programmers building these systems need not even speak the language they are working with.”[34] Going further, a few years later machines were programmed to achieve “generative” capabilities from these correlative capabilities. They are now able to create images, news articles and plausible essays on any topic.[35] Myriad activities and characteristics (demographic, economic, physical, etc.) can be correlated with one other and with individual identities.

These and many other cutting-edge achievements of deep learning warrant considerable study and discussion. But they root in a single capability that can readily be comprehended. As the head of the Innovation Information Office at the Defense Advanced Research Projects Agency (DARPA) put it, “[T]he big secret about neural nets is that they’re really just spreadsheets on steroids.”[36]

The common characteristic of correlating without conceptualizing—of identifying patterns and conformance with patterns without seeking, much less determining, causes of patterns—links the three intelligences discussed in this paper and notably distinguishes them from prevalent perceptions of human intelligence.[37] This contrast is amplified by a further fundamental characteristic differentiating these artificial intelligences from humans. There are limits to the diet of inputs that affect human perception and thinking,[38] but humans integrate inputs from a spectrum of senses and from memories and emotions. By contrast, each of the artificial intelligences assimilates data only in a limited format, according to narrowly defined rules. Reducing the rich variety of the world into digits, pixels, prices, and boxes checked on bureaucratic forms facilitates power, speed, and precision in machine, market, and bureaucratic operations. But while providing a standardized basis for correlative calculations, this reductionist method further narrows intelligences already blinkered from a lack of causative understanding.[39]

Markets, lower-level bureaucracies, and machine intelligences do not challenge the human monopoly on conceptual thinking. Instead, they shore up deficiencies in human correlative capabilities.[40] They are tools that correct for the limitations of an individual human being when assessing large amounts of data.[41]

Section 2: Artificiality & Dehumanization

If machines, like markets and bureaucracies, complement human capabilities, why does contemporary discussion so commonly consider them “artificial” and alien? This question is obscured, not clarified, by contrasts between carbon and silicon systems. It would be more, not less, disorienting and alienating if these intelligences were embedded in human bodies. The question is better addressed by considering it in the context of the history of bureaucracies and markets.

“Artificial” literally means not arising in nature. Put another way, it means it is not biological. That view, however, points in the wrong direction. The history of bureaucracies and markets suggests that the terms “artificiality” and “dehumanization” are often surrogates for unfamiliarity and disruption. In contrast with concerns about opacity, objections to these characteristics, intense during the period they are first encountered, diminish with time.

The artificiality and “dehumanizing” effects of bureaucracies and markets were keenly felt in their early days but are generally, matter-of-factly accepted in the twenty-first century. Hobbes referred to the state as “artificial man.”[42] Weber observed that bureaucracies replaced more personal forms of rule.

> … the more fully realized, the more bureaucracy “depersonalizes” itself, i.e., the more completely it succeeds in achieving the exclusion of love, hatred and every purely personal, especially irrational and incalculable, feeling from the execution of official tasks.[43]

Markets displaced personal relationships with impersonal transactions and stimulated the use of standardized, abstract units (like a meter or a quart) for measures that previously were more human (like “a foot” or a “heaping portion”).[44] The economic historian Robert Heilbroner traces changes that led to new concepts of land, labor, and capital as alienable abstract instruments of production. He captures the trauma of the transition:

> … the Middle Ages lacked the market; and lacking the market (despite its colorful local marts and traveling fairs), society ran by custom and tradition. The lords gave orders, and production waxed and waned accordingly. When no orders were given, life went on in its established groove…. [T]he change, long drawn out though it was, was not a peaceful evolution. It was an agonized convulsion of society, a revolution….
>
> The market system with its essential components of land, labor and capital was thus born in agony – an agony that began in the thirteenth century and had not run its course until well into the nineteenth. Never was a revolution less well-understood, less welcomed, less planned.[45]

This "agony" is important and has not entirely disappeared. But it is striking how it dissipates as the new intelligences are assimilated. Today the characteristics of bureaucracies and markets are hotly criticized and keenly contested. But the entities themselves are no longer decried as artificial. "Nowadays, nobody talks much about bureaucracy… one obvious reason is that we've just become accustomed to it. Bureaucracy has become the water in which we swim."[46] Similarly, "we do not often pause to reflect on [the] chains or webs" of market performance.[47]

This evolution underscores that humans are most inclined to perceive a system as "artificial" when it is new.[48] It reenforces Douglas Adams' three "rules that describe our reactions to technologies":

1. Anything that is in the world when you're born is normal and ordinary and is just a natural part of the way the world works.
2. Anything that's invented between when you're 15 and 35 is new and exciting and revolutionary and you can probably get a career in it.
3. Anything invented after you're 35 is against the natural order of things.[49]

As novelty diminishes and a technology recedes into the background, what was once regarded as artificial, unnatural, and dehumanizing comes to be accepted as familiar and empowering.[50] If readers of this paper look up from their computers or printed text (themselves "artificial"), they will find that except for people, pets, and plants, the environment around them is populated by things that are artificial. Our homes, roads, and clothes do not occur in nature,[51] but we have no anguish about their "artificiality" because we have been raised in environments where they are omni-present. Even amongst animals, we no longer ponder, as our ancestors must have, how dogs, "man's best friend," were bred from wolves or how unnatural writing seemed when first introduced into cultures that associated the acme of intelligence with oral communication and human memory.[52] Similarly, familiarity has bred acceptance of the automobile, nuclear weapons, the birth control pill, and "artificial light."[53]

This perspective should be comforting over the long term, but it is distressing considering the challenges for machine intelligence in the years immediately ahead. Machine intelligence is in its adolescence: It has transcended the impotence of infancy and has moved from toy problems to significant capabilities, but it has not yet been assimilated in the larger society. Powered by exponential expansions in computing capabilities and data sets and aided by improved and widely available algorithms,[54] intelligent machines are now likely to proliferate more rapidly and evolve more quickly than their predecessor bureaucratic and market systems. This stage in the evolution of machine intelligences is likely to be still more disorienting than it was for people to adjust to bureaucracies and markets—and those adjustments were, to use Heilbroner's word, agonizing.

In sum, the histories of bureaucracies and markets should sensitize to us the temporal aspects of judgments of artificiality and dehumanization. The histories of objections to bureaucracies and markets on grounds of artificiality bodes poorly for reactions to machine intelligence in the near future, but well for it once it is widely institutionalized. Opacity, however, is a more complex, enduring, and significant concern.

Section 3: Opacity

The "opacity," that is the lack of transparency, of intelligent machines is widely regretted. It is commonly viewed as contributing to the "inhumanity" of these machines, their propensities to disguise and propagate bias, their tendencies to "lack common sense," and to risks of their producing results that are not desired, expected, or controllable.[55] A comparison to markets and bureaucracies suggests, however, that human users of these systems have considerable tolerance for opacity. In fact, opacity is a characteristic of all intelligent systems and it is intrinsic to systems that are reductionist and reach results only by correlation. Starting with these observations, this section distinguishes positions along a spectrum in which, at one end opacity is readily accepted, at the other it is rejected, and in between it warrants reduction, not generally but for some specific and limited purposes.

Opacity as a Characteristic of the Intelligence that Emerges from Correlating Systems

When describing a program for "Explainable AI," the head of DARPA's Innovation Information Office observed that if a machine identifies an image as a cat, it cannot now well respond to the query "Why do you think it's a cat?"

> The system, if it could talk, would say, "Well, I did my calculations and at the end of those calculations, cat came out as highest probability." That's not very satisfactory. We'd much prefer the system to be able to respond to us, saying, "Well, it has ears, it has paws, it has fur, and it has these other features that lead me to conclude that it's a cat."[56]

But why is it much preferable to reduce opacity? The first section of this paper suggested that for many observers an opaque emergent quality is a prerequisite for calling something intelligent. If something is completely understood, it is "just computation." Furthermore, the attributes of everyday tools are frequently opaque, yet humans entrust the tools with their lives—such as when they drive automobiles. And opacity is not confined to machines: humans are often opaque to each other and even to themselves.[57]

The common correlative characteristics of bureaucracies, markets, and machine learning predisposes them to opacity. Correlations make conclusions credible; they do not make them comprehensible. In addition, because markets, bureaucracies, and machines typically assimilate a great many individual inputs and calculations, the sources of their results are difficult to trace, much less to comprehend.

The human response to this is quite variable. It was noted in Section 1, for example, that markets are valued as a means for establishing prices, but this process is opaque; prices are an emergent effect. On the other hand, there are circumstances in which opacity is not tolerated: regulators strive to compel transparency in important aspects of market operations. Similarly, the workings of bureaucracies are often opaque, but American laws like the Administrative Procedures Act establish some requirements for transparency. This variation in tolerance for opacity should sensitize those who consider opacity in machines that the problem is probably not best treated generally, but rather in a manner that distinguishes situations and establishes different priorities accordingly.

Mapping a Spectrum of Concern About Opacity in Bureaucracies, Markets and Machines

As a first approximation, we can distinguish four different kinds of interactions that humans have with artificial intelligences. Sometimes humans are "subjects" assessed and categorized by markets, bureaucracies, and machines. Sometimes they are "users" employing these intelligences without any need to understand their operations. As users, humans regard these artificial intelligences as Daniel Dennett described machine learning systems: they are "tools, not colleagues." However, there is a third set of interactions in which humans and intelligent machines do function as colleagues. In these "partnerships," decisions are shared and, as this discussion will suggest, a reduction in machine opacity is pursued. Sometimes humans have roles as "overseers" attempting to assure that artificial intelligences are operating properly. In this fourth role, opacity is a substantial problem.

Subjects

Most humans do not reflect on markets, bureaucracies, or machine programs from the perspective of Turing, judges in Turing's test, or DARPA program managers. Humans commonly experience these entities not as observers, but, reactively, as subjects. In these contexts, engagement with AI is often not chosen. Rather it is a necessity as when dealing with a bureaucracy to get a driver's license or pay taxes or when the operations of a market determine wages. Laypersons interacting with bureaucracies and markets do not ask, or care to learn much about, the construction or processes of these entities. They take the inhumanity of these systems as a given.

These reactions foreshadow many interactions with machine intelligence and point to measures of merit quite different from those in Turing's test. Cast in resented roles, the human frustration with bureaucracies and markets has been—and for machines likely will be—not whether the non-human systems think like humans, but whether they think about us as humans. We voice frustration at their opacity to us, but the cause of that frustration is our opacity to them.

Automated telephone trees illustrate this difficulty. If users have problems and requests that are well anticipated, then the systems work well. When they work well most who interact with them care not a whit about how they work. But if the machine does not comprehend a person's needs or, worse, renders results that seem unjustified, it leaves a human user indignant. In any system, undesired outcomes spark resentment, but an objectionable result is intensely inflammatory when it is the product of a system in which humans feel misjudged because they are misperceived. The primary objection to bureaucracies, markets, and machines is that after reducing humans to fragmented data points, they do not understand us. It is a secondary and less tractable problem that they do not understand how to make themselves understandable.

As we will see, in some contexts explainable AI is a useful goal, but in this context empathetic AI is a more important one.[58]

The core problem is not the opacity of machine, market, or bureaucratic processes. It arises from the reductionism foundational to these three artificial intelligences. These intelligences treat subjects only as data inputted through filters that eliminate human individuality. Careful structuring and transparency about these filters can facilitate fairness, but it will not eliminate insensitivity to individuality.[59] While correlation reduces comprehension, reductionism increases resentment.

Should societies strive to reduce the opacity of humans as perceived by machines? There are often undesired collateral consequences when bureaucracies, markets, and machines better understand the individuals with whom they interact. Movement in this direction sacrifices privacy, risks inappropriate discrimination between users, fans flames of concern about the ability of artificial intelligence to manipulate users (for example, to encourage gambling, shopping or political choices),[60] and intensifies fears that artificial general intelligence may evolve to manipulate its human masters. These issues are exemplary of a class of tradeoffs that will not be left to system designers and operators. They will prompt calls for regulation as discussed in Section 5 of this paper.[61]

Users

Market, bureaucratic, and machine intelligences attract voluntary users who seek to employ these instruments as sources of empowerment. Everyday observation shows laypersons entering markets, engaging with bureaucracies (for example, in corporations), and using machines without understanding the mechanisms that animate these systems. Opacity is generally not a material concern to those who employ these systems in a standardized way.

Society may require minimal user competence, as for example, with requirements that train, test, and certify new automobile drivers. But these drivers are not required, and rarely aspire, to understand how an automatic gearbox executes commands. To the contrary, for many, an attractive feature of modern design is that it disguises complexity. The focus is on choices (such as drive or reverse?) presented with maximum simplicity, rather than on the mechanisms that shape or implement that choice.

Just as we have seen that familiarity quiets concerns about artificiality, so for these users, familiarity also quiets concerns about opacity. Unfamiliarity creates anxiety about what is opaque, even for voluntary users. But the history of bureaucracies and markets suggests that this anxiety dissipates (though not completely) as new techniques become commonplace.

Partners

DARPA's program managers accurately capture another type of human-machine interaction when they refer to "the emerging generation of artificially intelligent partners." Partnership arises in human-machine interactions when human responsibility is not merely to use the machine but to reconcile or integrate its outputs with considerations outside the capacity of the machine.[62] As an agency of the U.S. Department of Defense, DARPA is particularly concerned with these situations in military combat, such as when a machine identifies an incoming aircraft as a hostile military airplane. Human oversight is commonly required before a missile launch against this target because the human can add a judgment about other considerations (Is this a wartime situation? Are there other options? And so forth.).

Often this role for a human is mis-described as an assertion that a decision is so important, it should be reserved for humans. The implication is that humans have superior, independent capabilities. In fact, though, humans—burdened by biases,[63] slower reaction and thinking speeds,[64] limitations in cognitive capabilities,[65] the influence of social and organizational pressures,[66] and the distorting effects of fatigue and emotion[67]—do not seem to be less prone to errors than machines. There are numerous vivid illustrations of humans causing catastrophic failures by rejecting correct machine judgments.[68]

The best argument for adding humans is that by compelling a partnership humans and machines are likely to perform better than machines or humans alone.[69] In addition, the combination is less vulnerable to malevolent subversion: requiring both machine and human intelligences is akin to requiring two-factor identification as a means of enhancing cybersecurity.

For this partnership between human beings and machines to function properly, the machine must be designed and the human systems trained in a manner sensitive to the strengths, weaknesses, and systemic consequences of both.[70] It is widely recognized, for example, that it is important to assure that machine outputs are continuously updated and informative, but do not overload the human. Concomitantly, there must be a human understanding of the machine, so that there is an appropriate amount of "trust" and a grasp of the strengths and weaknesses of the machine's data inputs, processing capabilities and programmed priorities. It is particularly essential that humans grasp that the correlative "intelligence" of the machine is probabilistic, not conclusively determined.[71]

Humans have—and should have—a low tolerance for opacity in these partnerships. One way of addressing this is to require human partners to have more training and certification than is required for those described in the preceding section as users. For example, humans demand a larger human responsibility for the use of a plane than for the use of a car and concomitantly require that a pilot be more extensively trained and tested than a driver. When dealing with complex transactions interacting with bureaucracies or markets, ordinary users commonly hire—indeed are often required to hire—professional intermediaries (lawyers, accountants, mechanics, etc.) trained to manage their interactions and provide a measure of reassurance that they are operating properly.[72]

A number of initiatives are taken to improve the "partnership" (though that term is not usually employed) between individuals on one hand and markets and bureaucracies on the other. Legally compelled, standardized and structured disclosures, for example, help human decision-makers (predominantly professionals) assess corporate financial information as inputs for their own decision-making.[73] It seems predictable and desirable that analogous methods will be applied to important machine learning systems when they are partnered with human intelligence.

Overseers

Those charged with designing or regulating artificial intelligences are in a different position from users, subjects, and partners. For users, subjects, and partners the goal is parsimonious: to inform them as much as, but no more than, needed. By contrast, those charged with assessing and mitigating systemic risk—here labeled overseers—must comprehend, not circumvent, complexity. While trust is a prerequisite for use, distrust is an imperative for oversight.

Opacity is a powerful enemy of oversight. Overseers can better understand machines, bureaucracies, and markets than users, subjects and partners because they can bring to bear more expertise, time, and information. But this does not mean they can completely dispel opacity. Present traumas associated with cybersecurity illustrate how difficult this is. It has been asserted that

> Von Neumann machines are deterministic and introspectable. You can fully describe them mathematically, and understand every step in their computing process. And if you put in one set of inputs, you can rely on them to always generate the same output.[74]

Nonetheless, numerous cybersecurity breaches have repeatedly demonstrated that the opacity of these systems is too great to allow designers, testers, and regulators to root out vulnerabilities that lurk within software, hardware and data. Furthermore, experience has shown that biases and values lie hidden in the algorithms that drive these systems.

Because machine learning systems are computational systems, these problems also will burden those who attempt to oversee intelligent machines. Beyond this, overseers will confront problems inherent in artificial intelligence. As described earlier, intelligent systems produce emergent results—by definition these results cannot be completely traced. And intelligent systems evolve. An assessment of their performance in test situations cannot be taken as an assurance about their performance at a later date and in unanticipated situations.

Accordingly, it seems likely that unexpected and undesired outcomes will accumulate as intelligent machines proliferate. Some of these will be traumatic. The history of markets and bureaucracies shows that demands for regulation and transparency have followed public trauma as day follows night. The 2008 financial crisis provides a representative example. This crisis was attributed by many to the opacity of "sub-prime mortgage derivatives" that misled those investing in the market, and to the limited regulation of bond and credit markets as compared to the stock markets.[75] Regulation ensued. Similar accounts could be given about the Sarbanes-Oxley regulations that emerged after the 2001 Internet Bubble, about the creation of the U.S. Securities and Exchange Commission (SEC) after the stock market crash of 1929 and numerous iterations in between these events.

If those responsible for the development of machine intelligence learn from this history, they can anticipate and moderate the demand for oversight by developing their own systems of professional self-regulation. Moreover, proponents of machine intelligence should recognize that regulatory oversight of bureaucracies and markets stimulated the development of career paths for auditors, investigators, lawyers and other professionals skilled at working with overseers and with the technologies they have sought to regulate.[76] These responses to opacity are not technical. They are investments in human capital.

Efforts to control undesired behavior in markets and bureaucracies suggest another lesson: control of intelligent systems requires continuous assessment, not just certification at the moment these controls are introduced. Though intelligent evolving systems demand testing and certification before they are put in service, key determinations about their performance are made from repeated observations after they have been put in operation. Lower-level bureaucrats and soldiers, for example, as well as civil servants, are opaque and evolve as they gain experience. Therefore, though they are trained, tested, and certified their power and responsibility is expanded primarily on the basis of experience.

Learning from this analogy, those responsible for testing and approving intelligent machines should move beyond typical methods of “verification and validation.”[77] These methods test machine performance against specified standards and presume that outcomes from development and then operational testing reasonably predict subsequent performance. Absent accidents, most organizations assign follow-up to other components charged with assuring that capability has not eroded, for example from wear or poor maintenance. With intelligent systems, whether “artificial” or “natural”, validation should never be complete, power should be only conditionally delegated, records of activities must be continuously reviewed, and authorizations for action expanded or withdrawn in light of experience. Overseers charged with evaluating intelligent machines should go beyond methods familiar from evaluating materiel and meld them with methods their organizations use for assessing personnel.

Furthermore, it is important to recognize that bureaucracies, markets, and intelligent machines are opaque not just because of the opacity of their systems, but also because of the data on which they depend. The AAA characterization of bonds in the 2008 sub-prime crisis should be a reminder of the powerful effects of mislabeled data. Dan Geer emphasizes how this raises the challenge of comprehending AI beyond even the challenge of comprehending computing machines. In machine AI “we no longer have source code—we have source data. This makes integrity, provenance, attribution, and so forth into first tier concerns where heretofore they were not."[78] Data is the lifeblood of markets, bureaucracies and machine learning systems. The next section considers its significance not only within each of these systems, but in tying them together.

Section 4: Bureaucracies, Markets, and Intelligent Machines as Data Dependent and Interdependent Systems

All intelligent systems depend on data. However, different ways of acquiring data have bred important differences between the three intelligences. Bureaucracies have historically secured data by command, markets by seduction, and machine systems by simulation or scavenging. These differences incentivize co-operation and crossbreeding between the systems. Patterns of access to data—fundamental to the power of each of these intelligences —are changing. As they do so, they create new relationships between markets, bureaucracies, and machines and new challenges for regulating these entities individually and collectively. This section explores these issues.

Data as a Source of Power for Bureaucracies and Markets

State bureaucracies were accumulating information long before students of machine intelligence labeled data "the new oil."[79] James O. Scott's 1998 classic *Seeing Like a State* identified creation and standardization of data as the means by which the state got "a handle on its subjects and their environment." By regularizing and registering "last names … weights and measures … population registers, [land] tenure" even "language and legal discourse," bureaucracies "took exceptionally complex, illegible and local social practices … and created a standard grid whereby it could be centrally recorded and measured."[80]

Before Scott, Friedrich Hayek, Charles Lindblom, and others recognized that the genius of markets was that they seduced individuals into accumulating information, assessing its significance, and then publishing that assessment in the form of offers to buy or sell a commodity, including their own labor. Markets stimulated, orchestrated, and rewarded this process; they then assimilated inputs to arrive at "efficient prices."

Responding to twenty-first century developments, Marion Fourcade and Kieran Healy have observed how digital software places market "data imperative[s]"[81] on steroids.

> Like Scott's administrative designs, digital economy's classificatory architecture allows market institutions to apprehend their clients, customers, or employees through new instruments of knowledge, efficiency and value extraction. Markets have learned to 'see' in a new way ... It used to be that the state was the only organization with the resources to identify and track individuals across many contexts and settings. No longer.[82]

Making similar observations, Shoshana Zuboff has meticulously documented and fiercely criticized "an information civilization" in which "surveillance capitalism unilaterally claims human experience as free raw material for translation into behavioral data.... Everything must be illuminated for counting and herding."[83]

> [T]he "cookie" – bits of code that allow information to be passed between a server and a client computer – was developed in 1994 at Netscape, the first commercial web browser company. Similarly, "web bugs" – tiny often invisible graphics embedded in web pages and email and designed to monitor user activity and collect personal information – were well known to experts in the late 1990s....
>
> Google brought new life to these practices... integrating a wide range of mechanisms from cookies to proprietary analytics and algorithmic software capabilities in a sweeping new logic that enshrined surveillance and the unilateral expropriation of behavioral data as the basis for a new market form.[84]

Machines Require Data and Amplify the Power of Bureaucracies and Markets to Manipulate Data

Intelligent machine systems do not have independent power to command or seduce. For most of the history of "artificial intelligence" as a field, designers of machine systems have had to simulate data or scavenge it from other sources.[85] Simulation can be helpful for testing and training.[86] Over the longer term, it is possible this capability will create virtual worlds that are so rich they can substantially reduce machines' need for real world data. But for at least the near future analysis of more complex dynamic systems (as for example, the activities of human beings) require real-world data.

To date, the dominant collectors of this data have been bureaucracies and markets. Machine intelligence brings unprecedented speed and processing capabilities to these systems. In turn, bureaucracies and markets provide data to those who develop and manage machine intelligences. The results are marriages made in heaven, or hell for those who fear any or all of the triad of machine, market, and bureaucratic power.

The Challenge of Sharing Data

Within the triad, there is conflict as well as complementarity. For example, the progress of machine intelligence is significantly dependent on access to data, but bureaucrats maximize their power by controlling information. Consequently, bureaucratic resistance to sharing data frequently becomes a critical impediment to the development of machine capabilities.

The problem is most evident in government bureaucracies and particularly acute in national security agencies with traditions of secrecy.[87] Otherwise admirable blueprints for progress have been slow to come to grips with this impediment in part because the issue is political, not technical.[88] The Department of Homeland Security's "Artificial Intelligence Strategy" contents itself with the technical: "DHS will assemble a cadre of internal and external computational capacity, data storage, and security experts across DHS Components to survey the current state of the Department's AI ready infrastructure to make recommendations on how it can be improved."[89] Reviewing artificial intelligence in the intelligence agencies and the Department of Defense, the National Security Commission on Artificial Intelligence "Final Report" decries present "weak data practices."[90] However, it devotes little more than a paragraph in the main body of this 270-page report to data sharing and its sole related recommendation is to call for "Data architecture composed of a secure, federated system of distributed repositories linked by a data catalog and appropriate access controls that facilitates finding, accessing, and moving desired data across the DOD."[91]

To get to the core of the challenge of data sharing, a reader would have to attend to a footnote in the National Security Commission report that recognizes, "This hinges on implementation of the DOD's new data strategy."[92] That largely hortatory document begins to get to the right point, but offers only a pollyannaish view: "DOD is making the cultural shift from the need to know (i.e., information withholding) to the responsibility to provide (i.e., information sharing)."[93] It is hard to square this with this author's and many other observers' observations[94] or with the National Security Commission's conclusion that within DOD "[t]he data that is needed to fuel machine learning (ML) is currently stovepiped, messy, or often discarded. Platforms are disconnected."[95]

More generally, there is an imperative for public investments for compiling, storing, standardizing, sharing, and securing data. Efforts are biased towards research in machine technologies and demonstration projects that exploit available databases. Building databases is like building infrastructure—it is unglamorous, has high near-term costs and demands creative skills quite different from those predominantly celebrated in our technical communities.[96] (Nobody ever won a Nobel Prize for building a database.[97])

Moreover, in the United States efforts at synergies and control over data run uphill on a terrain filled with crevasses because markets and state bureaucracies are regarded as separate "private and public" sectors with awkward, even hostile and suspicious, relations between them. Privacy concerns, bureaucratic and corporate rivalries, fear that sharing data with the government will lead to government regulation, and corporate zeal to protect intellectual property, compound this separation. The separation diminishes both control and opportunity.[98]

Over time, access to data may become less of a problem. Digitization and the networking of an ever-growing majority of activities is increasing usable data for machines. Machine intelligence will improve in proportion to systems' access to digital footprints left by portable devices and vehicular movements, financial transfers, security camera videos, Nest thermostat settings, and countless other devices connected in an Internet of Things. At the same time, simulation and a proliferation of virtual environments may provide increasingly robust alternatives for generating data for machines. Though the three intelligences will no doubt remain intertwined, these developments are likely to weaken machine dependence on markets and bureaucracies, increasing the independent power of machines and those who control them.

The maturation of machine intelligence is following a path that may be metaphorically compared to that of a human. At first, infants develop intelligence largely through improvements in processing power— they are developing and connecting synapses;[99] a child augments this with simulation in the mode of "play;" adolescents take a big further step by accessing a world of information through many independent sources.[100] Viewed this way, machine intelligence is entering adulthood, with an expansion of its capabilities only hinted at to date.[101]

The centrality of data is emphasized as we move away from a conceptual framework that predominantly compares machine intelligence with the individual human mind—an instrument with limited capacities for processing data. Bureaucracies, markets, and machine intelligence all absorb and depend upon systematic inputs of data in amounts no human can fully comprehend. These systems meet around pools of data as animals gather at shared water holes. But while animals have generally stable patterns of interaction, the artificial entities are evolving rapidly. Machine intelligence is changing most quickly, humans least so, and bureaucracies and markets somewhere in between. The challenge is not one of managing machine intelligence in isolation, but of managing the set of artificial intelligences as they evolve together. Seen this way, data are dominant sources of power and control.

A Comment on Artificial General Intelligence

Access to data is central to concerns about an hypothesized "artificial general intelligence," a super-intelligent machine ability that Elon Musk and others have called "our biggest existential threat."[102] Nick Bostrom's formative discussion published seven years ago, identified an "[a]n AI takeover scenario" as "of paramount importance."[103] This scenario supposes a future machine with an intelligence so powerful that it could "assert itself against the ... world."[104]

This essay is focused on present realities, not on machine intelligence as it might come to be in some future state. A discussion of machine, market and bureaucratic data dependency should note, however, that bureaucracies and markets have long been considered as forces that could "assert [themselves] against the world." We know from experience that their power is proportionate to their capacities for gathering data and their vulnerability to data limitations and distortions. So it is likely to be with any feared future machine super-intelligence.

Intelligence is not omniscience. Intelligence is impotent without information. Even a superintelligence knows nothing without data. Even if it simulates data (randomly or from a model), it is flying blind unless it is provided data to match its simulation with the real world. Even if given dominion over weapons, it needs knowledge of those weapons and targeting information to use them effectively.

Moreover, the history of bureaucracies and markets teaches that day-to-day control of human beings is exercised less by force and more by the manipulation of rewards and punishments. Bostrom's dystopian future envisions an intelligent machine with skills that would include planning, strategizing and "social and psychological modeling, manipulation, rhetoric and persuasion," enabling it to "leverage external resources by recruiting human support."[105]

The road to the hell that Bostrom fears is likely to be paved by data about human beings. Consequently, concerns about machine power should be as much about data as about machines. And these concerns must consider the entire reservoir of information that flows between bureaucracies, markets and machines.[106]

Section 5: The AI Ecosystem and Its Regulation

Modern societies are more capable and resilient because they developed two rather than a single type of artificial intelligence to absorb data and produce decisions. In the modern world, some problems have been better dealt with by the right hand of bureaucracies; others by the left hand of markets; some require two hands.[107] Capitalism, for example, rests on the premise that markets are excellent mechanisms—and, by and large, bureaucracies are poor mechanisms—for determining prices.[108] By contrast, even ardent capitalists would regard it as anathema to use a market in place of a bureaucracy charged with adjudicating whether the characteristics of an applicant (for example, for disability assistance) or the conduct of a defendant had met certain legal standards. Markets and bureaucracies co-exist because, like knives and spoons, they are different tools suited to different uses. Over time humans have developed rough rules for choosing, coordinating, and employing these tools. How will development of a third artificial intelligence change existing patterns and practices?

The Two-Body System

Markets and bureaucracies typically employ different processes and norms to serve their different functions.[109] Both systems take human judgments as inputs but process them quite differently. Bureaucracies aspire to constrain individual human judgments. They accomplish this by channeling personal views through procedures that demand rationalization by deduction from rules.[110] "Faceless bureaucrat" is a term of derision, but this characteristic is typically desired—it is manifested by the symbol of blindfolded justice, by covering judges' street clothes with robes, and by eliminating personal pronouns in favor of phrases like "this Court" and "the Social Security Administration." By contrast, markets are auctions—processes designed to encourage humans to scavenge information and then to elicit resulting conclusions as contributions to a composite judgment. The composite pretends to be no more than the sum of unabashedly self-interested, imperfect and idiosyncratic human judgments.[111]

Concomitantly, the two intelligences have developed contrasting standards for screening data. In stock markets, regulators attempt to exclude corrupt inputs. But with this exception, markets are encouraged to reflect inputs that are vast, varied and often obscure. Changes in weather are quickly assimilated as inputs in commodity prices; changes in taste are reflected in the value of products of the fashion industry and in the value of corporations that produce these products; macro-changes in political and economic environments are reflected throughout the market. Even before these adjustments occur, predictions about them are valid inputs. Markets do not pretend to value only disinterested or well-informed judgments.

By contrast, data considered by bureaucracies are strictly controlled. The only items that are intended to be considered are those that are appropriately placed on the scales of justice.[112] Those items must also be visible. Administrative decisions must be based on items in "the record." "Off-the-record" discussions are proscribed. The resulting focus excludes considerations that with a broader view may be deemed important to achieving a "just" result.[113]

Consistency and predictability are differently valued and differently achieved by the two systems. If the workings of markets were predictable, they wouldn't be needed. Markets are valued as tools to produce emergent effects.[114] By contrast, predictability and consistency are prerequisites for bureaucracies. While markets fluctuate with the weather, it would be misfeasance or malfeasance if a bureaucracy changed its decisions in response to the political or meteorological climate of the day.[115]

The rules for achieving the dexterity that coordinates the two are not rigid. Sometimes the systems are combined, as when ceilings on the price of utilities or housing are established by bureaucracies intruding in otherwise free markets, and when markets (for example, the New York Stock Exchange) are overseen by bureaucratic regulators (e.g. the SEC), managed by a corporate bureaucracy and composed of competing bureaucracies (corporations).[116]

Adjustments caused by machine intelligence as a newcomer to this ecosystem

Certainly, there is much to be determined, but already machine intelligences enhance the capabilities, alter the behaviors, and compete with bureaucracies and markets. Because markets are designed for speed and plasticity and bureaucracies are engineered for ponderous predictability, it is not surprising that markets are absorbing machine intelligence more rapidly than bureaucracies. But both are indisputably changing. In turn, the trajectories of the development and use of machine intelligences are being shaped by the gravitational pull of markets and bureaucracies.

The turbulence from this three-body problem can be observed and its intensification predicted with more confidence than its outcomes. The first two decades of this century have provided examples of this change that would previously have been hard to imagine. Google (now Alphabet) recognized that widespread use of its machine intelligence for internet search[117] enabled it to create a market matching the supply of user attention to the demands of advertisers.[118] That achievement is exposing it to increasing regulatory demands from government bureaucracies. Similarly, Uber and Airbnb married AI technologies to market techniques (for example, eliciting user ratings to develop market evaluations of drivers and landlords) and now struggle with government efforts to make them subject to bureaucratic regulation.[119]

The computational and memory capabilities of machine intelligences will themselves challenge bureaucracies and markets. Since bureaucracies are designed to insulate decisions from human idiosyncrasies, if machines can equal, not to mention exceed, human capabilities then why use humans at all?[120] And why should markets endure when machines can now assimilate more data, run more simulations, and discern patterns better than the crowd of humans who contribute to market judgments?[121] Alpha Go proved the principle, and the fact that some three-quarters of American stock trading is now self-initiated by machines[122] demonstrates the power of its application. Accordingly, machine intelligences threaten their predecessors as markets and bureaucracies threatened and ultimately displaced feudal systems.

There is room, nonetheless, for bureaucracies, markets and machines to coexist in a further evolved ecosystem.[123] Paradoxically, bureaucracies may be valued for their humanness and because they are the most likely tools for controlling machines. Markets, at least for the foreseeable future will be advantaged by the fact that they engage a great many human data gatherers and decision-makers (otherwise known as buyers and sellers) to assimilate a largely unrestricted range of exogenous information.[124] Machines are like blinkered horses. Put them on a road and they will pull heavy loads straight ahead. Humans can benefit (though they also risk distraction) from a wider aperture for sensing their environment and reconsidering their direction.[125]

These advantages are intensified in contexts in which the goal is to understand and predict human behavior. In many situations, the required intelligence is not simply to assess moves with inanimate pieces (as in Go) but with human pieces that may respond to direction but also wander off on their own volition in response to stimuli presently impossible to quantify comprehensively.[126] A fundamental question is whether and when, in assessing human behavior, the computational capacity of the computer will overtake the wisdom of the crowd.[127]

Operational and Regulatory Controls

In their infancy, new technologies, whether intelligent or not, do not attract much regulatory attention because they are not important to the operations of society.[128] This changes as dependence on these technologies increases. Missteps then catalyze demands for regulation, and these demands eventually lead to the creation or adaptation of bureaucratic mechanisms for oversight. The histories of markets and of bureaucracies repeatedly reflect this evolution: markets are regulated by new or expanded bureaucracies and bureaucracies are regulated by new, higher layers of bureaucratic management or by other bureaucracies (including courts, investigative agencies, and organizations with shared responsibilities).

However, technological innovation occurs faster than the human consensus required to create regulatory systems. This is because remedies commonly trail causes, and because human systems are organic and organic growth is slower than mechanical creation. A fence can be built quickly, a garden must be nurtured. Consequently, risks rise as a new technology proliferates.

The historic challenges of controlling markets and bureaucracies should sensitize policy-makers to the magnitude and urgency of the regulatory effort. In the third section, for example, it was suggested that effective test and evaluation of intelligent machines systems will require rules of engagement, oversight, control, and audit like those created to control human beings in military and bureaucratic organizations. In military services the systems for management of soldiers, Marines, sailors, and airmen and women were constructed and revised with the benefit of decades, indeed centuries, of experience.[129] Equivalent systems for the control of machine intelligence will be similarly demanding but need to be constructed much, much more rapidly. The challenge is not only to invent a system of bureaucratic intelligence that can control machine intelligence but also to do it at a speed comparable to the speed of proliferation of machine intelligence.

Numerous cross-currents will impede attempts to construct this system. One, previously discussed, is that machines will be simultaneously disrupting bureaucracies. Another is that regulation will create problems of its own. It will produce a backlash resulting in new controls intended to force regulators to be quicker, more or less risk prone, more or less insulated from those they regulate, etc.[130] A third is that regulation is likely to be fragmented and even competitive. Already, initiatives to control intelligent machines are being debated in US government agencies as diverse as the SEC, the Federal Trade Commission, the Federal Communications Commission, the Federal Aviation Administration, and others.[131]

The history of bureaucracies and markets suggests that It is not predictable how this will play out, but it is predictable that the future of machine intelligence will significantly depend on lawyers, economists, auditors, accountants, inspectors, and others.[132] The sources, training, experience, and consequent perspectives of this human cohort will be critical.[133] At present, investments in research and development greatly outstrip investments in educating the community of lawyers, economists, auditors, and others who will regulate machine intelligence. Current implicit priorities could be phrased as tons for the development and acquisition of technology, pounds for training those who use these technologies and ounces for those who will regulate it. Proponents of machine intelligence will be well-advised to adjust this balance so that regulators are better informed about the capabilities, limitations and biases of systems of machine intelligence.

The speed, power, and proliferation of machine intelligences will outstrip human experience with bureaucracies and markets. The challenges will be concomitantly greater. On the other side of the balance sheet, bureaucracies and markets arose in societies in which they were shockingly new and untrammeled by equivalent powerful artificial intelligences. The present newcomer, machine intelligence, enters an already populated landscape. It will be wise to learn from experience with bureaucracies and markets, to focus on how machine intelligence will be employed in the context of bureaucracies and markets, and to use bureaucracies to help tame machine intelligence and sustain its use for human benefit.

Conclusion

Markets and bureaucracies, powerfully present in everyday life, are largely absent from the literature on artificial intelligence. This essay has assessed experience with these long-standing systems to cast light on the newcomer of machine intelligence. It shows how viewing the three systems as a set of artificial intelligences can supplement and correct views drawn from a predominant paradigm focused on comparing machine intelligence to an individual human mind.

Markets, bureaucracies, and intelligent machines are reductionist. These tools absorb and process large amounts of data, but they do so by stripping those inputs of context and qualification. They are also mechanisms for assessing correlations without causative insight. These characteristics combine to make each of the three systems remarkably smart and remarkably stupid, remarkably powerful and remarkably vulnerable.

Considering the three tools together yields descriptive, analytic, predictive, and prescriptive insights. Some attributes of machine intelligence are normalized when seen as common consequences of reductionist, correlative systems. Others are highlighted as unique. A sound perception of each of the three systems requires understanding how it interacts with the others.

A common dependence on data, for example, brings the three systems into an ecosystem in which each member complements, amplifies, competes with, and creates hybrids with the others. Balances in this ecosystem are evolving and seem likely to shift as the Internet of Things and simulated environments generate information streams largely independent of bureaucratic and market systems.

Widening the aperture to consider the three intelligences highlights opportunities for considering the history of markets and bureaucracies as indicators of potential developments in human interactions with machine intelligence. For example, early in their development, markets and bureaucracies were criticized because of their perceived artificiality, “inhumanity,” and opacity to those enmeshed in them. These concerns dissipated with time and are likely to do so for machine intelligence.

By contrast, history suggests that some other concerns are likely to intensify. Enthusiasts for markets and bureaucracies have perceived those systems as value-neutral means of efficient information processing, allocation, and coordination. Proponents of machine intelligence laud similar capabilities in these systems. However, the histories of bureaucracies and markets suggest that this perspective cannot be sustained. Over time, markets and bureaucracies have repeatedly been subjected to regulation to reset values often unconsciously embedded within their operations and to control abuses and unforeseen or undervalued externalities. Similar efforts at regulation should be expected in response to machine intelligence.

Near-term regulation is likely to focus on machine intelligence as it operates in markets and bureaucracies. These efforts are likely at first to follow well-established paths for assuring quality control in machines. They will, however, need to be reconsidered and new tools developed to account for the speed, processing power and adaptive character of machine intelligence. Bureaucratic regulation of machine intelligence may not be best achieved (as presently envisioned) through improved systems of test and evaluation. Instead, they may prove to be more like the challenge of training, auditing, incentivizing and supervising lower-level administrators in a bureaucracy. This supervision will be demanding not only in itself, but also because at the same time as bureaucracies attempt to control machines, machine intelligence will be remaking bureaucracies.

In sum, the shift in definition described here widens the aperture of thinking in a way that moderates some familiar issues, highlights others, and suggests new perspectives. Most fundamentally, it emphasizes that machine intelligence is a tool that compensates for human limitations in detecting correlations, that other tools exist for similar purposes, and that control of the suite of artificial intelligences is both complicated and facilitated by the multiplicity of these tools. The present challenge is not just to manage machine intelligence, it is to manage an ecosystem of machine, bureaucratic, and market intelligences that together complement, but also challenge, human capabilities.

Author

Richard Danzig a Senior Fellow at the Johns Hopkins University Applied Physics Laboratory, a Trustee of the RAND Corporation, a Director of the Center for a New American Security, and a Director of Saffron Hill Ventures (a European investment firm). In recent years, he has been a member of President's Intelligence Advisory Board, the Secretary of Defense's Defense Policy Board, the Homeland Security Secretary's Advisory Council, the Aspen Strategy Group, the Toyota Research Institute Advisory Board, the Cyber Resilience Forum of the National Academies of Sciences, Engineering, and Medicine and the Reed College Board of Trustees. Dr. Danzig served as Under Secretary of the Navy from 1993 to 1997, and then as Navy Secretary until January 2001.

Acknowledgments

There is a symmetry between a message of this paper and the method by which it was written and rewritten. While these pages present my thinking about the characteristics of bureaucracies and markets as collective intelligences, my thinking has itself been improved by collective intelligence. As the sole author I bear sole responsibility for this paper's errors, but I have no doubt that its strengths significantly derive from the generosity and creativity of others who read drafts, participated in a workshop, and talked with me about issues.

DARPA sponsored most of this work, but the views, opinions, any findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. John Launchbury and John Everett, then Director and Deputy Director of DARPA's Information Innovation Office, encouraged and funded me to develop the original concept. Their successors, Bill Scherlis and Jen Roberts, sustained this support and contributed ideas and criticisms. Matt Turek took on the most central role as program manager and engaged the valuable help of Julie Fitzgerald, as his technical assistant. Phil Root and Ted Senator, members of DARPA's Defense Science Office, contributed supportive suggestions. Mike Hopmeier and Yong Kim at Unconventional Concepts provided contract administration. In addition, Mike contributed his substantial insight and his colleague Lori Settle did a remarkably good job of creating workshop notes.

Workshop participants, in addition to those already named, included: Greg Allen, Jeff Alstott, Jim Baker, Peter Bloom, Andrew Burt, Teddy Collins, Allan Dafoe, Matt Daniels, Thomas Dietterich, Craig Fields, Tim Hwang, Chris Inglis, Neil Johnson, David Krakauer, Jason Matheny, Melanie Mitchell, JF Mergen, Deirdre Mulligan, Dewey Murdick, Philip Nelson, Jose Hernandez-Orallo, Divya Siddarth, Philip Thomas, Mike Walker and Glen Weyl.

Several workshop participants provided extraordinary help by reading and commenting on drafts in correspondence unconnected to the workshop. I am especially grateful to Peter Bloom, Melanie Mitchell and Mike Walker for their generous contribution of ideas throughout the effort.

Some friends who did not participate in the workshop provided individual invaluable help. Graham Allison, Philip Bobbitt, Bob Butler, Chris Castelli, Eliot Cohen, Steve Lipner, Gary McGraw, A, Steve Dell, Dan Feldman, Laura Freeman, Alberto Gaitan, Ben Garfinkel, Dan Geer, Peter Levin, Andrew McKay, Anup Malani, Alex Montgomery, Dan Patt, Dennis Patterson, Heather Roff, Rasheed Sabar, Ted Senator, V. Siddhartha, and Eric Schmidt contributed criticism and insights. Henry Farrell generously shared his thinking about markets and bureaucracies as information processing entities, leading to reading of some of his work and that of others cited in the footnotes. Jim Miller read more than one draft and helped the work not only substantively, but also by his enthusiasm for it. Ralph Semmel and Christine Fox organized a rewarding discussion at Johns Hopkins University Applied Physics Laboratory. I am grateful to Johns Hopkins and particularly to Ralph and Christine for their consistent support of my work as an APL Senior Fellow.

Finally, Andrew Imbrie, Helen Toner, and Igor Mikolic-Torreira at the Center for Security and Emerging Technology contributed close readings of the document and Corey Cooper, Maura McCarthy, Matt Mahoney, Dewey Murdick and Lynne Weil played important roles bringing it to publication.

As a lawyer, I reflected some on whether my requests to many of these people for more than one reading of this paper unfairly exposed them to double jeopardy or went beyond that to constitute cruel and unusual punishment. Their generous responses reflect their personal as well as intellectual qualities. I am intensely grateful to all those I have mentioned.

Document Identifier: doi: 10.51593/ 20210046

Endnotes

[1] This minimal definition may be applied and extended in different ways. The essays in David Klahr and Kenneth Kotovsky, *Complex Information Processing: The Impact of Herbert A. Simon* (United Kingdom: Psychology Press, 2013), illustrate different approaches derived from Simon's insight that intelligence may be thought of as complex information processing. Simon's colleague James G. March particularly developed this theme in the context of organizations. For example, his paper "Rationality, Foolishness and Adaptive Intelligence," *Strategic Management* Journal 27, no. 3 (2006): 201 begins:

> Organizations pursue intelligence. That is, they can be described as seeking to adopt courses of action that lead them over the long run to outcomes that they find satisfactory, taking into account any modifications of hopes, beliefs, preferences, and interpretations that occur over time, as well as conflict over them. The pursuit of intelligence is an ordinary task. It is neither mysterious nor unusually difficult. It is carried out by ordinary organizations in ordinary ways every day in ways that permit most of them to survive from day to day.

A systems biologist emphasizes the adaptive aspect of this perspective. For David Krakauer, President of the Santa Fe Institute, intelligence encompasses "[t]he procedures for arriving at adaptive decisions based on approximate and noisy information …" David Krakauer, "The Computational Systems of the World," *BioScience* 64, no. 4 (April 2014): 351-354, https://academic.oup.com/bioscience/article/64/4/351/248408.

[2] This analysis could be extended to include other forms of collective intelligence. Legislatures, for example, have attributes of bureaucracies and markets, but are a distinct form of organization. Corporations like Uber are discussed in Section 5 of this paper as examples of hybrids that combine bureaucratic, market and machine intelligence.

This paper may be read in the context of recent discussions of socially distributed cognition in problem solving and in the evolution of human culture. See, for example, Thomas W. Malone, *Superminds: The Surprising Power of People and Computers Thinking Together* (New York, NY: Little, Brown, and Company, 2018), Annie Murphy Paul, *The Extended Mind: The Power of Thinking Outside the Brain* (Boston, MA: Mariner Books, 2021), and Timothy M. Waring and Zachary T. Wood, "Long-term gene–culture coevolution and the human evolutionary transition," in *Proceedings of the Royal Society B* 288, no. 1952 (June 2021), https://doi.org/10.1098/rspb.2021.0538. Waring and Wood comment, "First proposed by Maynard Smith & Szathmáry, evolutionary transitions are thought to unfold via a shift in the dominant level of selection from competitive individuals to well-integrated functional groups….These transitions exhibit a common set of patterns, including new divisions of labour, the loss of full individual autonomy and reproductive control, and the rise of new routes of information transmission." See also Eörs Szathmáry, "Toward major evolutionary transitions theory 2.0," in *Proceedings of the National Academy of Sciences of the United States of America* 112, no. 33 (April 2015), https://doi.org/10.1073/pnas.1421398112 and Francisco Brahm and Joaquín Poblete, "The Evolution of Productive Organizations," *Nature Human Behaviour* 5 (September 2020): 39-48, https://extranet.sioe.org/uploads/sioe2019/brahm_poblete.pdf. Alan Dafoe et al have recently proposed recognition of a field of "cooperative AI," to "study problems of cooperation through the lens of artificial intelligence and to innovate in artificial intelligence to help solve these problems. Whereas much AI research to date has focused on improving the individual intelligence of agents and algorithms, the time is right to also focus on improving social intelligence: the ability of groups to effectively cooperate to solve the problems they face." Alan Dafoe et al., "Open Problems in Cooperative AI" (DeepMind, 2020), https://deepmind.com/research/publications/2021/Open-Problems-in-Cooperative-AI.

[3] Ludwig Wittgenstein, *Philosophical Investigations* (Hoboken, NJ: Wiley-Blackwell, 1953), sections 66 and 67. The concept of fuzzy definitions invoked in footnote 7 derives from Wittgenstein.

[4] James Flynn, *What is Intelligence: Beyond the Flynn Effect* (Cambridge, MA: Cambridge University Press, 2009), 51.

[5] When commissioned by NASA to define "life," Cleland and Chyba concluded that no adequate definition could be constructed given present deficiencies in our understanding. They offered the analogy of early efforts to define "water." Only after "we had an understanding of the molecular nature of matter ... could [we] identify water in such a way that all ambiguity disappears: water is H_2O. This identification holds regardless of whether the water is in any of its familiar solid, liquid, or vapor phases, and it will hold equally well for water in less familiar high pressure solid phases." Carol E. Cleland and Christopher F. Chyba, "Defining 'Life'" in *Origins of Life and Evolution of the Biosphere* 32 (2002): 390, http://www.aim.univ-paris7.fr/enseig/exobiologie_PDF/Biblio/Cleland%20and%20Chyba%20_2002.pdf. Cleland later expanded this analogy:

> [M[edieval alchemists classified many different kinds of substances as water, including nitric acid (which was called "aqua fortis"). They did this because nitric acid exhibited many of the sensible properties of water, and perhaps most importantly, it was a good solvent. It wasn't until the advent of molecular theory that scientists could understand why nitric acid, which has many of the properties of water, is nonetheless not water. Molecular theory clearly and convincingly explains why this is the case: water is H_2O— two hydrogen atoms and one oxygen atom. Nitric acid has a different molecular composition.

Astrobiology Magazine Staff, "Interview with Carol Cleland" in "Life's Working Definition: Does it Work?," NASA, https://www.nasa.gov/vision/universe/starsgalaxies/life%27s_working_definition.html.

Richard Feynman wisely admonished that "extreme precision of definition is often not worthwhile, and sometimes it is not possible – in fact mostly it is not possible." Richard P. Feynman, *The Meaning of It All: Thoughts of a Citizen Scientist* (New York, NY: Basic Books, 1998), 20. See also 4-5. Marvin Minsky's seminal 1985 discussion, *The Society of Mind* (New York, NY: Simon & Schuster, 1988) cautioned "It often does more harm than good to force definitions on things we don't understand."

[6] George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago, IL: University of Chicago Press, 1980), 10. See also Vinciane Despret, *What Would Animals Say If We Asked Them the Right Questions?* (Minneapolis, MN: University of Minnesota Press, 2012, tr. from French by Brett Buchanan, 2016). Approaching the question, for example, "Do chimpanzees mourn?," Despret emphasizes that we must ask what mourning means for us:

> [W]e must, so as not to exclude chimpanzees from the start, put our own concerns to the test… I am not suggesting that chimpanzees will propose a new theory of mourning…But they call for us to reactivate our lifeless versions, they oblige us to rethink, they put our … versions to the test…[C]alling one [version] 'dominant' privileges a certain kind of story, solicits a certain kind of attention to some behaviors rather than others, makes imperceptible the relation to other possible versions (173–5).

[7] Definitions also will vary among users. "Concepts … are bundles of regularities tied together by family resemblance, collections of varyingly held properties or traits which are united in some instrumentally useful way which justifies the unification. When we attach word-handles to these bundled concepts, in order to wield them, it is frequently though not always for the purpose of communicating our concepts with others, and the synchronization of these bundles across decentralized speakers, while necessary to communicate, inevitably makes them a messy bundle of overlapping and inconsistent senses—they are 'fuzzy,' or 'inconsistent,' or 'polysemous.'" Suspended Reason, "Conceptual Engineering: The Revolution in Philosophy You've Never Heard Of," *Lesswrong*, June 2, 2020, https://www.lesswrong.com/posts/9iA87EfNKnREgdTJN/conceptual-engineering-the-revolution-in-philosophy-you-ve. "The word 'intelligence' is indeed likely to give rise to misconceptions, partly owing to … overtones and partly to the suggestion implicit in an abstract noun that the word refers to 'something…'" Alice W. Heim, "Intelligence: Its Assessment" in Richard K. Gregory (ed.), *The Oxford Companion to the Mind* (Oxford, UK: Oxford University Press, 1987), 379.

[8] U.S. Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy* (Washington, DC: Department of Defense, 2018), https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF. The Congressional Research Service observes, "Although there is no official U.S. government definition of artificial intelligence (AI), AI generally refers to a computer system capable of human-level cognition." Kelley M. Sayler, "Defense Primer: Emerging Technologies" (Congressional Research Service, updated June 8, 2021), https://crsreports.congress.gov/product/pdf/IF/IF11105.

[9] Nathan Benaich and Ian Hogarth, "State of AI Report," October 1, 2020, slide 5 https://docs.google.com/presentation/d/1ZUimafgXCBSLsgbacd6-a-dqO7yLyzIl1ZJbiCBUUT4/edit#slide=id.g557254d430_0_32.

[10] Turing wrote that he turned to this test because of his frustration with efforts to define terms like "think" and "machine." "Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words." A.M. Turing, "Computing Machinery and Intelligence," *Mind* (October 1950), https://www.abelard.org/turpap/turpap.php/. In an interview, Marvin Minsky said that Turing never intended this test to be more than "one way" of assessing machine intelligence. See Singularity Weblog, "Marvin Minsky at AI: The Turing Test is a Joke!" *YouTube*, July 12, 2013, https://youtu.be/3PdxQbOvAlI?t=1416. This suggests that the subsequent undesirable influence of the test is a fault of those who have made it central, rather than a fault of Turing himself.

Dennett makes the nice observation that Turing's test puts humans and computers behind a screen to counter biases that predispose us to find intelligence only in humans and not in machines. "The insight underlying the Turing test is the same insight that inspires the new practice among symphony orchestras of conducting auditions with an opaque screen between the jury and the musician.... His point was that we should not be species-chauvinistic, or anthropocentric, about the [physical characteristics] of an intelligent being, for there might be inhuman ways of being intelligent." Daniel C. Dennett, *Brainchildren: Essays on Designing Minds* (Cambridge, MA: The MIT Press, 1998), 5, 13. But while the Turing's test's blinded process reduces some bias, observer judgments about what is human will still be affected by bias as to class, race, culture etc. Worse, the test is fundamentally anthropomorphic in its standard. It is as though having placed a cellist behind the screen, the judges assessed the cellist not according to the qualities of his or her music, but how well that music approximated the sound of a violin. Humans are not the only instrument in the orchestra of intelligence (though they may at present be the most versatile).

A focus on human intelligence also presupposes a single standard for human intelligence. We know that there are emotional as well as cognitive intelligences and that both intelligences have varied methods, capabilities, distortions, and results. Yet worse, the Turing test commonly presupposes (without saying so) a benchmark stereotypical adult with cultural characteristics and training usually matching those of the observer who of course brings her or his own biases. The Turing test is accordingly a measure of the matches between the idiosyncrasies of the observer and the machine.

There is no universal human. Though it is widely thought that no machine now meets the Turing test, for many observers a modern machine would probably prevail in such a test if it pitted a computer programmed to converse in the observer's language against a foreign tourist or a child. As conceived and applied in the mid-20th century, the Turing test unconsciously reflected the biases of the group that regarded it as a standard. Implicitly their goal was for a computer's output to be indistinguishable from that of a male assimilated in Anglo-American society.

[11] Valuable descriptions of bureaucracies referenced in this essay include James Q. Wilson, *Bureaucracy: What Government Agencies Do and Why They Do It* (New York, NY: Basic Books, 1989) and Ludwig Von Mises, *Bureaucracy* (New Haven, CT: Yale University Press, 1944). Both build upon (though differ at points from) the seminal work of Max Weber. For markets, Charles E. Lindblom, *The Market System: What It Is, How It Works, and What to Make of It* (New Haven, CT: Yale University Press, 2001) provides a classic treatment. Frederick Hayek's influential 1945 article remains well worth reading and reference is made to it in this paper for its description of commodity markets as allocational systems. Friedrich A. Hayek, "The Use of Knowledge in Society," *The American Economic Review* (September 1945): 519-530, https://www.jstor.org/stable/1809376?seq=1#metadata_info_tab_contents.

Economists distinguish commodity markets and matchmaking processes. Both create matches, but in the first "price does all the work"—products are standardized, substantial numbers of participants tender bids and auction mechanisms impersonally determine a price and match buyers and sellers. Match-making markets are more individualized, as for example, in applications for college admission or employment in a corporation. This paper focuses predominantly on the former. Alvin E. Roth, *Who Gets What – and Why* (Boston, MA: Mariner Books, 2015) provides a valuable discussion of matchmaking markets.

[12] Other characteristics may be asserted as important to bureaucracies and markets or fundamental to Intelligence. For example, bureaucracies and markets may be viewed as coordinating mechanisms and machine intelligence may be similarly employed (as for example, in orchestrating swarms of drones or plotting efficient delivery routes). Intelligence might be defined as including an ability to learn and a capability for autonomous action. Others may wish to explore these characteristics to amplify or limit the comparison presented in this essay.

Readers interested in AI and learning may consult Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning* (Cambridge, MA: The MIT Press, 2016) which commends the "succinct definition" of learning offered by T.M. Mitchell in his book *Machine Learning* (New York, NY: McGraw-Hill, 1997), 96: "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E." It may be argued that bureaucracies and markets differ from machines in their reliance on humans to do this learning, but this is an observation about whether bureaucracies and markets are artificial, not whether they learn. It should be clear that they do learn, though much more slowly than machines. Furthermore, the artificiality of a composite market or bureaucracy is not negated by the presence of human components any more than the artificiality of a pair of shoes is negated by the shoemaker who made them or a human who wears them. (And do the shoes meet Goodfellow's definition of learning because as they are broken in their performance "improves with experience?")

Those interested in autonomy as a differentiating characteristic of machine intelligence may consider Kai-Fu Lee's distinction: "The developed world has largely replaced raw human muscle with high-powered machines. But while these machines are *automated* they are not *autonomous.* While they can repeat an action, they can't make decisions or improvise according to changing conditions." Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (Boston, MA: Mariner Books, 2018), 128-129. Bureaucracies and markets seem to meet Kai-Fu Lee's criteria of autonomy.

[13] This thought may be captured in different ways. For example, a systems biologist views computation as "The procedures for arriving at adaptive decisions based on approximate and noisy information ..." Krakauer, "The Computational Systems of the World," 351-354. A National Academy of Sciences report approaches the point from another direction: "[T]he challenges for massive data go beyond the storage, indexing, and querying that have been the province of classical database systems (and classical search engines) and, instead, hinge on the ambitious goal of inference. Inference is the problem of turning data into knowledge, where knowledge often is expressed in terms of entities that are not present in the data per se but are present in models that one uses to interpret the data." National Research Council, *Frontiers in Massive Data Analysis* (Washington, DC: The National Academies Press, 2013), 3.

[14] "As originally observed by Adam Smith, the economy is a complex system in which agents interact with one another to produce aggregate behavior whose emergent behaviors are qualitatively different from those of the individual actors." David C. Krakauer et al., "Intelligent Data Analysis of Intelligent Systems," in *Advances in Intelligent Data Analysis IX* (2010): 13, https://link.springer.com/book/10.1007/978-3-642-13062-5.

This attribute of opacity is discussed in Section Three.

[15] "[E]very time we figure out a piece of it, it stops being magical; we say, 'Oh, that's just a computation." Attributed to Rodney Brooks then Director of MIT's Artificial Intelligence Laboratory, in Jennifer Kahn, "It's Alive!," *WIRED*, March 1, 2002, https://www.wired.com/2002/03/everywhere/. This view is commonly taken to reflect constantly rising expectations, but this discussion suggests a more fundamental point: in practice we regard intelligence as emergent.

[16] A 90-minute documentary recounting this achievement is at "AlphaGo," DeepMind, https://deepmind.com/research/case-studies/alphago-the-story-so-far. See DeepMind, "AlphaGo – The Movie," *YouTube*, March 13, 2020, https://www.youtube.com/watch?v=WXuK6gekU1Y.

[17] Siraj Raval, "Move 37 Explained," https://www.youtube.com/watch?v=vl9BIIT7ovg. The discussion embodying this quotation begins at the three-minute mark.

[18] Cade Metz, "How Google's AI Viewed the Move No Human Could Understand," *WIRED*, March 14, 2016, https://www.wired.com/2016/03/googles-ai-viewed-move-no-human-understand/: "[T]his surprisingly skillful machine made a move that flummoxed everyone from the throngs of reporters and photographers to the match commentators to, yes, Lee Sedol himself. "That's a very strange move," said one commentator, an enormously talented Go player in his own right. "I thought it was a mistake," said the other."

[19] David Krakauer, the President of the Santa Fe Institute, sees auctions through the same lens, but slightly differently. His call for a science of "emergent engineering" argues that this is most likely to come "from the domains of biological and social life." He would presumably regard bureaucracies as classically engineered systems designed to minimize error "where collective dynamics are predictable and controllable." Auction mechanisms move away from "overly complicated, over-regulated centralized valuation mechanisms.... [they] predict and control a behavior rather than shoehorn complexity into regimes where the classical engineering axioms hold sway." David C. Krakauer, "Emergent Engineering: Reframing the Grand Challenge for the 21st Century," in David Krakauer (ed.), *Worlds Hidden in Plain Sight: The Evolving Idea of Complexity at the Santa Fe Institute, 1984-2019* (Santa Fe, NM: SFI Press, 2019), 349, 352, 354-55.

[20] Francis Spufford, *Red Plenty* (Minneapolis, MN: Graywolf Press, 2010) is an extended rumination on this point. Gunnar De Winter, "AI Playing the Stock Markets," *Medium*, June 10, 2020, https://medium.com/predict/ai-playing-the-stock-markets-c71429621372 provides a description of some real-world efforts to develop machine systems that predict the stock market.

[21] Famously, for example, in his novel *The Trial* (1937) in which the accused is confronted with a physical environment of unmarked passages and stairways that mirror proceedings without apparent cause, duration, purpose, or process. Thus for instance, "the legal records of this case, and above all, the actual charge sheets, were inaccessible to the accused and his counsel, consequently one did not know in general, or at least did not know with any precision, what charges to meet in the first plea; accordingly, it could only be by pure chance that it contained really relevant matter." Franz Kafka, Edwin Muir and Willa Muir tr., *The Trial* (New York, NY: Alfred A. Knopf, 1947), 144.

[22] "[B]ureaucracy is not the simple, uniform phenomenon it is sometimes made out to be." Wilson, *Bureaucracy*, ix and see for example 113ff, contrasting a McDonald's franchise with a Department of Motor Vehicles office.

[23] Sometimes also referred to as "deep neural networks." The method was pioneered by DeepMind, a British entity acquired by Google in 2014. For an authoritative description see the foundational paper by three originators of the approach, Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature* 521 (2015): 436-444, https://doi.org/10.1038/nature14539. Footnotes 35 and 37 and the accompanying text discuss generative adversarial networks and transformers as examples of recent cutting-edge applications of Deep Learning.

[24] "AI scientists in recent years have largely focused on the specific type of machine learning known as deep learning…" Erik J. Larson, *The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do* (Cambridge, MA: Harvard University Press, 2021), 134.

[25] Judea Pearl and Dana Mackenzie, *The Book of Why: The New Science of Cause and Effect* (New York, NY: Basic Books, 2018) offers a lucid extensive exploration of the contrast between correlative (or associative) and causal explanation. The examples in this paragraph are discussed in their introduction. Pearl and McKenzie are unequivocal as to present machine intelligence: "That's all that a deep-learning program can do: fit a function to data" (17). They believe that a "strong AI" is eventually "an achievable goal" (11). However, "In technical terms, machine-learning methods today provide us with an efficient way of going from finite sample estimates to probability distributions, and we still need to get from distributions to cause-effect relations" (362).

[26] There is of course much uncertainty and debate about how humans think. For example, Erik J. Larson argues that humans have a third capability, "abductive inference" and that this intuitive ability to conceive hypotheses is central to distinguishing systems like Deep Mind from human thought. Larson, *The Myth of Artificial Intelligence*, e.g. 25–6 and 99ff. Conversely, Judea Pearl and Dana Mackenzie take a narrow view. In *The Book of Why*, they assert that "human intuition" is simply "grounded in causal, not statistical, logic" (19). See also 189: "[H]uman intuition operates under the logic of causation, while data conform to

the logic of probabilities and proportions.” For our purposes it suffices to say that most humans regard themselves as causative thinkers and the absence of causative thought in machines, markets and bureaucracies provokes a sense of alienation in humans when they encounter these systems.

[27] Lindblom, *The Market* System, 40. And see 140ff on markets as mechanisms for establishing efficient prices.

[28] W. Richard Scott and Gerald F. Davis’s *Organizations and Organizing: Rational, Natural and Open Systems Perspectives* (United Kingdom: Routledge, 2006), 56, quoting DiMaggio and Powell who themselves credit March and Simon with this perception.

[29] Max Weber as quoted in Scott and Davis’s Organizations and Organizing: Rational, Natural and Open Systems Perspectives.

[30] Melanie Mitchell, “Abstraction and Analogy-Making in Artificial Intelligence,” arXiv preprint arXiv:2102.10717 (2021), https://arxiv.org/pdf/2102.10717.pdf.

[31] Judea Pearl, “The Limitations of Opaque Learning Machines,” in John Brockman (ed.), *Possible Minds: Twenty-Five Ways of Looking at AI* (Westminster, UK: Penguin Books, 2019), 17, 15-16.

[32] Goodfellow, Bengio, and Courville, *Deep Learning*, 95.

[33] “China’s use of AI-powered surveillance technologies to repress its Uyghur minority and monitor all of its citizens foreshadows how authoritarian regimes will use AI systems to facilitate censorship, track the physical movements and digital activities of their citizens, and stifle dissent. The global circulation of these digital systems creates the prospect of a wider adoption of authoritarian governance. But liberal democracies also employ AI for internal security and public safety purposes. More than half of the world’s advanced democracies use AI-enabled surveillance systems.” The National Security Commission on Artificial Intelligence, Final Report (Washington, DC: NSCAI, March 2021), 27.

[34] Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford, UK: Oxford University Press, 2014), 15.

[35] For descriptions and examples, see OpenAI’s sites “Generative Models,” OpenAI, https://openai.com/blog/generative-models/, “Dalle-E: Creating Images from Text,” OpenAI, https://openai.com/blog/dall-e/ and “Better Language Models and Their Implications,” OpenAI, https://openai.com/blog/better-language-models/.

[36] John Launchbury, “A DARPA Perspective on Artificial Intelligence,” *MachineLearning* (TechnicaCuriosa), March 19, 2017,

https://machinelearning.technicacuriosa.com/2017/03/19/a-darpa-perspective-on-artificial-intelligence/.

[37] The relationship between causal and correlative understanding in human beings is worthy of exploration as is the question of whether there are aspects of human thought beyond these two categories. Pearl and Mackenzie seem to argue for a single category, asserting that "human intuition is grounded in causal, not statistical, logic." Pearl and Mackenzie, *The Book of Why*, 19. See also 189: "[H]uman intuition operates under the logic of causation, while data conform to the logic of probabilities and proportions." "[H]uman intuition is grounded in causal, not statistical, logic.... [H]uman intuition operates under the logic of causation, while data conform to the logic of probabilities and proportions." Pearl and Mackenzie, *The Book of Why*, 19, 189. It is possible that the distinction between correlative and causative intelligence may turn out to be illusory or at least less meaningful than it presently seems. As described in footnote 35 and the accompanying text, machines are now programmed to use deep learning to produce outputs (for instance, generating functioning code or fairly convincing poetry) that appear to be creative." Human assessments about causation may be rationalizations of judgments that are actually correlative.

While Simon Baron-Cohn, a professor of psychology and psychiatry, sees human reasoning as pattern-seeking he also argues for a critical third capability: empathetic ability. In his book, *The Pattern Seekers: How Autism Drives Human Invention* (New York, NY: Basic Books, 2020), Baron-Cohn focuses on autism and describes it in a manner very like this paper's description of the correlative workings of artificial intelligences. His discussion emphasizes a deficit in empathy that characterizes autistic individuals. His footnote 47 to chapter 2 has potential implications for the ability of machines to learn empathy. The footnote describes studies in which "some autistic females... become 'obsessed' with films or novels... This may be one way in which they 'camouflage' their autism, effectively learning theory of mind from ... books or movies, even if implementing it in real time in the social world remains challenging" (p. 145). Kazuo Ishiguro's novel, *Klara and the Sun* (2021) presents a compelling portrait of a machine programmed to learn empathy. (The results are all the more complex because Ishiguro uses his art to induce readers to empathize with the machine). Sayaka Murata's novel *Convenience Store Woman* (New York, NY: Grove Press, English translation 2018) portrays a young woman who has difficulties with empathy and relationships but learns a way of relating as a result of training and imitation in the highly ritualized interactions expected of her as a convenience store clerk. Together these novels suggest that some empathetic ability may be acquired by machines and in bureaucracies through observing and copying patterns.

[38] We don't, for example, perceive ultraviolet light. Animals and machines are particularly valuable complements to humans when they harvest data that otherwise escapes us.

[39] George Dyson, *Turing's Cathedral: The Origins of the Digital Universe* (New York, NY: Pantheon Books, 2012), 3: "To a digital computer, the only difference that makes a difference is the difference between zero and one." Of course, human assimilation of inputs is also reductionist, as when neurons at the back of an eye respond to photons or sound is transmitted by variations in air pressure detected by membranes in the ear. The comparison may warrant more extended discussion in other contexts.

[40] Emphasizing "the limitations of the human mind," Pedro Domingos writes that "Most of the brain's hardware (or rather, wetware) is devoted to sensing and moving and to do math we have to borrow parts of it that evolved for language. Computers have no such limitations..." Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York, NY: Basic Books, 2015), 31. This perspective about computation is not unlike Marshall McLuhan's perception that modern media have been constructed to amplify the limited capacities of humans to communicate. Marshall McLuhan, *Understanding Media: The Extensions of Man* (New York, NY: McGraw-Hill, 1964).

[41] The most notable modern defense of markets by Frederick Hayek "did not rest primarily upon the supposed optimum attained by them, but rather upon the .... [c]omputational limits of human beings." Herbert A. Simon, *The Sciences of the Artificial* (Cambridge, MA: The MIT Press, 1st ed., 1969; all references are to the 3rd ed., 1996), 34, citing Hayek, "The Use of Knowledge in Society," 519-530. See also Simon's comments at 41: "We can think of a decision as produced by executing a large computer program... No single person or group need be expert on all aspects of the decision."

Henry Farrell displays keen insight about Simon, Hayek and Charles Lindblom's contributions in an interview, Sophie Roell, "The Best Books on the Politics of Information recommended by Henry Farrell," *Five Books*, https://fivebooks.com/best-books/politics-of-information-henry-farrell/. Remarkably, Davis and Scott, though they are appreciative of other work by Simon, make no reference to *The Sciences of the Artificial* but briefly reference Weick's assertion that "human beings organize primarily to help them reduce the information uncertainty they face in their lives." Scott and Davis, *Organizations and Organizing*, 104. Malone, *Superminds* provides a modern argument to the same effect.

[42] Thomas Hobbes, *Leviathan* (1651), Ch. XVI. Hobbes' use of the word "artificial" is quoted here to indicate the novelty of the collective activities he was pondering in the seventeenth century. As with any vocabulary used long ago, the term came embedded in a culture and a body of thought distant from our own. For a discussion of some intricacies, noting also that Hobbes had several different concepts of artificiality and that these were intertwined with developing concepts of principals and agents, see David Copp, "Hobbes on Artificial Persons

and Collective Actions," *The Philosophical Review* 89, no. 4 (October 1980): 579-606, https://www.jstor.org/stable/2184737?origin=crossref&seq=1.

[43] Max Weber, *Law*, 231, as quoted in Reinhard Bendix, *Max Weber: An Intellectual Portrait* (Berkeley, CA: University of California Press, 1978), 427. The quotation continues, "In the place of the old-type ruler who is moved by sympathy, favor, grace and gratitude, modern culture requires for its sustaining external apparatus, the emotionally detached and hence rigorously 'professional' expert." "Separation of office and incumbent" is fundamental to bureaucratic organization. Bendix, *Max Weber: An Intellectual Portrait*, 424. "Before this time, control of government and markets depended on personal relationships and face to face interactions." James R. Beniger, *The Control Revolution: Technological and Economic Origins of the Information Society* (Cambridge, MA: Harvard University Press, 1986), 7.

[44] "Traditional measures [like the foot] were 'human' in many respects. They were expressive of man and has work; they depended at times upon his will, which depended upon his character and attitude toward fellow humans... [They were replaced by] 'dehumanizing' measures [like the meter]..." Witold Kula and R. Szreter (tr.), *Measures and Men* (1986), 122-123. Hannah Arendt decries "earth alienation," which is "perhaps clearest in the development of the new science's most important mental instrument, the devices of modern algebra, by which mathematics 'succeeded in freeing itself from the shackles of spatiality, that is from geometry which, as the name indicates, depends on terrestrial measures and measurements." 264-265, quoting Edwin Arthur Burtt, 44. Arendt and Kula might have a yet stronger example in time-keeping, as the world moved from local time-keeping where noon marked the sun's highest point, to standardized time zones where the sun's zenith might be a half hour before or after 12 o'clock.

[45] Robert L. Heilbroner, *The Worldly Philosophers: The Lives, Times and Ideas of the Great Economic Thinkers* (New York, NY: Touchstone, 4th ed., 1972), 27, 31. Francis Spufford offers a memorable description: "Then the markets and the things made turned alike into commodities, and the motion of society turned into a kind of zombie dance, a grim cavorting whirl in which objects and people blurred together till the subjects were half alive and the people were half dead." Spufford, *Red Plenty*, 66.

[46] David Graeber, *The Utopia of Rules: On Technology, Stupidity, and the Secret Joys of Bureaucracy* (Hoboken, NJ: Melville House, 2015), 3-4. Graeber offers data on the frequency of the word "bureaucracy" in support of this proposition. (Peak use of the word was in 1973 and since then has been in "a slow but inexorable decline.")

[47] Lindblom, *The Market* System, 37. The 1911 Encyclopedia Britannica provides a marker of the novelty of large international markets when it matter of factly observes "What is true of the cotton market is also true to some extent of all

markets, though few markets are so highly organized or show such large transactions as those for cotton." *Encyclopedia Britannica* Vol. 17 (1911): 731.

[48] Thus, Arthur Clarke's famous "law" that "Any sufficiently advanced technology is indistinguishable from magic," *Profiles of the Future: An Inquiry into the Limits of the Possible* (London, UK: Orion Publishing Group, 1962).

[49] Douglas Adams, *The Salmon of Doubt*, as quoted at "*The Salmon of Doubt* Quotes," Goodreads, https://www.goodreads.com/work/quotes/809325-the-salmon-of-doubt-hitchhiking-the-galaxy-one-last-time. Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control* (Westminster, UK: Penguin Books, 2019) suggestively supposes that when machine driven cars are perfected, "the reckless and antisocial act of driving a car oneself may be banned altogether" (p.67). Herbert Simon observes "A forest may be a phenomenon of nature; a farm certainly is not." Herbert A. Simon, *The Sciences of the Artificial*, 3.

[50] "[W]e use them without noticing them." W. Brian Arthur, *The Nature of Technology: What It Is and How it Evolves* (New York, NY: Free Press, 2009), 150.

[51] "The world we live in today is much more a man-made or artificial world than it is a natural world... A forest may be a phenomenon of nature; a farm certainly is not." Herbert A. Simon, *The Sciences of the Artificial*, 2-3.

[52] "[M]ost Greek literature was meant to be heard or even sung – thus transmitted orally – and there was a strong current of distaste for the written word even among the highly literate...We cannot assume an immediate (and modern) reverence for the written word that was elusive even as late as the fourth century." Rosalind Thomas, *Literacy and Orality in Ancient Greece* (Cambridge, UK: Cambridge University Press, 1992), 3, 47-48. When musical notation supported the evolution of polyphonic harmonies, "the monks did not have to memorise the new music, which was much more complex than the traditional plainchant. To help them read the notes, they were allowed to bring candles into the choir. Older monks deplored this innovation, which they said would rot the novices' memories." Seb Falk, *The Light Ages: The Surprising Story of Medieval Science* (New York, NY: W.W. Norton & Company, 2020), 66.

[53] In a celebrated article thirty years ago, Mark Weiser observed that the computer, then just beginning to proliferate, would soon cease to be novel. "The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.... Such a disappearance is a consequence not of technology, but of human psychology." As prior examples, he offered writing, "perhaps the first information technology," and electric motors. Mark Weiser, "The Computer for the 21st Century," *Scientific American* (September 1991), https://www.lri.fr/~mbl/Stanford/CS477/papers/Weiser-SciAm.pdf.

[54] "Thanks to the ever-increasing availability of more massive datasets, increased computing power, improved machine learning algorithms, and improved open-source code libraries and software development frameworks, things that used to be nearly impossible, such as automated facial recognition, are now possible. Programs that used to have terrible performance, such as automatic translation, now have significantly better performance. Finally, AI systems that used to be extremely expensive to develop, such as imagery classification, are often now affordable and sometimes even cheap." Greg Allen, *Understanding AI Technology* (Washington, DC: Joint Artificial Intelligence Center, 2020), 10-11, https://www.ai.mil/docs/Understanding%20AI%20Technology.pdf.

[55] Frank Pasquale, *The Black Box Society: The Secret Algorithms that Control Money and Information* (Cambridge, MA: Harvard University Press, 2015) is a leading book-length statement of this concern. Pasquale argues that "the contemporary world ... resembles a one-way mirror, Important corporate actors have unprecedented knowledge of the minutiae of our daily lives, while we know little to nothing about how they use this knowledge ..." (9).

[56] Launchbury, "A DARPA Perspective on Artificial Intelligence."

[57] Marvin Minsky, a pioneer in artificial intelligence, observed that "[y]ou ride as a passenger in someone else's car not knowing how that driver works. Most strange of all you drive your body and your mind, not knowing how your own self works." Minsky, *The Society of Mind*, 56. Minsky offers the processes of everyday speech as examples: "What kinds of agents choose your words so that you can express the things you mean? How do these words get arranged into phrases and sentences, each connected to the next?" Minsky, *The Society of Mind*, 50. Daniel Kahneman writes that "[m]ost impressions and thoughts arise in your conscious experience without your knowing how they got there. You cannot trace how you came to the belief that there is a lamp on the desk in front of you, or how you detected a hint of irritation in your spouse's voice on the telephone ..." Daniel Kahneman, *Thinking, Fast and Slow* (Westminster, UK: Penguin, 2011), 4. "[W]e do not have access to how our minds work simply by studying our own consciousness." Jonathan St B.T. Evans, *Thinking and Reasoning: A Very Short Introduction* (Oxford, UK: Oxford University Press, 2017), 10.

[58] This observation may also cast light on some of the challenges of self-driving cars. Human drivers will injure other human drivers and bystanders, but it is presumed that they avoid this not only because of legal consequences, but also because they empathize with those they put at risk. A machine that merely calculates risk and penalty is objectionable, no matter the rationality of the calculation.

[59] Indeed, in some instances improvements in sensitivity may involve embracing more opacity, as for example when we assure a right to trial by jury, a step that makes judgments more broadly sensitive, but also more opaque.

[60] See Natasha Dow Schüll, *Addiction by Design: Machine Gambling in Las Vegas* (Princeton, NJ: Princeton University Press, 2014) and Richard Danzig, "An Irresistible Force Meets a Moveable Object: The Technology Tsunami and the Liberal World Order," in R. Nicholas Burns, *The World Turned Upside Down: Maintaining American Leadership in a Dangerous Age* (Aspen Institute, 2017), 101ff.

[61] Section 5 also discusses another path to improving the match between human and machine intelligence. Bureaucracies, markets and their clients employ intermediaries like "customer support" where trained humans help untrained humans express themselves in terms comprehensible to the machine. Large service sectors like law, accounting and brokerage firms exist to bridge the gap between ordinary humans and the reductionist and opaque bureaucracies and markets.

[62] Conjoint human and machine operations are extensively and optimistically explored in Malone, *Superminds*.

[63] A classic paper demonstrated that resumes with "very white sounding names" yielded 50 percent more invitations to interview than identical resumes with "very African-American sounding names." Marianne Bertrand and Sendhil Mullainathan, "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review* 94, no. 4 (September 2004): 991ff.

[64] John Cartlidge et al., "Too Fast Too Furious: Faster Financial-Market Trading Agents Can Give Less Efficient Markets," in *ICAART-2012: Proceedings of the Fourth International Conference on Agents and Artificial Intelligence,* 2 (Agents) (2012), https://research-information.bris.ac.uk/en/publications/too-fast-too-furious-faster-financial-market-trading-agents-can-g.

[65] The extraordinary work of Tversky and Kahneman documented a great variety of recurring cognitive errors. Their foundational papers were Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," Science 185, no. 4157 (1974): 1124-1131 and Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47, no. 2 (1979): 263-91. Recent renderings of this work may be found in Kahneman, *Thinking, Fast and Slow* and Daniel Kahneman, Olivier Sibony, and Cass R. Sunstein, *Noise: A Flaw in Human Judgment* (New York, NY: Little, Brown, and Company, 2021).

[66] Charlie Munger, "The Psychology of Human Mismanagement," https://fsmisc.s3.ca-central-1.amazonaws.com/great-

talks/The+Psychology+of+Human+Misjudment.pdf provides a wonderful inventory of how organizations suffer from and often encourage human misjudgment. The Defense Intelligence Agency provided a more focused, less anecdotal, example from an experiment in which it showed that as key sources of information were removed, humans only modestly adjusted confidence in their judgments when compared with machines. The results were said to validate the general rule that "[o]verconfidence is human and a particular trait among highly-functioning expert humans, one that machines don't necessarily share." Patrick Tucker, "Artificial Intelligence Outperforms Human Intel Analysts In a One Key Area," *Defense One*, April 29, 2020, https://www.defenseone.com/technology/2020/04/artificial-intelligence-outperforms-human-intel-analysts-one-key-area/165022/.

[67] See for example, Maryam Kouchaki and Isaac H. Smith, "The Morning Morality Effect: The Influence of Time of Day on Unethical Behavior," *Psychological Science* 1 (January 2014): 95-102, https://pubmed.ncbi.nlm.nih.gov/24166855/.

[68] I discuss these in Richard Danzig, "Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority" (Center for a New American Security, May 30, 2018), 13-16. The shooting down of an Iranian civilian airline because of the mistaken conclusion of a U.S. Navy Captain in the face of a computer system conclusion that it was not a military aircraft is a much-discussed instance. Danzig, "Technology Roulette," 14-16. A crash of Boeing's Max 737 commercial airliner some 120 days after the publication of "Technology Roulette" provided a further illustration of this phenomenon. The software system provided recommendations that were accurate, but some pilots could not assimilate them. Alan Levin, "Boeing Failed to Predict that a Slew of 737 Max Warning Alarms Would Confuse Pilots, Investigators Say," *TIME*, September 26, 2019, https://time.com/5687473/boeing-737-alarm-system/. See also, William Langewiesche's account of another crash, "The Human Factor," *Vanity Fair*, October 2014, https://archive.vanityfair.com/article/2014/10/the-human-factor.

[69] Malone, in *Superminds*, writes: "Many people assume that computers will eventually do most things by themselves and we should put 'humans in the loop' in situations where they're still needed. But I think it's more useful to realize that most things now are done by groups of people, and we should put computers into those groups ... [W]e should move from thinking about putting humans in the loop to *putting computers into the group*" (p.75; italics in the original). John Gruber et al. put this view more technically, arguing for "not simply asking the aggregate questions of substitution between humans and machines" but instead studying "the detailed interaction between humans and AI tools where different aspects of the production function can be either augmented or replaced by technology" (p.2)/ Their study of how machine intelligence combines with expert advice in recommending health care insurance plans is an illuminating study of partnership. John Gruber et al., "Managing Intelligence: Skilled Experts and AI in

Markets for Complex Products" (National Bureau of Economic Research, June 2020), https://www.nber.org/papers/w27038. Among other things they observe a "general phenomenon in which AI-based tools are particularly good at accounting for quantifiable aspects of decisions but rarely account for the universe of welfare relevant aspects of a decision" (15). Human expertise fills that gap.

[70] Strengths can have adverse systemic consequences as when the speed of some machines spawns instability or inequity. Neil F. Johnson et al., "Abrupt Rise of New Machine Ecology Beyond Human Response Time," *Scientific Reports* 3 (2013), https://www.nature.com/articles/srep02627; Cartlidge et al., "Too Fast Too Furious;" and Neil F. Johnson, "To Slow, or Not to Slow? New Science in Subsection Networks," arXiv preprint arXiv:1706.08667 (2017), https://arxiv.org/pdf/1706.08667.pdf.

[71] "[C]omputer science has traditionally been all about thinking deterministically, but machine learning requires thinking statistically." Domingos, *The Master Algorithm*, 9.

[72] Compliance officers are frequently designated in an effort to decrease deviant performance in bureaucracies and markets. For example, commodities markets require traders to operate through brokerage firms and require the firms to "diligently supervise the handling by its partners, officers, employees and agents (or persons occupying a similar status or performing a similar function) of all commodity interest accounts…." 17 CFR § 166.3 – Supervision at https://www.law.cornell.edu/cfr/text/17/166.3.

[73] For example, periodic reporting requirements, generally accepted accounting procedures (GAAP), immediate compulsory disclosure of all material information and a welter of related requirements are imposed on the 2800 companies that comprise the New York Stock Exchange. For a brief layperson's summary see Cam Merritt, "Disclosure Requirements for the New York Stock Exchange," Zacks, https://finance.zacks.com/disclosure-requirements-new-york-stock-exchange-3854.html. See also, "What is GAAP," Accounting.com, September 7, 2021, https://www.accounting.com/resources/gaap/.

[74] Thomas Smith, "The Explosion of New Architectures Is Fundamentally Changing Computing: Killing Von Neumann," *Medium*, February 10, 2020, https://medium.com/swlh/the-explosion-of-new-architectures-is-fundamentally-changing-computing-f69b7faae89d.

[75] "Goldman Sachs created a security so opaque and complex that it would remain forever misunderstood by investors and rating agencies: the synthetic subprime mortgage bond-backed CDO or collateralized debt obligation." Michael Lewis, *The Big Short: Inside the Doomsday Machine* (New York, NY: W.W. Norton & Company, 2010), 72. Lewis' lucid layperson's account observes:

An investor who went from the stock market to the bond market was like a small, furry creature raised on an island without predators removed to a pit full of pythons. It was possible to get ripped off by the big Wall Street firms in the stock market, but you really had to work at it.... The stock market was not only transparent, but heavily policed.... The presence of millions of small investors had politicized the stock market. It had been legislated and regulated to at least seem fair.

The bond market, because it consisted mainly of big institutionalized investors, experienced no similarly populist political pressures. Even as it came to dwarf the stock market, the bond market eluded serious regulation.... Bond technicians could dream up ever more complicated securities without worrying too much about government regulation ... (61–2).

[76] The U.S. Securities and Exchange Commission, the Federal Trade Commission, the Federal Drug Administration, the Federal Aviation Administration, the National Highway Traffic Safety Administration and other such entities are examples of regulatory agencies that require transparency from the corporate bureaucracies they oversee. All have developed disclosure requirements not only to benefit consumers who are making market decisions. They also, perhaps more so, require this transparency because it is a prerequisite for their effective oversight. Similarly, the Administrative Procedure Act (requiring for example provision and retention of public records and statements of reasons for a decision) protects those affected by bureaucratic decisions, but just as significantly it provides the basis for courts to oversee bureaucracies.

[77] Heather M. Wojton et al., "Operational Testing of Systems with Autonomy" (Institute for Defense Analyses, March 2019), https://www.ida.org/research-and-publications/publications/all/o/op/operational-testing-of-systems-with-autonomy and Laura Freeman, "Test and Evaluation for Artificial Intelligence," Insight 23, no. 1 (March 2020): 27-30 argue for this perspective. Freeman observes that a system "is not 'final' during the deployment. It changes dynamically as algorithms learn and the environment changes..." (28). She continues:

For fielded systems not leveraging AI, follow-on testing or regression testing occurs only if the system receives any new improvement and then only if the improvement is substantial enough to impact performance. Integrating AI into a system continually evolves our knowledge of the system's ability to accomplish tasks. For systems learning off-line there is the chance they will encounter new data in the field producing an unexpected or new outcome. Additionally, systems allowed to learn online continually evolve as the algorithms process new data. Therefore, testing and evaluation must be continuous (28-9).

[78] Dan Geer, email to the author, June 24, 2021. As policy-makers concerned with cybersecurity have occasion to note, source code is policed by buyers and sellers operating in markets, rather than by public regulation.

[79] After saying that "AI has been around for decades. So, why has everyone been talking about it constantly in recent years?," Greg Allen lists "More Massive Datasets," as the first transformative factor in the development of machine intelligence. Allen, "Understanding AI Technology," 9. See also p.7: "Algorithms and computing hardware are also important, but nearly all ML systems run on commodity computing hardware, and nearly all of the best algorithms are freely available worldwide. Hence, having enough of the right data tends to be the key."

[80] James C. Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed* (New Haven, CT: Yale University Press, 1998), 2. See also Kula, *Measures and Men*.

[81] Marion Fourcade and Kieran Healy, "Seeing Like a Market," *Socio-Economic Review* 15, no. 1 (January 2017): 13.

[82] Fourcade and Healy, "Seeing Like a Market," 10, 19.

[83] Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (London, UK: Profile Books, 2019), 2, 4, 211.

[84] Zuboff, *The Age of Surveillance Capitalism*, 86–7. Google, in Zuboff's accounting, became preeminent in "hunting, capturing and transforming" data… Google's founders transformed the youthful Dr. Jekyll into a ruthless, muscular Mr. Hyde determined to hunt his prey anywhere, anytime, irrespective of others self-determining aims." Zuboff, *The Age of Surveillance Capitalism*, 80-81. See also, Bruce Schneier, *Data and Goliath: The Hidden Battles to Capture Your Data and Control Your World* (New York, NY: W.W. Norton & Company, 2015), documenting the proposition that "Corporations and governments alike have an insatiable appetite for our data…" (7). Note however that Google recently announced a campaign to eliminate cookies in its systems. See "Building a more private web: A path towards making third party cookies obsolete," *Chromium Blog*, January 14, 2020, https://blog.chromium.org/2020/01/building-more-private-web-path-towards.html.

[85] As previously noted, for example, machine translation was given a boost by harvesting public records of the European Union that provided parallel texts in English, French and German. This is discussed in Bostrom, *Superintelligence*, 15. Facial recognition has advanced by the compilation of faces photographed "in the wild." See for example, "UTKFace: Large Scale Face Dataset," *GitHub*, https://susanqq.github.io/UTKFace/. Synthetic data can be used for training and testing. See for example, Laboratory for Information and Decision Systems, "The

Real Promise of Synthetic Data," *MIT News*, October 16, 2020, https://news.mit.edu/2020/real-promise-synthetic-data-1016.

[86] Three CSET authors caution that "[I[t is sometimes suggested that government organizations will only be able to benefit from the AI revolution if they can digitize, clean, and label large amounts of data. While this suggestion holds merit, it is inaccurate to think of all progress in AI as contingent upon such conditions… AI innovation in the government sector (and beyond) can still happen without massive investments in big data infrastructures." Husanjot Chahal, Helen Toner, and Ilya Rahkovsky, "Small Data's Big AI Potential" (Center for Security and Emerging Technology, September 2021), 5, https://cset.georgetown.edu/publication/small-datas-big-ai-potential/. Their contention that increasing amounts can be done with sparse data sets has merit, but, as they acknowledge, shouldn't be taken to negate the benefit of large data sets.

[87] These bureaucracies of course lack a profit motive. In national security bureaucracies, inertial resistance is especially strong because security concerns can justify hiding data lest adversaries learn too much about us or about our sources and methods for understanding them.

[88] While this paper was in its final draft, the Department of Defense improved its focus with the promulgation of a memorandum from the Deputy Secretary of Defense, *Creating Data Advantage* (Washington, DC: Department of Defense, May 5, 2021), https://media.defense.gov/2021/May/10/2002638551/-1/-1/0/DEPUTY-SECRETARY-OF-DEFENSE-MEMORANDUM.PDF, stating that "Data is a strategic asset… [L]eaders must ensure all DoD data is visible, accessible, understandable, linked, trustworthy, interoperable, and secure." The vice chairman of the Joint Chiefs of Staff followed with implementing memos intended to catalyze greater sharing. Jackson Barnett, "Hyten Signs New Requirements to Ensure Military Services Make Data Accessible," *FedScoop*, June 23, 2021, https://www.fedscoop.com/hyten-signs-new-data-requirements-accessibility/.

[89] *U.S. Department of Homeland Security Artificial Intelligence Strategy* (Washington, DC: Department of Homeland Security, December 3, 2020), 7, https://www.dhs.gov/sites/default/files/publications/dhs_ai_strategy.pdf.

[90] National Security Commission on Artificial Intelligence, "Final Report," 24.

[91] National Security Commission on Artificial Intelligence, "Final Report," 63.

[92] National Security Commission on Artificial Intelligence, "Final Report," fn. 13, 71.

[93] U.S. Department of Defense, *Executive Summary: DOD Data Strategy* (Washington, DC: Department of Defense, September 30, 2020),

https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF.

[94] The Commandant of the Marine Corps offered this unvarnished assessment of his service when he assumed command two years ago:

> We do not currently collect the data we need systematically, we lack the processes and technology to make sense of the data we do collect, and we do not leverage the data we have to identify the decision space in manning, training, and equipping the force. Where we have individual leaders and organizations that are trying to adopt the best practices in data science and data analytics, it is often accomplished through the heroic efforts of a few individuals rather than the organized and sustained effort required to transform how we sense, make sense, and act.

38th Commandant of the Marine Corps, *Commandant's Planning Guidance* (Washington, DC: Department of the Navy, July 16, 2019), 14, https://www.hqmc.marines.mil/Portals/142/Docs/%2038th%20Commandant%27s%20Planning%20Guidance_2019.pdf?ver=2019-07-16-200152-700.

The former head of data operations for two years at Project Maven, then the leading Department of Defense artificial intelligence program, wrote:

> One insight stands out from lessons learned over four years of bureaucratic bushwhacking. It speaks to an unfortunate cultural tendency, something that senior leaders could abolish if they realized how pervasive and problematic it is: Many data brokers in the defense community still treat data like currency. This was a No. 1 blocker at Project Maven, hands down. The team lost months of time waiting for partner organizations to release data from archives or grant access to data streams. The timeline for some of the team's requests, even when initiated by senior leaders, could be measured not in days or weeks, but by the passing of seasons.

Jaim Coddington, "Want better AI for the DOD? Stop treating data like currency," *C4ISRNET*, May 7, 2021, https://www.c4isrnet.com/opinion/2021/05/07/want-better-ai-for-the-dod-stop-treating-data-like-currency/.

[95] National Security Commission on Artificial Intelligence, "Final Report," 62.

[96] The problem is discussed in Nithya Sambasivan et. al, ""Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI," *Google Research* (2021), https://research.google/pubs/pub49953/. Note as an example, that Goodfellow, Bengio, and Courville, *Deep Learning*, referred to as an exemplary basic text in this paper and throughout the field of AI, credits data bases only as something that "has made machine learning much easier…" (19), while "The increase in model size over time, due to the availability of faster CPUs

the advent of general purposes GPUs..., faster network connectivity and better software infrastructure for distributed computing, is one of the most important trends in the history of deep learning" (21).

[97] The professionally focused Turing Award is more sensitive to this point, approximately once each decade honoring leaders who have contributed "in theory and practice" to the development of databases.

[98] Tim Hwang describes how this disadvantages the US when competing with authoritarian regimes that "may have an easier time acquiring data for training ML applications, in part because they may already maintain an existing infrastructure for ubiquitous surveillance that enables easy data collection. Moreover, there may be no strong legal mechanisms to protect citizen privacy or prevent the state from compelling companies to provide access to their data." Tim Hwang, "Shaping the Terrain of AI Competition" (Center for Security and Emerging Technology, June 2020), 3, https://cset.georgetown.edu/publication/shaping-the-terrain-of-ai-competition/.

[99] In the "12- to 14-week-old embryo, nerve cells are proliferating at the rate of about 15 million per hour.... After about 18 months of age, no more neurons are added, and the aggregation of cell types into distinct regions is roughly complete." Sandra Ackerman, *Discovering the Brain* (Washington, DC: National Academies Press, 1992), Chapter 6, https://www.ncbi.nlm.nih.gov/books/NBK234146/. Ackerman observes that "pruning of excess connections—clearly a process of great importance for the shape of the mature brain—continues for years."

[100] "[T]he amnesia of infancy ... makes us assume that all our wonderful abilities were always there inside our minds, and we never stop to ask ourselves how they began and grew." Minsky, *The Society of Mind*, 21. "We regard our knowledge as 'obvious' only because we cannot remember how hard it was to learn." Minsky, *The Society of Mind*, 27.

[101] I am indebted to Chris Inglis and Andrew Burt for insights leading to this formulation.

[102] Samuel Gibbs, "Elon Musk: Artificial Intelligence Is Our Biggest Existential Threat," *The Guardian*, October 27, 2014, https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat. See also, Catherine Clifford, "Elon Musk: 'Mark my words — A.I. is far more dangerous than nukes,'" *CNBC*, March 13, 2018, https://www.cnbc.com/2018/03/13/elon-musk-at-sxsw-a-i-is-more-dangerous-than-nuclear-weapons.html.

[103] Bostrom, *Superintelligence*, 95. Bostrom's view is concisely conveyed in Nick Bostrom, "What happens when our computers get smarter than we are?," *TED*, March 2015,

https://www.ted.com/talks/nick_bostrom_what_happens_when_our_computers_get_smarter_than_we_are/transcript?language=en. A TED talk by Sam Harris further articulates these concerns: Sam Harris, "Can we build AI without losing control over it?," *TED*, June 2016, https://www.ted.com/talks/sam_harris_can_we_build_ai_without_losing_control_over_it.

[104] Bostrom, *Superintelligence*, 95.

[105] Bostrom, *Superintelligence*, 94. Quotations are from entries on Table 8. Bostrom is also concerned with the possibility that such a machine could "persuade its gatekeepers to let it out." Daniel Suarez's novel *Daemon* (London, UK: Quercus, 2009) provides a vivid account of how a computer can be programmed to manipulate humans by blackmailing them, hacking their bank accounts and many other means.

[106] A more extensive discussion might explore other constraints on a conceivable super-intelligent artificial general intelligence (AGI). For example, omniscience seems unachievable because data cannot always be comprehensive. Model builders repeatedly learn this lesson as variables either unknown or unmeasured intrude on otherwise convincing calculations. As another example, concerns about AGI commonly treat omniscience as though it were synonymous with omnipotence. Harris says ""we are in the process of building some sort of God" at Harris, "Can we build AI without losing control over it?" (Minute 13). (Not all discussions make this error. Russell, for example, notes this shortfall and others described in this footnote, but argues that it is not of great operational significance. Russell, *Human Compatible*, 36-39.)

Furthermore, even a perfectly intelligent, omniscient and all-powerful superintelligence would be limited by the intractable difficulties of multi-objective optimization. No matter how capable, an intelligence cannot simultaneously optimize everything. After reaching a Pareto boundary, it would need to choose which strategies and tactics to maximize. Accordingly, no one intelligence (even a machine intelligence) can assuredly dominate the others.

Additionally, there are problems that are impervious to intelligence either as a matter of theory or because their solution requires exponential time and computational resources. Humans are especially difficult to predict and control, in part because they are adaptive. (Human beings "have a tendency to change as soon as we have come to understand them." Krakauer, "Complexity: Worlds Hidden in Plain Sight," 231.) "Will there be a giant leap forward in proactive social engineering allowing the orchestration of precisely engineered individual- or societal-level outcomes? Probably not. Two reasons why this is unlikely are the stochastic (random), rather than deterministic, nature of human behavior and the stochasticity in the process by which behavior combines to produce society." Jessica Flack, "Introduction: Social and Economic Engineering," in David Krakauer and Caitlin McShea, *InterPlanetary Transmissions: Proceedings of the Santa Fe*

Institute's First InterPlanetary Festival: Genesis (Santa Fe, NM: SFI Press, 2019), 218, http://c4.santafe.edu/resources/FlackIPIntroSFIPress.pdf.

[107] Von Mises, *Bureaucracy*, explores the contrast at length (though Mises does not employ this metaphor). Calabresi and Bobbitt note that societies often move 'with desperate grace from one approach to another…" as they become disenchanted with the values or methods inherent in each. Guido Calabresi and Philip Chase Bobbitt, *Tragic Choices* (New York, NY: W.W. Norton & Company, 1978), 198.

[108] See generally Lindblom, *The Market System*, especially pages 140ff. And see Henry Farrell's discussion of Lindblom in Roell, "The Best Books on the Politics of Information recommended by Henry Farrell."

[109] Jane Jacobs' neglected classic, *Systems of Survival: A Dialogue on the Moral Foundations of Commerce and Politics* (New York, NY: Random House, 1992) explores differences in norms between what she calls the "commercial syndrome" and the "guardian syndrome." She associates the latter with military organizations, but "[a]ll the occupations associated with it are not heroic. A lot consists of humdrum bureaucratic work in government" (29).

[110] Von Mises, *Bureaucracy*.

[111] Markets "force costs information on every chooser, for each chooser must pay to effect a choice." Lindblom, *The Market System*, 142.

[112] Beniger describes Weber's concept of rationalization as centered on "one essential idea: control can be increased not only by increasing the capability to process information but also by decreasing the amount of information to be processed….. In short, rationalization may be defined as the destruction or ignoring of information in order to facilitate its processing…. One example from within bureaucracy is the development of standardized paper forms."

[113] See for example the discussion of *Sullivan v. O'Connor* in Richard Danzig and Geoffrey R. Watson, *The Capability Problem in Contract Law: Further Readings on Well-Known Cases* (St. Paul, MN: Foundation Press, 2004), 5ff.

[114] David Krakauer, the President of the Santa Fe Institute, sees auctions through the same lens, but slightly differently. His call for a science of "emergent engineering" argues that this is most likely to come "from the domains of biological and social life." He would presumably regard bureaucracies as classically engineered systems designed to minimize error "where collective dynamics are predictable and controllable." Auction mechanisms move away from "overly complicated, over-regulated centralized valuation mechanisms…. [they] predict and control a behavior rather than shoehorn complexity into regimes where the classical engineering axioms hold sway." Krakauer,

"Emergent Engineering: Reframing the Grand Challenge for the 21st Century," 349, 352 and 354–55.

[115] They are responsive to long-term changes, as for example to changes in attitudes towards racial equality or view about the significance of gender or sexual preference.

[116] The New York Stock Exchange is owned by The Intercontinental Exchange (ICE), "an American Fortune 500 company formed in 2000 that operates global exchanges, clearing houses and provides mortgage technology, data and listing services. The company owns exchanges for financial and commodity markets, and operates 12 regulated exchanges and marketplaces." See "Intercontinental Exchange," Wikipedia, accessed September 2021, https://en.wikipedia.org/wiki/Intercontinental_Exchange.

[117] Google was founded in 1998 as a search company using human programmed algorithms. In 2015 Google introduced a machine learning component, "RankBrain," into its system.

[118] "[W]e find ourselves in a different world, in which the scarce resource is not the capacity to publish, but the capacity to pay attention." Roell, "The Best Books on the Politics of Information recommended by Henry Farrell." "More than anyone else, Google and Facebook have learned to exploit the stream of data released by billions of users to produce marketable consumer targeting. Google is essentially an advertising company, with ad revenue just shy of 90% of its total as of 2014." Fourcade and Healy, "Seeing Like a Market," 16. Herbert Simon's "Designing Organizations for an Information-Rich World" in M. Greenberger (ed.), *Computers, Communications, and the Public Interest* (Baltimore, MD: The Johns Hopkins Press, 1971), https://digitalcollections.library.cmu.edu/awweb/awarchive?type=file&item=33748 is widely cited for its insight that while historically information scarcity was a dominant problem, in the modern age information abundance might make attention from users the more precious resource. See generally "Attention Economy," Wikipedia, accessed September 2021, https://en.wikipedia.org/wiki/Attention_economy.

[119] A suggestive essay by Tim O'Reilly, "Open Data and Algorithmic Regulation," in Brett Goldstein and Lauren Dyson (eds.), *Beyond Transparency: Open Data and the Future of Civic Innovation* (San Francisco, CA: Code for America Press, 2013) argues that reputational ratings are examples of feedback that can enhance the flexibility and effectiveness of bureaucratic regulations. See also, Nello Cristianini and Teresa Scantamburlo, "On Social Machines for Algorithmic Regulation," *AI & Society* 35 (2020): 645-662, https://doi.org/10.1007/s00146-019-00917-8.

[120] Tim Wu provides an admirable discussion focused however on judicial decisions demanding more creativity than required for lower level bureaucratic procedures. Tim Wu, "Will Artificial Intelligence Eat the Law? The Rise of Hybrid Social-Ordering Systems," *Columbia Law Review* 119, no. 7 (2019): 1ff.

[121] This issue is debated in Francis Spufford's novel *Red Plenty* e.g. at 83–4 and 291–7 (Spufford imagines efforts to transform the Soviet Union by using computers, rather than markets or politically motivated central planners, to determine efficient prices). Henry Farrell discusses Spufford's insights in Roell, "The Best Books on the Politics of Information recommended by Henry Farrell."

[122] Samuelsson, "What Percentage of Stock Trading is Algorithmic?," The Robust Trader, updated September 7, 2021, https://therobusttrader.com/what-percentage-of-trading-is algorithmic/#:~:text=In%20the%20U.S.%20stock%20market,is%20generated%20through%20algorithmic%20trading.

[123] As, for example, the bureaucratic system of the Church and rudimentary markets coexisted with feudal systems.

[124] James Surowiecki, *The Wisdom of Crowds* (New York, NY: Anchor, 2004). Spufford's fictional account reaches the conclusion that computation cannot replace human crowd sourcing of prices.

[125] J.F. Mergen says "Humans live in many spaces. Machines in just one."

[126] Human, machine, bureaucratic and market intelligences need also to consider problems of execution. Essentially, no one loses at chess or Go because they can't move the pieces. At the other extreme, no one wins at basketball or golf simply because he or she comprehends the game. The different forms of intelligence are differently positioned in meeting these challenges.

[127] See Surowiecki, *The Wisdom of Crowds* and Spufford, *Red Plenty*.

[128] Characteristics locked in at this stage can create enduring difficulties. The recent history of cybersecurity provides an illustration. After computer software flourished in an unregulated market, later efforts at market and bureaucratic controls have not been able to compensate for foundational vulnerabilities in the design of the enormously proliferated technology.

[129] For an illustration of the breakdown of these mechanisms and the resulting consequences, see Jim Frederick, *Black Hearts: One Platoon's Descent into Madness in Iraq's Triangle of Death* (New York, NY: Crown, 2010).

[130] For example, the Food and Drug Administration has been prodded to be faster and more liberal, to make more use of "real world evidence" and to reduce requirements for proof through randomized clinical trials when approving drugs.

Recent legislation including these strictures include the Prescription Drug User Fee Act and the 21st Century Cures Act.

[131] The European Union has taken the lead in this respect as shown in European Union, *White Paper on Artificial Intelligence: A European Approach to Excellence and Trust* (Brussels: February 19, 2020), https://ec.europa.eu/info/files/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en. Recent developments are well described in Mark MacCarthy and Kenneth Propp, "Machines Learn That Brussels Writes the Rules: The EU's New AI Regulation," *Lawfare*, April 28, 2021, https://www.lawfareblog.com/machines-learn-brussels-writes-rules-eus-new-ai-regulation. The authors draw a correct bottom line at the outset: "The European Union's proposed artificial intelligence (AI) regulation, released on April 21, is a direct challenge to Silicon Valley's common view that law should leave emerging technology alone." More generally, see Anu Bradford, *The Brussels Effect: How the European Union Rules the World* (New York, NY: Oxford University Press, 2021).

The increasing engagement of one American regulator can be traced by examining the Federal Trade Commission, *Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues* (Washington, DC: Federal Trade Commission, January 2016), https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf; Andrew Smith (Director, FTC Bureau of Consumer Protection), "Using Artificial Intelligence and Algorithms," Federal Trade Commission, April 8, 2020, https://www.ftc.gov/news-events/blogs/business-blog/2020/04/using-artificial-intelligence-algorithms; and the Senate bill "To direct the Federal Trade Commission to require entities that use, store, or share personal information to conduct automated decision system impact assessments and data protection impact assessments," 116th Cong. (2019), https://www.wyden.senate.gov/imo/media/doc/Algorithmic%20Accountability%20Act%20of%202019%20Bill%20Text.pdf.

[132] A German project points to cutting edge investment in this work. "ExamAI – Testing and Auditing of AI Systems," ExamAI, https://testing-ai.gi.de/english.

[133] Randy Connolly, "Why Computing Belongs Within the Social Sciences," in *Communications of the ACM* 63, no. 8 (August 2020): 54-59, argues along similar lines.