The Great Promise of Digital Transformation

By Rod Collins, Innovation Sherpa at Salt Flats

In his most recent book, *The Seventh Sense*, Joshua Cooper Ramo relates the story of one the most closely guarded secrets during the early years of the Cold War: If the Soviet Union had engaged in a nuclear first strike, it was highly likely the United States would have been unable to respond. That's because the American field officers and their commanders in Washington would have had no way to communicate with each other. Consistent with the technology at the time, the American radio and telephone systems were highly centralized, which made them also highly vulnerable. One of the key structural problems of centralized systems is that each regional center has the potential to become a single point of failure that can disrupt the entire system, as often happens when air traffic across a nation is snarled because of unexpected weather at a major hub. Fortunately, this national security vulnerability was corrected with an innovative solution: the distributed network.

Recognizing the urgency of this challenge, Paul Baran, who at the time was with the joint venture between the U.S. Air Force and the Douglas Aircraft Company known as RAND, devised a way of building messaging systems without any central hubs. Each message would be able to find its own path from point A to point B. Thus, if any part of the system was disrupted, the remaining pathways in the network could resiliently adapt to route all the traffic in the system with minimal disruption. This structural shift from centralized systems to distributed networks, which solved a critical military problem in the 1950's, would turn out to be a harbinger of a dramatic phenomenon that would shape the early twenty-first century: digital transformation.

Digital Transformation Is More Than a Technology Revolution

There is no topic that is both more important and more confusing to business leaders than digital transformation. With the deluge of blogs and books on how the digital revolution is accelerating radical change, you would think that there would be more clarity about this pervasive phenomenon. Instead there is general sense of confusion reminiscent of the Buffalo Springfield lyric: "There's something's happening here; what it is ain't exactly clear."

What we do know is that none of us could live without what are now necessary gadgets that just a mere decade ago were figments of our imaginations. We also know that these marvels are changing our lives more profoundly than any of the progression of the twentieth-century gadgets spotlighted in Disney World's Carousal of Progress. What isn't clear, however, is the extent to which the technology revolution has not only transformed our gadgets but also the fundamental fabric for how the world works. And until business leaders understand the full extent of the profound changes spawned by the technology revolution, digital transformation will remain an elusive enigma.

The first thing we need to understand is that digital transformation is not just a technology revolution; it is far more importantly, the most significant socioeconomic revolution in human history. We are in the middle of an unprecedented inflection point in the development of

civilization—the transition from the first human epoch where centralized hierarchies that leveraged individual intelligence were the basis of social organization to the second human epoch whose social structures will be highly sophisticated distributed networks capable of rapidly leveraging human and artificial collective intelligence. This fundamental architectural shift from hierarchies to networks is the essential evolutionary dynamic of digital transformation and is changing the world rapidly and profoundly, and there is nothing that can stop this change.

In the next decade, we will experience two of the most consequential events in human history: the connection of all humans and things via a common digital network and the proliferation of human collective intelligence via artificial intelligence systems. This transition represents a seismic qualitative shift in the human experience because humanity itself will be transformed. The key building blocks for solidifying this shift are already in place; they just need to be configured.

The next thing we need to understand is how networks work. That's because, as Paul Mason presciently observes in his book *Postcapitalism*, "the 'intelligent machine' was not the computer but the network." In other words, it's not the gadgets that are intelligent, but rather the underlying networks of people and data that connect to the gadgets. What makes the gadgets so powerful is that for the first time we have the wherewithal to rapidly aggregate and leverage the global collective intelligence of human and data networks. And now that we have this capacity, an inevitable evolutionary shift has been set in motion, and when it is complete the fundamental architecture and the basic rules for how all our socioeconomic institutions work will be radically transformed.

When we think of architecture, what comes to mind are beautiful buildings or elaborate edifices. We rarely think of architecture as something that explains how societies or economies work. And yet without social architecture, much of what we experience as everyday life would not be possible.

A fundamental social architecture must answer two questions: 1) How does power work? and 2) How do things get done? In hierarchies, power belongs to those in charge and things get done through the application of centralized control mechanisms. Thus, hierarchical structures leverage the individual intelligence of the elite to organize the work of large numbers of unconnected people. In networks, however, power belongs to the connected and things get done through the application of collective intelligence dynamics that enable the self-organization of work among large numbers of people.

Hierarchies and networks are not equal alternative structures in a hyper-connected world. Networks tend to outperform hierarchies by a wide margin in terms of both intelligence and speed, as we learned in the summer of 2011 when Firas Khatib, a biochemist at the University

of Washington, turned to Foldit to solve a stubborn molecular puzzle that had stumped the world's best scientists for over a decade.

Foldit is a collaborative online video game developed by the University of Washington that enlists players worldwide to solve difficult molecular problems. What's most interesting about Foldit is that many of the more than 250,000 players have little or no background in biochemistry. There are no special requirements for joining the Foldit community—all comers are welcome.

The stubborn puzzle involved figuring out the detailed molecular structure of a protein-cutting enzyme from an AIDS-like virus found in monkeys. Cracking this puzzle could be the breakthrough needed to arrest the medical malady. When Khatib presented the molecular challenge to the Foldit community, what had evaded the world's best individual scientists for ten years was amazingly solved by the collective intelligence of a diverse group of online gamers within only ten days.

Although, it may seem counterintuitive, networks are far more effective and efficient than hierarchies because, by leveraging the distributed intelligence of the many rather than the smarts of the elite few, networks accelerate the path to knowledge. This acceleration is a byproduct a what is known as the network effect, which is achieved when networks reach a sufficient level of critical mass to give rise to the sequential evolution of three laws:

- *The law of connections*: the simple act of connection changes the fundamental dynamics for how power works by shifting the locus of power from elites to peers.
- *The law of self-organization:* When the power shift is complete, peers begin to self-organize their efforts in autonomous and often unexpected ways.
- *The law of collective intelligence:* When a network achieves an effective level of self-organization, it develops the capacity to rapidly aggregate and leverage its collective intelligence, often producing extraordinarily intelligent results at incredibly fast speeds.

As the fundamental architecture for how the world works rapidly shifts from hierarchies to networks, our public and private sector leaders are severely challenged because, as Ramo, notes, "We're at an extremely primitive point in our understanding of networks." Rapidly increasing our understanding of how networks work and how to lead them is the most important leadership challenge of our day. Leaders can no longer afford to build centralized organizations where supervisors with the legitimate authority to kill good ideas or keep bad ideas alive become legions of single points of failure. If leaders want to build resilient organizations that have the wherewithal to rapidly adapt to disruptive change, the first task of digital transformation is to learn how to build and lead highly effective distributed networks.

Collective Intelligence Is the Game Changer

The science fiction writer William Gibson once astutely observed, "The future has already arrived; it's just not evenly distributed." This is especially true today because, despite the increasing evidence of the ascendance of hyper-connectivity, our rapid transformation to a fully networked world remains hidden in plain sight. Although we use our connected iPhones to do Google searches as we step into an Uber car on our way to close a deal we made on eBay, we are in many ways oblivious as to just how radically the world is changing around us. Our tools may be networked and digital, but the world we carry around inside our heads is still very hierarchical and linear. It's an unsustainable situation, and we know how this ends: The future always wins, and it's just a matter of time before the future will become very evenly distributed.

The prime distinction between hierarchies and networks is that hierarchies are designed to leverage the "power of one," while networks naturally enable the "power of many." That is why networks are so much more powerful. Perhaps you may be thinking, if networks are so superior, why is it that hierarchies have shaped our social architecture for all this time? The simple answer is, until the digital technology revolution, we had no way to effectively bring people together into cohesive real-time networks. In the absence of this capability, the best we could do to coordinate the activities of large numbers of people was to build sophisticated hierarchical structures. The fundamental assumption underlying what was once the greatest human organizational innovation is that, by leveraging the individual intelligence of the elite leaders at the top, the whole organization is smarter than it otherwise would be if people were left to their own judgments. And for many centuries, this supposition was true.

However, the current technology revolution has spawned a new and very different innovation in organizational structure that has completely nullified this centuries-old assumption. In the hyper-connected network, the smartest organizations are not the ones with the smartest individuals, but rather those with the capacity to rapidly aggregate and leverage their collective intelligence.

In his seminal book, *The Wisdom of Crowds*, James Surowiecki provides numerous examples of where, under the right conditions, distributed groups are highly intelligent and consistently outperform even the smartest individuals among them. He describes how the sports bookmakers at the Mirage assure the reliable profitability of the betting operations at the Las Vegas hotel by relying on the collective judgments of the gamblers to set the betting lines, how Linus Torvalds defied logic by introducing the phenomenon that has come to be known as crowdsourcing to build the highly successful Linux operating system, how the World Health Organization rapidly deployed a global networked communications structure to rapidly solve the SARS threat before it could spread to pandemic proportions, and how Google, a late entry into a crowded field of upstarts, established quick dominance of the search engine market

when a pair of Stanford graduate students discovered a way to use the collective intelligence of the users to rank the pages.

Despite these compelling examples, tapping into the wisdom of the crowd is more the exception than the rule. Perhaps that's because accessing collective intelligence is not as easy as it may appear. There are many leaders who feel that they are tapping into this resource by gathering different perspectives into a room and managing a spirited discussion among the multiple points of view before making an executive decision. While they may be well-intentioned, this is not how collective intelligence works.

Surowiecki specifies four conditions that are necessary to access the wisdom of the crowd:

- *Diversity of opinion*: Having different perspectives—even eccentric notions—broadens the available information, provides the capacity for evolving ideas, makes it easier for individuals to be candid, and protects against the negative dynamics of shortsighted groupthink.
- *Independent thinking*: Each individual is free to express his or her own opinions without editing and without any pressure to conform to the beliefs of others in the group. Surowiecki makes the point that "paradoxically, the best way for a group to be smart is for each person in it to think and act as independently as possible."
- *Local knowledge*: To truly access collective intelligence, the group must be able to draw upon specialized and localized knowledge because the closer a person is to the problem or the customer, the more likely he or she is to have a meaningful contribution.
- *Aggregation mechanisms*: A distributed system can only produce genuinely intelligent results if there are processes or algorithms for integrating the content of everyone's observations and opinions.

Without all four conditions, accessing collective intelligence is not possible. That is why the leader who gathers different perspectives into a lively discussion is not tapping into the collective wisdom of the group. Although he or she may have access to multiple perspectives and have input from people with extensive local knowledge, chances are organizational politics is interfering with true independent thinking and when the leader is processing the consolidation of the information, there is clearly no aggregation mechanism. Google, on the other hand, has all four attributes. The billions of users assure diversity of opinion as well as sufficient local knowledge, people are free to exercise individual choice of the pages to view, and sophisticated algorithms serve as highly effective aggregation mechanisms.

Of the four conditions, perhaps the most important is the use of aggregation mechanisms. This is why so much of social media is dysfunctional. Popular sites such as Facebook and Twitter have the first three attributes, but are clearly devoid of aggregation mechanisms. And so, while

we have diversity of opinion, independent thinking, and a great deal of local knowledge, without a way to aggregate the different contributions, we have a cacophony of chaos that divides us into myopic tribes and reinforces a highly polarized climate in which compromise and, even more so, consensus becomes impossible. This is the dark side of our hyper-connected world.

The great challenge for social media sites going forward is to become platforms that contribute to the best—and not the worst—that humans have to offer. To do so they need to find ways to develop sophisticated aggregation mechanisms that are capable of accessing and leveraging the collective intelligence of their users to transform tribal positions into innovative solutions that promote the common good. Creating this collective intelligence capability, while admittedly not easy, is the greatest contribution that social media sites could provide humanity for building a better future.

If we are to create this better future, we will need to change the world we carry inside our heads—a world that has been ingrained since the first day we stepped into a schoolroom. There we learned that human intelligence was an attribute of individuals and that knowledge is advanced through a competition of ideas. There is nothing in our educational histories that has prepared us for a world of networked intelligence that has suddenly actualized what used to be a platitudinous sentiment: Nobody is smarter than everybody.

We hold onto old mindsets about human intelligence because, as Surowiecki points out, "One of the striking things about the wisdom of crowds is that even though its effects are all around us, it's easy to miss, and, even when it's seen, it can be hard to accept." Simply put, the phenomenon of collective intelligence has been hidden in plain sight because it defies all our beliefs about how intelligence works. But whether we believe it or not won't matter for much longer because we are on the threshold of one of the most consequential events that will reshape the human experience and accelerate the evolution of both human and artificial collective intelligence: the connection of all humans and things via the Internet of Things (IoT).

The Internet of Things Changes Everything

In 2006, Don Tapscott and Anthony D. Williams described in their book *Wikinomics* how a new phenomenon they called mass collaboration was going to change everything. They recognized that this unprecedented capacity for self-organization would give rise to powerful new models of production based on distributed peer-to-peer networks rather than centralized top-down hierarchies. Tapscott and Williams envisioned a world where this new way of organizing would eventually displace traditional corporate structures as the economy's dominant engine for wealth creation. At the time, many critics dismissed the two authors as being carried away by breathless hype and overstating the impact of the digital revolution. While these critics acknowledged the obvious reality of fast-paced technological innovation, they scoffed at the

notion that new technologies would radically change the fundamental dynamics for how our social structures work.

Given that more than a decade later the top-down hierarchy continues to remain the dominant organizational structure, it might be tempting to conclude the critics are right and that the notion that mass collaboration changes everything is indeed nothing more than hype. However, closure at this point might be premature because there's increasing evidence that we are on the cusp of a new second wave of the digital revolution, which promises to be far more transformative than the already disruptive first wave.

The first wave of the digital revolution emerged with the dawn of the new century when the Internet created new ways to connect people and get things done. One of the interesting developments of this first wave is that it hasn't affected all industries equally. If you are in the media, entertainment, retail, or communications industries, your world has been thoroughly transformed. Stalwart names, such as Border's, Blockbuster, Kodak, Tower Records, and the Encyclopedia Britannica, have been either disrupted or displaced by the upstarts Amazon, Netflix, Apple, Spotify, and Wikipedia. However, if you work in the healthcare, insurance, energy, food processing, or, until recently, the financial services industries, your world has not been heavily impacted by this first wave. For these core economic industries, digital transformation has been essentially limited to digitizing existing product models. Unlike the media and retail industries whose basic business models have been radically disrupted, the longstanding models of the core economic industries have remained essentially the same.

However, this is likely to soon change as we see early signs of the arrival of a new wave, which will leave no industry untouched. This second wave will be spawned by an incredibly powerful force, the Internet of Things, which will not only dramatically accelerate the capacity for mass collaboration, but as Tapscott and Williams foresaw, will indeed change everything.

Jeremy Rifken, in *The Zero Marginal Cost Society*, describes the emerging IoT as the first smart-infrastructure revolution in history because it will allow every human being and every thing to communicate with each other by connecting every machine, business, residence, and vehicle within a single comprehensive operating system. This infrastructure will rapidly take shape over the next decade as the number of sensors grows at exponential speed. In 2007, there were 10 million sensors, by 2013 we achieved 3.5 billion sensors, and by 2030, it is projected that 100 trillion sensors will connect to the IoT.

The IoT has already begun to create a superintelligent network that will allow humans to keep up with the pace of change by making everything more intelligent. According to Steve Case, the founder of AOL, the IoT transforms "the Internet from something we interact with to something that interacts with everything around us."

One of the most important attributes of the IoT will be its ability to recognize weak signals before they become strong ones, allowing humans to recognize patterns before they cause problems. Weak signals are patterns that a human mind doesn't notice because these signals are usually understated when they first emerge, and therefore, are likely to remain hidden in plain sight. Perhaps no industry will be more impacted by this attribute than the arguably biggest laggard in the first two decades of the digital revolution: healthcare.

While most hospitals today are full of smart devices, few of the sensors in these devices communicate with each other. Once these sensors are fully connected via the IoT, the practice of healthcare will be dramatically transformed. For example, the IoT will be able to warn patients—at home via Amazon's Alexa or on a smartphone—of blood clots before impending strokes or heart attacks. Sensors linked to electronic medical records will allow the IoT to quickly diagnose a patient's likely physical state to assist emergency medical personnel and expedite treatment. Skin patches will capture vital data, measuring heart rate, food consumption, and other factors, and will communicate this information to patients and providers through third-party apps. And these apps, using sensor data will become vital conduits of healthcare, providing reliable instant diagnoses.

Rather than waiting for the presentation of symptoms, the IoT will recognize the weak signals of cellular anomalies, notify both patients and doctors through apps, and even set up doctor appointments. As the IoT matures, it will automatically correct the medical malady and notify both patient and doctor of the aborted illness. When the IoT reaches this level of functionality, it will be humankind's first experience of an omniscient system that will know everything about everyone. The system will have "divine-like" qualities because it will be able to benevolently and autonomously intercede to maintain continual health.

Another core industry that is likely to be radically transformed is the transportation industry. As short as five years ago, most of us would have thought the notion of driverless cars was either science fiction, or at a minimum, decades away. Yet, today driverless cars are a reality as Google is currently testing these autonomous vehicles on open roads and traditional automakers such as Mercedes-Benz, Nissan, and Audi are revamping their business strategies around this transportation game-changer. But chances are most of us underestimate just how transformative driverless cars will be. That's because we assume that driverless cars will be individually driven, as autos are today—the only difference being the computer will be at the controls instead of a human operator. Although that may be the initial modus operandi, the ultimate operating system for these autonomous vehicles is likely to be shaped by the IoT.

Because each driverless car is a collection of smart sensor devices that can be interconnected into a holistic network, the IoT will be able to aggregate and leverage the collective intelligence distributed throughout the network to drive all the vehicles concurrently. In other words, no longer will individual drivers, whether human or automated, be making individual decisions by

anticipating what other drivers will do. Rather the IoT system will be making proactive decisions for all the driverless cars with the full knowledge of what each car is doing in real-time. This holistic driving capability has the potential to drastically reduce car accidents. It also means that driverless cars will essentially become a highly sophisticated automated mass collaboration system.

Healthcare and transportation are just two examples of how all industries will be revolutionized by the fundamental architectural shift from bureaucratic hierarchies to digitally transformed networks. Similar metamorphoses are already happening in financial services as blockchain technology eliminates the need for financial intermediaries and will likely happen in energy as individuals bypass corporations to form their own networks for exchanging solar energy.

The Internet of Things will change everything because by connecting every person and every thing into a single global network, the IoT will accelerate the network effect and set in motion the evolutionary sequence of the three laws of networks, described above: the law of connections, which shifts power from elites to peers; the law of self-organization, which creates the conditions for mass collaboration; and the law of collective intelligence, which produces extraordinarily intelligent results at incredibly fast speeds.

However, while the second wave portends much promise for the future capabilities of humanity, these wondrous possibilities are not guaranteed. That's because a byproduct of the acceleration of the network effect is the emergence of two massive jobs to be done, which if not handled properly, could very well wreak havoc on all of us.

The Two Big Jobs To Be Done

One of Peter Drucker's most popular and enduring business quotes is, "The best way to predict the future is to create it." This advice has never been truer than it is today as the technologies of digital transformation are changing all the rules for how the world works. Another famous Drucker quote is, "If you don't understand innovation, you don't understand business." That's because when your job is to create the future—which is the fundamental responsibility of the business leader—you better have a firm grasp of how innovation works.

Perhaps no one has advanced our knowledge on the workings of innovation more than Clayton Christensen, whose many books on the topic have become essential reading for twenty-first century business leaders. In his recent book *Competing Against Luck*, which he co-wrote with Taddy Hall, Karen Dillon, and David Duncan, Christensen emphasizes that the starting point for innovation is often uncovering what he calls "the job to be done." Whereas most attempts at business innovation often start with a product idea, Christensen urges business leaders to step back and take the time to uncover what problem customers are "hiring" their product to solve. If they fully understand that problem and use that knowledge to guide what products to make,

they will not only delight their customers, but they may very well create something that has never existed before, which is the essential task of innovation.

As the structural organization of the world continues its rapid shift from hierarchies to networks, the digital transformation that is driving this transition is not only creating powerful new capabilities, such as the Internet of Things, artificial intelligence, nanotechnology, and robotics, it is also creating unprecedented problems that need to be solved—or as Christensen would describe it, new jobs to be done. In particular, there are two pressing big jobs to be done, which if not solved, could unfortunately mutate digital transformation into digital destruction.

The first and most immediate big job to be done is to transition all IT systems to a new platform that completely extinguishes the "power of one." The urgency of this need was amplified by the recent Equifax security breach, which affected over 143 million people. What's most alarming is that the number of instances of these types of security failures have become so common, we have almost come to accept them as facts of life rather than as totally unacceptable problems to be solved.

The unpleasant reality is that this problem represents a clear and present danger because the power to create large scale catastrophes is reaching the hands of more and more people as connected technologies become more powerful and more ubiquitous. As connectivity expands, data is becoming increasingly more vulnerable and open to exploitation and irreparable loss of control, e.g., ransomware and increasing identity fraud.

The root cause of this problem is that conventional IT systems, which use hierarchical control mechanisms, are no longer "moated" structures once they are interconnected on the Internet. The combination of hierarchically structured controls in an increasingly networked world is toxic because any smart hacker can use the network to break into most, if not all, conventional IT systems, and once in the system, can easily figure out how to manipulate its single point of control mechanisms, such as passwords. In a hyper-connected world, traditional control mechanisms that act like singular keys to allow individuals to access data are rapidly becoming perilous liabilities.

The continued maintenance of centralized computing architectures in a hyper-connected world is unsustainable because with single points of control, companies are increasingly vulnerable to crashes, fraud, and security breaches. More ominously, the possibility of surreptitious and total surveillance by both government and nongovernment actors is now a reality. These circumstances will get worse, not better, unless we correct for the toxic mix between networked technologies and hierarchical IT structures.

The solution is to replace singular control mechanisms with collective intelligence dynamics and replace the "power of one" with the "power of many." In other words, we need a networked security architecture for a networked world that requires a consensus among a set of multiple actors to validate an action. This type of architecture would have likely prevented the Equifax breach.

The second big job to be done is to create a new economic engine to preserve the middle class as the number of jobs are dramatically reduced. Society is entering a new phase in which fewer and fewer workers are needed to produce and distribute all the goods and services consumed.

Jaron Lanier, the author of *Who Owns the Future?*, cites the examples of Kodak and Instagram to demonstrate this phenomenon. Kodak, at its peak, employed more than 140,000 people and was worth $28 billion. When Instagram was sold to Facebook for a billion dollars in 2012, it employed only 13 people. Instagram's value doesn't primarily come from 13 people, but rather from the millions of users who contribute to the network without being paid for it. Today's networked economy is a form of a feudal system, where many people contribute work but only a small number get paid. This has the net effect of centralizing wealth and limiting overall growth. An economy that fails to compensate the prime contributors of value is not a sustainable system.

Another troubling development is that, for the first time in human history, technology and automation are no longer replacing displaced jobs with better and higher paying jobs. As a consequence, the longstanding Industrial Age concept of a job may become obsolete.
While most of us understand that the technology revolution will rapidly displace all manual jobs with robots and 3D printers, few of us have recognized that automation has begun to displace knowledge and service jobs. If enough jobs are eliminated, the customer base will dwindle rapidly. And if there are no customers, the economy as we have known it may collapse.
Perhaps, this explains why Elon Musk and Mark Zuckerberg have both become proponents for the notion of universal basic income. However, such a notion is a short-sighted solution and an attempt to ignore Albert Einstein's sage insight that, "We can't solve problems using the same kind of thinking we used when we created them." We are surely capable of a more creative and a more reliable solution.

The real problem is not the disappearance of jobs, but rather the disappearance of the primary mechanism for compensating people for the value they contribute to the overall economy. The notion of a universal basic income is short-sighted because it separates compensation from value contribution. The real issue is that, in these early years of a digitally transformed world, nobody knows how to value contributions in a networked information economy. The traditional concepts of economic value are rapidly becoming obsolete in a hyper-connected world. Plant, property, equipment and labor are no longer the meaningful categories of value. They are being replaced by data, information, intelligence, and contribution.

If we want to preserve the middle class, we need to change how the economy works. We need to create new economic mechanisms that can accurately, fairly, and sufficiently measure and distribute economic value across the broad population of all network contributors. In the new economy, most traditional labor will indeed be accomplished by robots, 3D printers, and artificial intelligence, and the economic value of individuals will no longer happen in the context of a job, but rather in the context of their contributions to networked activity. Because traditional accounting models were not designed to measure the value of data, information, intelligence, and contribution, we will also need a new accounting paradigm designed for a post-digital economy.

Developing a new practical paradigm of economic value as well as the means to equitably calculate the contributions of individuals who participate in value-generating networks is the biggest economic challenge of our times and the necessary solution to preserving the middle class. Fortunately, while digital transformation has created these perilous circumstances, it has also enabled the creation of a new relatively unknown technology that may solve both of these two big jobs and whose underlying dynamics are strikingly similar to the organizing principles of a highly innovative company.

Blockchain May Be the Future of IT

Morning Star is not your usual company. That's because the 400-person California-based agribusiness has no supervisors. Rather than relying on the intelligence of an elite few, Morning Star is a highly successful self-managing peer-to-peer network that has skillfully leveraged the "power of many" to sustain its position as the world's largest tomato processor.

From the beginning, the company's founder, Chris Rufer, built his innovative enterprise on two core principles. First, individuals should keep their commitments to others. People at Morning Star are not handed assignments. Instead they negotiate Colleague Letters of Understanding with their co-workers, and, to assure everyone is honoring their commitments, the metrics associated with these agreements are published bi-weekly.

The second principle is that no individual should use force against others or their property. This means that no single person has the authority to issue an order or the ability to unilaterally fire another person. For the people at Morning Star, what they do and with whom they work are always collective decisions.

For over three decades, these two principles have served as a solid foundation to support an innovative management technology that is disrupting the way we build and lead human organizations.

Recently, there has been increasing buzz about another innovative technology—this time in the financial services world—that many believe has the potential to disrupt the way we build and organize IT systems. This technology is blockchain. Like Morning Star, blockchain is designed as a peer-to-peer network which, upon close examination, also follows Rufer's two core principles.

Blockchain is the creation of an anonymous individual or group of individuals who, using the pseudonym Satoshi Nakamoto, published a short paper in 2009 that outlined an unconventional peer-to-peer system that allows users to directly transact business without the need for any intermediaries.

Blockchain is a distributed ledger system that uses a network consensus to record and execute transactions. It's best known as the platform for the Web currency Bitcoin. Blockchain's most distinguishing characteristic is that no single agent has the ability to execute control over system activity. Or to use Rufer's words, no individual can engage in coercive activity against another person or their property.

To understand how blockchain works, consider this analogy—which while admittedly simplified—conveys the basic sense of this paradigm shift in systems architecture. Imagine you are attending an auction, along with 300 other people, to bid on the numerous treasures and heirlooms from the estate of a recently departed collector. Let's also imagine that the 300 auction participants are a blockchain community. As the auction proceeds, there is a particular painting that you would like to bid on, but you can't afford to spend more than $2,000. As the bidding proceeds, you find yourself in a competition with another participant who bids the painting up to $2,500 and wins the bid.

Because the auction is using blockchain, recording this transaction requires the majority of the 300 people in the room to agree that the particular painting was sold for $2,500 to the competing bidder and to affirm he has the cash to pay for the painting. Once consensus is reached, the transaction is grouped and recorded with other bidding transactions into a block, which is permanently timestamped and connected into a chain with other blocks of transactions from the auction, hence the name blockchain.

When a block is connected to the chain, it is immutable and can never be altered. In addition, these blocks are not recorded in a single central ledger, but rather into a distributed ledger, which means that all of the participants have their own individual copies of the ledger. This makes it difficult for a single individual to commit fraud because all copies would need to be changed to pull off the counterfeit transaction. As Lanier succinctly puts it, "You can fake an ID, but you can't fake a thousand concurrent views of the person you are falsely pretending to be."

Any modification to a transaction has to be recorded as a separate immutable entry in a new block that references the timestamped original transaction. This means that, if after the event,

the winning bidder, who also happens to be the brother-in-law of the auctioneer, tries to persuade the auctioneer to accept a lessor amount for the painting and change the recorded transaction because "we're family," the auctioneer would be unable to do so because adjusting the record would require a completely new transaction that would need to be agreed upon by the blockchain community. Obviously, the majority of the auction participants are not going to affirm a false entry.

This game-changing systems architecture is likely to revolutionize the way we build IT systems because it has the potential to eradicate most hacking and fraud activity and would provide a solid solution for the first big job to be done discussed above: to transition all IT systems to a new platform that completely extinguishes the "power of one."

Another innovative characteristic of blockchain is that it has the capacity to create what are known as "smart contracts." In typical business arrangements, expectations and agreements between parties are usually memorialized in written contracts, with the understanding that each of the parties can trust the others to act according to the terms of the contract. If that trust is breached, the parties have the option to sue each other in legal proceedings, which are often prolonged and costly. Smart contracts eliminate the possibility of breaches by transforming the foundation of trust from reliance upon the good intentions of others to credence on a basic attribute of the blockchain system

With smart contracts, agreements and expectations are built into the system, making it difficult, if not impossible, for one party to violate the terms of a contract. Thus, any time there is a transaction, recording it on blockchain requires an affirmation that the action is consistent with the terms of the smart contract. In other words, Rufer's principle that individuals should keep their agreements with others is a core pillar of the basic architecture of blockchain systems.

Blockchain also provides a foundation for solving the second big job to be done of creating a new economic engine to preserve the middle class as the number of jobs are dramatically reduced. As mentioned above, blockchain is the architecture used for the cryptocurrency in Bitcoin. This means that blockchain has the capacity to calculate value. At this time, the mathematical dynamics for calculating cryptocurrency are relatively basic. However, as the evolution of the Internet of Things, the emergence of artificial intelligence, and the phenomenon known as Moore's Law continue to accelerate the technology revolution, there is a strong possibility that blockchain systems architecture will develop the capability to create sophisticated algorithms that can accurately calculate the economic value that people contribute by their participation in network activity. This value is likely to take the form of a cryptocurrency that will be valid tender in specific economic markets.

If you are wondering how this form of currency might work, let's look at an early precursor of this type of value creation with which we are all familiar: frequent traveler points. Frequent traveler points are a form of currency that have real economic value in specific markets. In addition to booking flights and hotel rooms, these points can be used for a wide range of approved retail purchases from dedicated catalogues. Similarly, once algorithms have the wherewithal to calculate the relative value of our participation on social media sites or our contributions to the quality of Wikipedia articles and Google searches, we can be fairly compensated in cryptocurrency for our activity in these new economic platforms. When we have this capability, the pervasive problem of the inherent wealth inequality in the current structure of the Internet will begin to resolve itself as more people have a vehicle for wealth creation in a digitally transformed economy.

While blockchain is clearly in its infancy and has many bugs to be worked out and applications to be developed, the network-based architecture of this innovative technology promises to go a long way in making sure the benefits of digital transformation significantly outweigh the dangers and that these benefits are shared by all as we continue to discover new ways to create wealth.

An Evolutionary Leap in Human Intelligence

"May you live in interesting times," is an old English expression whose enigmatic meaning can make people wonder whether they've been offered a blessing or a curse. Regardless of the apparent well-wisher's intentions, interesting times are oftentimes a blessing for some and a curse for others. The difference depends upon how capable and how fast people are in recognizing and embracing new ideas and new opportunities.

We, today, are clearly living in interesting times, perhaps the most interesting times in the history of human civilization. If our interesting times are to become a blessing, we will need to dramatically increase our understanding of networks and how they work because digital transformation is about to go into overdrive with the coming emergence of the Internet of Things. Once the IoT becomes the fundamental fabric of our everyday lives, what has been hidden in plain sight will become painfully obvious to those who fail to expand their knowledge of networks.

We are on the threshold of an evolutionary leap in the human species because the IOT will create new dimensions of human intelligence that have never existed before. Thanks to Moore's law, sensors will become ubiquitous. They are already in our cars, home appliances, street lamps, medical devices, and smart phones. They will soon be in our clothes, the walls in our homes, the buildings where we work, our food, our pets, and even our bodies. As these sensors proliferate, they will become inevitably interconnected by the IoT and will propagate a

plethora of data at rapid speeds in permutations and combinations that have never been available to us before.

When this happens, the IoT will become a platform for highly sophisticated artificial intelligence (AI) systems that will substantially exceed the capabilities of any single individual. These AI systems will be able to process the universe of sensor information at the speed of Google searches. They will be able to recognize weak signals and identify patterns across the data that would normally escape even the brightest experts among us. And using the technologies of machine learning, these AI systems will interact with humanity to advance our understanding of the world around us and to enhance the quality of our lives, especially as AI systems become the foundation for a new era in medicine fundamentally oriented around sustaining health rather than treating illnesses.

Currently, there's much ambivalence about the future of AI and whether it's a blessing or a curse, especially among those who are concerned that advanced AI could become hostile and upend humankind as the dominant species. However, if we look at AI from the perspective of a network mindset rather than the more familiar hierarchical worldview, we will begin to understand why AI systems are far more likely to be extensions of humanity than they are to be separate and competitive entities.

Hierarchies are designed to leverage individual intelligence, and thus promote the notion that the most intelligent should naturally rise to the dominant position at top of the pyramid. This explains why those with a limited understanding of networks are fearful of AI. They dread that AI will become the smartest entities on the planet and will use this intelligence to exercise coercive control over humankind. However, the phenomenon of digital transformation, which is providing the technology that makes AI possible, is also accelerating the practicality and the proliferation of the peer-to-peer network. Networks, in contrast to hierarchies, naturally leverage collective intelligence and are more inclined to foster mass collaboration rather than fractious competition.

The unprecedented emergence of technologies and social structures that have the capacity to rapidly aggregate and leverage collective intelligence is a Cambrianesque leap in the evolution of human intelligence. Rather than being separate entities, AI systems following the evolutionary laws of networks—the law of connections, the law of self-organization, and the law of collective intelligence—have the potential to be the ultimate manifestation of human intelligence.

I first learned of the extraordinary power of collective intelligence about twenty years ago when I was tasked with improving the performance of a complex business alliance of thirty-nine companies that had sustained two decades of anemic growth. As our team considered various strategic options, it became clear to us that the root of our growth stagnation was our own

management approach. We recognized that we led this alliance as if it were a typical hierarchical firm, and as a result we spent a substantial amount of time arguing with the various organizations about who was in charge. Maybe, we began to surmise, we need to learn how to lead a network. This would mean developing a very different set of skills based around building consensus rather than giving directives.

An important part of our new approach was to design an innovative meeting format (which today is known as the Collective Intelligence Workshop) that would minimize the usual debates that characterize the typical business meeting and to employ sophisticated facilitated meeting techniques to achieve rapid consensus among the disparate partners in our business alliance. These meetings were highly successful, and it wasn't long before we realized that in designing our "no debate" meetings, we had fortuitously and unexpectedly stumbled into a process that effectively and rapidly aggregated the collective intelligence within the room. And once we discovered this powerful phenomenon, we were able to use it to solve previously unsolvable problems.

What impressed us immediately about collective intelligence is that it is extraordinarily intelligent and incredibly fast. We recognized that having a process to converge the best thinking from divergent perspectives gave us the capacity to create breakthroughs that no one of us could do alone. And given our track record of endless parochial debates, we were stunned that issues that would normally fester for months, or sometimes years, could be resolved in the span of a few hours. But perhaps the most distinctive attribute of this newfound asset is that it greatly diminishes, if not eliminates, the natural cognitive biases that the psychologists Daniel Kahneman and Amos Tversky convincingly proved plague individual judgment and decision making.

While the history of humanity has been a progressive evolution of more advanced technologies and civilizations, this human development has been marred at times by inhuman atrocities. Most of these atrocities can be traced to the manifestation of some form of human bias acted out in judgments or decisions made possible by the oppressive exercise of control that is an inherent hazard in hierarchical forms of social organization. One of the social benefits of collective intelligence structures is that, by effectively integrating all the various diverse perspectives, they essentially eliminate the wherewithal for single individuals to wrestle enough control to impose their individual biases on captive audiences.

Over the next decade, tandem developments in the Internet of Things and blockchain technology will come together to enable a human-machine symbiotic network that will morph into a superintelligent artificial intelligence system exponentially smarter and faster than any of us individually or all of us combined. The continued rapid emergence of the peer-to-peer network as the dominant form of social organization, together with the use of collective intelligence dynamics to shape the development of artificial intelligence systems will ignite an

extraordinary leap in the social exercise of human judgment and human decision-making. If this happens, our interesting times are highly likely to become a profound blessing.

The Human-Machine Symbiosis

Over three days from February 14 – 16, 2011, the American public was introduced to the amazing capabilities of artificial intelligence during the telecast of the *Jeopardy!* IBM Challenge, when IBM's Watson faced off against Ken Jennings, the holder of the record for the game show's longest winning streak, and Brad Rutter, *Jeopardy!'s* all-time money winner. This much-publicized human-versus-machine competition turned out to be no contest at all, as Watson easily beat its human challengers with a score of $77, 147 to Jennings's $24,000 and Rutter's $21,600.

With this remarkable victory, the public learned what was once considered science fiction was suddenly scientific reality: computers could interact with humans in intelligent conversation. Watson wasn't a linear program following a series of pre-planned fixed steps, like a dishwasher. It was an intelligent machine that could absorb spontaneous human communication and could respond sensibly. And what was most impressive is that it could do so far smarter and faster than the long-running game show's two greatest champions.

This new reality, while incredibly awe-inspiring, was also unsettling for many. Now that we had first-hand experience with the quality and speed of AI, we could understand why some were concerned that, sometime in the near future, superintelligent robots might develop to the point that they could overtake or even subjugate humans. However, as noted above, this thinking may be shortsighted because it reflects the prevailing hierarchical mindset, which assumes power is a function of being in charge and, therefore expects humans and machines to behave as separate and competing entities embroiled in a battle to see who comes out on top. If this hierarchical mindset continues to shape how we think and act, this concern could eventually morph into what many would consider a clear and present danger.

If we are too slow in transitioning IT systems control structures from linear hierarchical architecture to a more robust network architecture, the day could likely come when one madman would be able to shut down an electrical grid, use IoT to spread a deadly virus, or possibly unleash a weapon of mass destruction. And if this were to happen, perhaps the concerns about an independent-minded AI might materialize should it harness its formidable intelligence and exercise its power of judgment to stop the madness or even eliminate the threat. Just as Jennings and Rutter were easily defeated by IBM's Watson, if human intelligence and machine intelligence remain separate entities, then these fears may prove true. However, it doesn't have to be that way as another highly publicized contestant who lost to AI discovered.

Garry Kasparov is considered by many to be the greatest chess player of all time. During his active career between 1986 and 2005, he was the world's #1 ranked chess player for 225 out of 228 months. Like Jennings and Rutter, he knows how it feels to lose to a machine because Kasparov has the distinction of being the first world champion chess player to lose a match to a computer, when he was defeated by IBM's Deep Blue in 1997. Kasparov, however, had a somewhat counter-intuitive and creative reaction to his drubbing. He decided to apply an age-old adage, "If you can't beat them, join them," and engaged in an interesting experiment where he paired human chess masters with machines to compete against other machines in a series of matches. In every instance, the human-machine combination defeated the solitary machine. What we learn from Kasparov that we didn't see on *Jeopardy!* is that humans can become more powerful—and perhaps even more human—when they collaborate rather than compete with machines.

In his book, *The Innovators*, Walter Isaacson asserts that the most important development of Digital Age innovation is the emergence of a new form of human-machine symbiosis that is dramatically transforming the essential orientation of all systems from programming to learning. This insight is significant because it reinforces the pressing need for IT systems builders to shift their control architecture from hierarchically programmed structures to networked learning structures, especially if we are serious about curbing all the data breaches.

The notion of a human-machine symbiosis is not a new creation of the Digital Age. This phenomenon traces its roots as far back as the Hunter Gatherer Age when humans first built tools to ease the burden of physical work. This symbiosis, which incrementally evolved through the Agrarian Age over several thousand years, catapulted in both form and scale with the emergence of the Industrial Revolution. The proliferation of mechanical inventions, the advent of mass production, and the rise of bureaucracies and corporations reformulated the fundamental dynamics of the human-machine symbiosis. Rather than machines being merely tools, as they were throughout the Agrarian Age, the machine became the dominant metaphor for the worldview that defined the context of everyday social and economic life in industrialized societies.

This mechanistic worldview was reflected in the fundamental organizational design principles of Frederic Taylor, whose Scientific Management model became the template for the command-and-control structures that have defined the practice of management for well over a century. Accordingly, the basic orientation of this management approach is prescribed programming where workers are expected to carry out fixed plans, and where controls and incentives are put in place to make sure employees don't deviate from the program. This ultimate form of the top-down hierarchical architecture often resulted in Borg-like entities where large numbers of people interacted with each other in rigidly prescribed ways. Most of us have been so socialized into this mechanistic worldview that we fail to recognize that the human-machine symbiosis of the last two centuries has favored the machine over the human. That's why the fundamental

dynamics of all systems are grounded in programming. In many ways, at least in our social architecture, this pervasive orientation toward programming resulted in humans unwittingly became more machine-like.

However, with the recent emergence of the phenomenon of digital transformation, we are witnessing a radical reformulation of the human-machine symbiosis as machines are becoming more humanlike. Rather than being a threat to humanity, this reformulation of the human-machine symbiosis could very well be the renaissance of humanity.

Isaacson observes that today's computer technology "augments human intelligence by being tools both for personal creativity and for collaborating." This means that as machines become more humanlike, we have the opportunity to partner with machines in ways that will greatly accelerate our capacity to learn. Isaacson notes, "that no matter how fast computers progress, artificial intelligence may never outstrip the intelligence of the human-machine partnership." Consequently, the symbiotic relationship that combines the strengths of both humans and machines could usher in a new era of enhanced human learning and intelligence. Isaacson points to the example of the Google search engine, which rapidly collates the individual judgments of billions of people to provide sensible search results.

It isn't the intelligent malevolent machine that presents the greatest danger to human civilization; it is the singular malevolent individual who misuses the power of singular control to wreak havoc in our hyper-connected world. A human-machine symbiosis that's built on a platform of networked collective intelligence could enhance the human experience far beyond our wildest expectations by eliminating the capability for single individuals to engage in large scale coercive actions and mitigating any concerns about AI overtaking humanity. But accomplishing the first big job to be done will only be possible if we complete the tasks of digital transformation and invent the new tools needed to embrace a new mindset, create a new economy, and build a new world that leverages the "power of many" and eliminates the "power of one."

The Shift from Workers to Gamers

As the powerful forces of digital disruption continue to accelerate, we are likely to see the emergence of a pressing world-wide social problem: a dramatic and permanent reduction in the number of jobs across all industries and institutions. Within the next two decades, it's highly likely that advances in robotics and the proliferation of home 3D printers will eliminate most manufacturing jobs. Amazon is currently experimenting with cashierless stores, a harbinger of the coming massive reduction in service jobs. And, perhaps most surprising and unexpected, will be the extensive eradication of professional jobs as knowledge workers are displaced by the smarter and faster capabilities of artificial intelligence. Needless to say, there will be a massive sense of urgency around the second big job to be done.

It shouldn't surprise us that, as we enter a new economic age, jobs may become obsolete because the notion that people receive economic value for providing time in mass production enterprises is an invention of the Industrial Revolution. For the most part, there were no hourly wages or annual salaries in the Agrarian Age.

Technology revolutions breed economic revolutions, which often spawn new ways to create economic value and new vehicles to distribute the new value. The invention of the job was the new vehicle for the Industrial Age. While we have seen how the Digital Age is creating new and massive opportunities for wealth, what remains to be seen is how we can better distribute that wealth.

The expected dramatic reduction in traditional employment doesn't necessarily mean that work disappears. Work will continue to be an important dimension of day-to-day life, but its context is likely to be very different in a digitally transformed world. Like the Foldit example referenced earlier, our economic contributions are likely to happen within the context of a radically different economic role: the gamer.

While the role of the gamer may seem foreign—or even far-fetched—being a contributor is something to which we can all relate because so many of us are already contributors to various Internet venues, such as Google, Facebook, and Instagram. As we shift economic paradigms, what needs to change is the opportunity to be paid for these contributions. The role of the gamer may be a tangible way to effect that change and may become a very lucrative pathway to middle class wealth in the very near future.

For those born before the 1990's, the idea that participation in games could become the conduit for middle-class wealth might seem bizarre. That's because most of us view games as recreational pastimes that serve as diversions from real-world work. Games are what we do when we want to entertain ourselves or we need to keep our children occupied. In the past, games were about sitting around a table playing cards or board games. Games also included playing on sports teams in basketball or softball leagues or participating in online Fantasy Football leagues. Unless you are a professional athlete or a professional gambler, for most adults, games have not been considered serious activities.

However, those born after 1990 and who have grown up with the Internet have a very different perspective. They understand that games can be powerful and highly productive platforms for mass collaboration. Most of us would be surprised to learn that, according to Jane McGonigal, the author of *Reality Is Broken*, gamers spend more time today compiling collective intelligence than anyone else. Because one of the distinguishing attributes of networks is their capacity to leverage the extraordinary speed and intelligence of the crowd, then understanding the

thinking and the activities of those who are most involved in compiling collective intelligence is a good way to begin to acclimate to the new realities of the coming economy.

McGonigal suggests new developments in online video games can make significant contributions to boosting the productivity of many business organizations. That's because today's online games are increasingly collaborative efforts and have two things in common with successful businesses: "a clear goal and actionable steps toward achieving that goal." McGonigal cites the real-world example of how one business used gaming technology to quickly solve what appeared to be an insurmountable problem.

In the summer of 2009, the *Guardian* newspaper had obtained leaked documents that pointed to widespread expense fraud among the members of the British Parliament (MPs). To appease public outrage, the MPs engaged in a bit of gamesmanship by providing the newspaper with an unsorted electronic data dump of more than a million expense forms for claims from the previous four years. Knowing it was impossible for their reporters to make sense of the data dump, the *Guardian* decided it would crowdsource the effort by developing a game called "Investigate Your MP's Expenses," and invite all British citizens to participate.

The response was astounding. Within a few days, more than 20,000 participants were able to quickly sort through the records and identify several probable perpetrators among the MPs. The collaborative work of the gamers led to the resignation of at least twenty-eight MPs, criminal proceedings against four MPs and the repayment of 1.12 million British pounds by hundreds of other MPs.

While the *Guardian* gamers volunteered their contributions, this example shows how games can be a highly efficient way of working. Once business leaders grasp the power of this new model, they may use similar gaming applications to solve what have been intractable problems and be willing to pay gamers for their contributions. According to McGonigal, very big games represent the future of collaboration and may be the best hope for solving the most complex problems of our time.

Another likely candidate for gaming applications is blockchain technology. As blockchain evolves and becomes more ubiquitous as the fundamental platform for IT, we will need to build a robust Proof of Truth mechanism that better validates the content of transactions to replace the current cumbersome and energy-intensive Proof of Work mechanisms. Transforming blockchain miners into gamers may be the practical vehicle for incentivizing the effort necessary to make a Proof of Truth mechanism work on a very large scale.

Don and Alex Tapscott, the authors of *Blockchain Revolution*, describe other ways that games could be used to create economic value. For example, Nike could use blockchain and gamer technology to generate and store data on a distributed ledger that both the company and its

customers could monetize as agreed in smart contracts. Nike could offer a tiny piece of its shares with every pair of its shoes, if the consumer would agree to activate the smarts in the shoes, or even sync the shoes to other wearables such as a heart monitor.

These are just some of the innovative ways that all of us—rather than just the founders and the investors—could share in the new wealth created by digital transformation. And while being a gamer may seem like a strange way to make a living, if both paid the same, which would you rather do: follow somebody else's orders as a compliant worker or have fun playing games?

Moving Beyond the Hazards of Human Bias

In 2002, Daniel Kahneman was awarded the prestigious Nobel prize in economics. This extraordinary feat was even more impressive because Kahneman is a psychologist. Although he was recognized by the *Economist* in 2015 as the seventh most influential economist in the world, Kahneman holds no formal academic credentials in the field of economics. How does a psychologist so earn the respect of colleagues in a completely different field that, without any formal training in their field, he is embraced not just as one of their own but is considered one of their stars? The answer is rather simple: convincingly demonstrate that one of the prime assumptions of a discipline is completely wrong.

In his recent book, *The Undoing Project*, Michael Lewis popularized the story of how Kahneman and his close friend and colleague, Amos Tversky, collaborated over the course of several decades to understand the working dynamics of human judgment and decision-making. They closely examined the ways in which people make decisions under conditions of uncertainty and found that their judgments and behavior were consistently at odds with the longstanding assumption of economic theory that takes for granted that people act rationally by seeking to maximize their gains.

The extensive research of the two psychologists revealed that, when making judgments, people were strongly influenced by a dynamic they called "framing." For example, if the subjects were presented with the sequence "A |3 C," people would identify the middle figure as the letter "B." However, if presented with the sequence "12 |3 14," people would identify the middle figure as the number "13." Similarly, Kahneman and Tversky discovered that, contrary to the prevailing utility theory, people could shift from risk avoiding to risk seeking behavior simply by presenting them with a different description of the same reality. This means that framing can lead to errors in decision-making and cause people to behave in ways that clearly do not maximize gains, and in some instances, actually result in unwitting losses. How does this happen?

In his best-selling book, *Thinking, Fast and Slow*, Kahneman asserts that humans engage in two different thinking modes in their day-to-day lives. He refers to these ways of thinking by the nondescript names System 1 and System 2. System 1 is fast thinking, which operates automatically with little or no effort. It is highly proficient at identifying causal connections between events, sometimes even when there is no empirical basis for the connection. System 2, on the other hand, is slow thinking and involves deliberate attention to understanding details and the complex web of relationships among various components. Whereas System 1 is inherently deterministic and undoubting, System 2 is probabilistic and highly aware of uncertainty and doubt. Needless to say, these two ways of thinking are contextually very different.

Tversky summed up the human thinking paradox well when he concluded that "Man is a deterministic device thrown into a probabilistic universe." Given the limited capacity of the human brain, the time it takes to do System 2 thinking is misaligned with the speed needed to make practical decisions. For example, if you suddenly find yourself in an unfamiliar place late at night in the presence of a complete stranger with no one else is in sight, you will need to make some practical decisions rather quickly. Doing a detailed background check on the stranger is not possible, so you will have to make a fast judgment about whether this unknown person is likely friendly, hostile, or indifferent. You will rely upon your experience and intuition to quickly examine the clues in front of you to make a decision whether to ignore the person, engage in a conversation, or flee as fast as you can.

This type of fast thinking is far more common than slow thinking because the high degree of ambiguity and the rapid pace of events that is the context of our day-to-day lives makes System 2 thinking highly impractical, which is why Kahneman contends, "the intuitive System 1 is more influential than your experience tells you, and it is the secret author of many of the choices and judgments you make."

Utility theory assumes that people are rational agents because, for the most part, we see ourselves as levelheaded people who make informed and deliberate choices. Unfortunately, this perception is more illusion than fact. Kahneman and Tversky discovered two important tendencies about human thinking: "We can be blind to the obvious, and we are also blind to our blindness." In other words, despite the confidence we feel from our System 1 thinking, there's a great deal that we don't know, but more importantly, there's a great deal that we don't know that we don't know. And when that happens, humans are naturally prone to hold unconscious biases.

A common example is confirmation bias, which is the tendency to immediately interpret new information in a way that confirms an individual's preexisting beliefs or opinions. As a byproduct of System 1 thinking, we build mental narratives to make sense of the continuous flow of information and events that we need to rapidly process in our daily lives. These mental

narratives are the soil of confirmation bias and can often cause us to make confident decisions that are completely wrong. Even the most intelligent among us are not immune from this tendency. Medical doctors, like all other professional experts, are prone to construct mental narratives based upon their individual and shared experiences. However, these narratives, while useful most of the time, can often get in the way of accurate diagnoses. Lewis cites research by the Oregon Research Institute that found that an algorithm was more effective at diagnosing cancer than a group of doctors and outperformed even the single best doctor.

One of the reasons that algorithms outperform experts is that algorithms don't have the limitations of the human brain. Thanks to Moore's Law, computers have a seemingly infinite capacity to store and process information, which means that one of the truly unprecedented benefits of burgeoning AI systems is that we may soon have the ability to process System 2 thinking at System 1 speeds.

In addition to speed, another factor that distinguishes AI is that its underlying structure is networked. Because the aggregation mechanisms in the algorithms of AI networks will have instant access to the full diversity of data that represent the independent thinking and the local knowledge of different perspectives, these machine learning systems will meet the criteria identified by Surowiecki for harvesting collective intelligence. This is important because one of the attributes of collective intelligence is its ability to integrate diverse and even opposing perspectives into workable holistic solutions that move beyond the limits of human biases.

Amir Husain, the author of *The Salient Machine*, makes the point that AI is not just another technology, it is a new form of intelligence. With its vast capacity to store information and its ability to rapidly process and retrieve information at the speed of Google searches, AI doesn't need to engage in the heuristic short-cuts that are prevalent in Systems 1 thinking. Without the limitations of the human brain, AI is far more capable of thinking probabilistically and holistically, and is, thus, capable of weighing the relative content of multiple perspectives in a matter of seconds. In other words, the great potential of artificial intelligence is that it could effectively put an end to the flawed human biases that often plague decision-making in our social organizations.

For centuries, hierarchical structures have been the near universal template for how we have designed our social organizations. Because top-down hierarchies naturally leverage the individual intelligence of the elite few, they are susceptible to—and even amplify—human bias. The foundational theory of hierarchical structures is that, by giving the supposedly smartest people who rise to the top of these organizations the authority to command and control the work of others, then the organizations will be smarter than they otherwise would be. Unfortunately, smart people are not immune from human biases, and when they make decisions based on narratives that out of touch with what's actually happening, the consequences can be drastic. The blindness of the financial experts who gave us the Great

Recession and the failure of traditional media companies to grasp the significance of the digital revolution are two examples of how narratives that were once useful guides can suddenly become pathways to disaster.

The evolutionary leap in the fundamental fabric of human social organization is the essential element of digital transformation. As the context for how humans think and act together rapidly shifts from top-down hierarchies to peer-to-peer networks, and as our social leaders become more proficient in their understanding of how to build and lead networks, we will no longer be blind to the obvious, and more importantly, we will no longer be blind to our blindness because the emerging human-machine symbiosis made possible by our newfound capacity to leverage the extraordinary amplitude and the incredible speed of collective intelligence will go a long way toward moving us beyond the hazards of human bias.

The Dawn of the Second Human Epoch

The evolutionary leap made possible by digital transformation is arguably the most consequential socioeconomic event in human history. We are at the beginning of an unprecedented inflection point in the development of civilization—the transition from the first human epoch where centralized hierarchies that leveraged individual intelligence were the basis of social organization to the second human epoch whose social structures will be highly sophisticated distributed networks capable of rapidly leveraging human and artificial collective intelligence.

This fundamental architectural shift is far bigger than most of us are prepared for because the vast majority of us are lacking in an understanding of the dynamics of networks. We are much more conversant in the mechanics of the hierarchical structures that have shaped the evolution of the human experience since the dawn of civilization. Unfortunately, our knowledge deficiencies in the ways of networks cause us to significantly underestimate both the magnitude of the inflection point and its related exponential rate of change.

This growing knowledge gap is problematic because hierarchies and networks are neither equal nor interchangeable structures. As demonstrated throughout this article, networks tend to outperform hierarchies by a wide margin in terms of both intelligence and speed because, by leveraging the collective intelligence of the many rather than the individual smarts of the elite few, networks dramatically accelerate the path to knowledge. The natural proclivity of networks to leverage collective intelligence is the great game changer and, arguably, the single most important consequence of digital transformation because it is a new and far more powerful form of intelligence.

The Great Promise of Digital Transformation

By Rod Collins, Innovation Sherpa at Salt Flats

The connection of all humans and things via a common digital network and the proliferation of human collective intelligence via sophisticated artificial intelligence systems will result in a Cambrianesque leap in the human species because humanity itself will be transformed as we gain access to the extraordinary speed and capacity of collective intelligence through a metamorphic human-machine symbiosis. It is this metamorphosis that will radically change all the rules for how the world works and usher in the dawn of the second human epoch.

An epoch is an extended period of time marked by distinctive features or events and is usually comprised of several successive ages. The first human epoch is known as the Holocene Epoch. It began about 12,000 years ago with the emergence of humans as the dominant species. Until recently, the entirety of human civilization has evolved within the context of this single epoch over the course of three ages: the Hunter-Gatherer Age, the Agrarian Age, and the Industrial Age.

One of the consistent characteristics across the three ages is the persistence of the top-down hierarchical model of social organization. While our social structures have evolved with each age as we have morphed from tribes to bureaucracies, the one common proposition that persevered across the ages was the unquestioned assumption that, for human organizations to work, somebody needs to be in charge. Whether you were a tribal chief, a monarch, or a CEO, power belonged to those in charge and they had the wherewithal, either through force, law, or position, to exert their will over large numbers of people. Even modern-day democracies are a variant form of the top-down hierarchy, the difference being that we get to elect the people who will be in charge.

The digital revolution is a watershed moment in human history because its technological innovations have made it possible for the peer-to-peer network to be what it never could before: a practical form of social organization. Digital technology has provided the ubiquitous means for people to self-organize within the context of hyper-connected networks in ways none of us could have imagined a mere two decades ago. For example, who would have thought that the world's largest reference work would be a self-organized effort of volunteers working without assignments and without pay? None of us could have conceived of an enterprise such as Wikipedia. After all, how would anything get done if no one was in charge? Surely, such an enterprise would be powerless and doomed.

As we now know, things did not turn out as we might have expected. The Wikipedia self-organized network didn't just survive, it thrived and created a level of productivity far beyond anyone's expectations. And, without any intention to do so, it ended the reign of the 244-year-old market leader within a mere decade. How did this bossless structure become so powerful?

Perhaps the most important thing to understand about this new form of social organization is the way it radically transforms the way power works. In networks, power belongs to those who

are connected, not to those who think they are in charge. And when a network takes firm hold in the market place, it tends to operate far smarter, far faster, and far more efficiently than hierarchical counterparts. The unprecedented benefit of networked collective intelligence structures is that they open up the door to the practical ability to do, in Kahneman's terms, System 2 thinking at System 1 speeds and, thus, greatly reduce the ill effects of individual human biases that are the inherent hazards in the top-down hierarchical structures that have defined the first human epoch.

The continued rapid emergence of the peer-to-peer network as the dominant form of social organization in combination with the use of collective intelligence dynamics to shape the development of AI systems will ignite an extraordinary leap in the social exercise of human judgment and human decision-making. Contrary to the concerns of those who are fear that AI will somehow create an army of highly intelligent individual agents to rise up and subjugate humans, AI holds the potential to become a powerful extension of human intelligence. That's because the ultimate manifestation of AI will not be a collection of mechanical agents, but rather a real-time collective intelligence system that holistically integrates the best of both human and machine learning. In other words, AI—especially if it's built on a blockchain platform—holds the great potential to transform individual agency as we have known it.

For example, as described earlier, driverless cars will not be a collection of individual agents navigating the road the same way human drivers do today. Instead, each autonomous vehicle will be part of a holistic collective intelligence network that will allow every car to navigate the road with the real-time knowledge of what every other vehicle in the vicinity is doing. This is possible because AI networked intelligence is not limited by the information storage capacity of the human brain and is capable of instantly transferring the complete knowledge in the AI system to any of the individual vehicles. This collective intelligence system will not take away the freedom for the individual vehicles to make decisions and act. Rather it will inform those decisions to optimize choices in ways that do not harm the overall system. That's why the number of accidents will drop dramatically at the same time that speeds will probably increase. Individuals will be free to make choices that will do no harm but will be prohibited by the collective intelligence dynamics within the system from willfully crashing into another car.

Individual agency will be transformed because our systems will no longer rely upon traditional control mechanisms to execute activity. This means that no single individual agent will be able to seize unilateral control over any activity, which will greatly diminish, if not eliminate, the fraud and hacking so prevalent today. Instead, every action in networked blockchain systems will need to be validated by collective intelligence protocols in real time before an action can be executed. While this may be hard to imagine, none of us thirty years ago could have conceived that we would build a system that could search the entire world's information stocks and present us with intelligent results in a fraction of a second. Today we take Google searches for granted.

The shift from hierarchical individual agency to networked collective agency, which is the inevitable consequence of a hyper-connected world, will be the single most consequential event in human civilization, and it will remain so until the second human epoch runs its course. How we build and lead organizations, the way intelligence works, how we create and distribute value, and what it means to be human will all rapidly evolve in ways that, at this point, are beyond our imaginations.

However, for an insight into how the world may work very differently and how humanity might be transformed, chances are our new ways will reflect the wisdom of the two foundational principles, referenced earlier, which were used by Chris Rufer, the founder of Morning Star, to design his self-organized networked company: 1) no individual should use force against others or their property and 2) individuals should keep their agreements with others. The socialization of these two simple rules in the context of a highly intelligent human-machine symbiosis that moves us beyond the hazards of individual human bias will indeed transport us into a second human epoch that will leverage the remarkable power of AI collective intelligence to radically transform humanity for the better. This is the great promise of digital transformation.

Rod Collins (@collinsrod) is the Innovation Sherpa at Salt Flats and the author of Wiki Management: A Revolutionary New Model for a Rapidly Changing and Collaborative World (AMACOM Books). Rod Collins moderates collective intelligence and design thinking workshops for clients of Salt Flats, as well as leads the Innovation Leadership Experience training which equips leaders with the skills, tools, and competencies they need to lead highly resilient organizations in these digitally disruptive time. For more information on services at Salt Flats, visit www.saltflats.co.