THE IMPACT OF AI ON BUSINESS STRATEGY

By Bill Achtmeyer, Bob Goad, and John Trustman

There is no shortage of current opinion pieces on the likely impact of artificial intelligence (AI) and the current "hot" AI topic, generative AI (GenAI). The work ranges from serious and high quality to fatuous, and the distinction is further complicated by the fact that the better pieces tend to be written by the more technical writers for more technical readers. Often these pieces are suffused with complex algebraic expressions difficult for anyone without a mathematics or engineering background to fully comprehend. A working knowledge of linear algebra and calculus is a pretty high bar for the average reader. If "gradient descent" seems a more suitable topic for a ski trip than a first-order iterative optimization algorithm for finding a local minimum of a differentiable function to you then stay with us here!

Some of the best work in AI has featured breakthroughs and accomplishments in the creative arts, music, and painting created by GenAI which is clearly spectacular, if not – understandably – original. The very nature of GenAI favors derivative works, something "in the style of" someone. There are very strong pieces on the impact of AI on business as well. Most of those pieces feature the likely impact of AI on various professions, workflows, and tasks. This work is tactical rather than strategic. "Strategy" with a modifier other than "business" or "corporate" in our context, or perhaps "war" or "grand" in others, is not strategy but rather tactics. "Marketing strategy" may be important, even critical, but it's tactics, not strategy. It's part of the "how" question, not the "what" or "whether" questions. That's not to say either that marketing isn't strategic, as it often is, or that marketing efforts don't need a strategy, which they imperatively do. But marketing is a tactic in achieving the business' strategy. The same applies for the rest of the business strategy adjectives – financial strategy, HR strategy, logistics strategy, sales strategy, technology strategy, just to mention a few. All important and each often critical, but tactics, not strategy. By "strategy", we mean something like:

A general framework for achieving a specific outcome for an entity with constrained resources under competitive conditions of uncertainty.

We also need to do a clear job of what we mean by "artificial intelligence". The complexity we face here is that in any field as (relatively) old, dynamic, and varied as artificial intelligence, many technologies and more approaches have claimed the label. We find the relevant distinction to be Andrej Karpathy's. Karpathy was a founding member at OpenAI where he has returned and was the Senior Director of AI at Tesla. He coined the term, now in wide use, "Software 2.0" which refers to code not written by human engineers but written rather by an optimization function of an evaluation criterion against a set of observed data. It might take the machine a (relatively) long time and a lot of energy to find the solution, but the machines don't get tired.

"Software 1.0 is code we write. Software 2.0 is code written by the optimization based on an evaluation criterion." – Andrej Karpathy

---

Artificial Intelligence is the proper name for the software that the machines produce. If the humans are developing the heuristics and writing the algorithms, it's Software 1.0. If the machines are doing this work, even if the machine's algorithms are written by humans teaching the machines to use logical reasoning, it's Software 2.0. And when the machines come up with better heuristics and algorithms on their own, it's working very, very well.

This is by no means the only definition of AI we encounter, but it is the definition we use as we think it's the clearest. It sets a fundamental boundary between Software 1.0 applications, even great ones, including a lot of business intelligence, analytics and optimization programs, and Software 2.0. This distinction is critical, because what you tell the software you want it to do is fundamentally different between the two generations. Ironically, this can become critical in intermediate situations where layering in classic human heuristics can speed up the time to a solution at the cost of limiting or even interfering with what the AI can develop on its own. While humans are better at certain types of reasoning and calculation, computers are better at others, and the shortcuts that people take are often suboptimal when extended to AI.

AI and Games

If that's entirely clear, you can skip this section. Otherwise, consider four games – tic-tac-toe, checkers, chess, and go, each played entirely without a chance or luck element. Looking to games is perfectly reasonable – the academic classification of strategy is "behavioral game theory", which is game theory with the requirement removed that all players behave rationally. And, as we have come to learn, game theory is the fundamental math behind both conflict and economics.

Simple Games

In tic-tac-toe (or naughts and crosses if you are reading this in Europe) players alternate playing pieces – typically Xs and Os – on a three-by-three grid until someone gets three in a row to win or all the places are filled, which is a draw. We call the game simple because there are "only" 26,830 possible game configurations if you take symmetry into account – playing first in any corner is essentially the same thing, you could just rotate the board. If you don't take symmetry into account the number goes up to over 250,000, which might be a lot for a person but is quite trivial for a computer to solve by brute force. Staring at the board for a few minutes and playing a few games, most smart people can figure out a winning strategy on their own. Either way, the game has been solved; a good player, human or computer, will never lose and will win so long as their opponent makes a mistake. Although you can use AI to play the game, you don't have to.

Checkers (draughts for the Euros) is harder, but it too has been solved with Software 1.0. Checkers is a strategy board game for two players which involves diagonal moves of uniform game pieces and mandatory captures by jumping over opponent pieces. The game is played on the dark squares of an eight-by-eight board of alternating colors. Pieces reaching the far rank are "kinged" or promoted and can move backwards as well as forwards. Checkers has approximately 5×10^{20} (that's 500 quintillion in the US short scale and 500 trillion in the European long – either way a very large number) possible positions with somewhat more possible moves as kings can cause positions to be reached in multiple ways. If you played a move a second, it would take you 265 billion years to play all of the possible moves, which is quite a bit longer than we believe the universe has existed. (Admittedly, some of the moves would be illegal, but it's still a very long time.) By the 1990s computers programmed by humans were beating mere humans at checkers and won their first championship in 1994 using algorithms based on expert play. That's human heuristics with algorithms that the machine is executing – Software 1.0.

But computers can play a lot faster than humans and by 2007, even on computers that were incredibly slow by today's standards, the game had been solved, with all possible moves and positions evaluated so that the best an opponent to a proper checkers computer program can achieve is a draw, no AI required.

Complex Games

Chess is harder. Chess is qualitatively different from tic-tac-toe, checkers, and go (which we will address next) in that the rules are anything but simple. Played on the same board as checkers but using all of the squares, chess has multiple varieties of pieces, each with multiple capabilities and limitations. We include it here because it is the most widely known complex game of pure intellectual skill with over 600 million players world-wide. It is estimated there are more than 10^{40} legal positions in chess, which is a very, very large number. The number is sufficiently large to have a name, it's the Shannon Number. This is not a game which can be solved by Software 1.0 with today's computers or likely any we will see in our lifetimes. But a Software 2.0 approach has already yielded significant results. The best Elo chess rating by a human ever is held by Magnus Carlsson at 2,882. (Bobby Fischer's was 2,785.) In contrast, today's Stockfish chess engine ranks at 3,550. Stockfish is a combination of human algorithms and an AI neural network, NNUE, which has helped it improve its Elo rating by several hundred points since its incorporation.

In terms of computational complexity, go makes chess look simple. Played by some 40 million people predominantly in east Asia, the game is played by two players alternating placing black or white stones on a nineteen-by-nineteen grid, attempting to surround areas of the board. Although the rules are straightforward and the pieces homogenous, there are over 10^{170} possible moves in Go making it computationally unapproachable from a human algorithmic standpoint. But DeepMind (now part of Google) AlphaGo and successors have far exceeded even the best human performance using a Monte Carlo tree search algorithm to find its moves based on knowledge previously acquired by an artificial neural network. AlphaGo was trained by playing thousands of games against human players; AlphaGo Zero was trained by playing itself, and then played against the strongest human player in the world. The current version, MuZero, learns without even being taught the rules. This is Software 2.0.

AI and Business

The critical question posed by recent developments in AI is not whether but how AI will affect you and your business. Hyperbole aside, AI is very real. Ignore at your peril; ignore for too long and the damage will likely be irreparable. Rivals who embrace AI will end up with a sustainable competitive advantage over those who wait. On the other hand, excessive spending and management distraction could be detrimental. The solution to the "wait equation" will be dictated industry by industry and competitor by competitor.

For context, consider prior industrial revolutions. The shift from human to machine labor is remembered as a revolution because it changed the fundamental way that goods were produced, but what made it revolutionary from a strategic perspective was that it changed the sources of competitive advantage. By disrupting cost structures and introducing new products not possible with reliance on human labor, the industrial revolution changed the nature of competition. It was revolutionary, but it is viewed as revolution largely because the entrenched leaders did not emerge as such. Why not? The leaders going into the period had all the strategic advantages we have come to recognize – the highest relative market share, the best cost structure, the best profitability and were best positioned to be able to invest in the new technology. And yet not only did they did thrive, many failed to survive at all. The same has been

largely true of subsequent industrial revolutions. In the penultimate episode, the introduction of the internet, there has been relatively less upheaval, but few of the pre-internet leaders have survived in leadership positions despite overwhelming advantage at the outset. Even Microsoft, who has fared much better than the other leaders, had its since regained leadership position largely evaporate as it failed to recognize the threat and opportunity.

AI will have a more pervasive and closer to ubiquitous impact than the internet. There can be reasonable debate as to whether AI and the internet are part of the same industrial revolution, the digital or information revolution, or separate. But that debate is somewhat incidental. Certainly, the breakthrough in today's AI builds solidly on the back of the explosion of information and information availability enabled by the internet as well as the availability of high-performance hardware. Software 2.0 builds on the foundation of Software 1.0. No assembler, no C; no C, no C++; no C++, no Python; and no Python, no PyTorch or TensorFlow. The evolutionary roots of most revolutions are well established.

Given their structural advantages, we need to understand why some, and likely many, leaders will not successfully transition. The core issue is that when the drivers of competitive advantage shift it is much more difficult for entrenched leaders to recognize and realize the opportunities of that shift. Experience suggests that leaders are more likely to see the shift as a threat rather than an opportunity. They will, properly, recognize that shifting to the new basis of competition will likely put them at a diminishing advantage and may well place them at a disadvantage as the investments, skills, and instincts that they have developed will all be challenged or made obsolete. The natural though counterproductive reaction is to resist that change. Business self-cannibalization is bad but often preferable to just being eaten by aggressive competitors.

The AI Adoption Cycle

Artificial intelligence will affect different industries and functions in differing ways. We'll likely see the same progression of impact, described below, as we've seen with other technological innovations. While this progression is "natural", its nature is based on risk management rather than value creation. This will create strategic opportunities at each stage of transition as well as creating systemic leadership disadvantages which will likely lead to disruption as leaders will be less inclined to implement the painful business model realignment made possible and then necessary at each stage. While it may be possible to ignore some of the changes for a period like the frog in slowly boiling water, eventually you get cooked.

An historical perspective suggests that technology driven breakthroughs unfold in a predicable sequence:

- Initial adoption will be in specific cases employing AI to automate a single task, either a standalone task or a task which is part of a well understood value chain Examples include asking a customer for the nature or subject of their inquiry. AI has already largely replaced humans as the first point of contact in customer service inquiries, typically asking the customer to describe the nature of their issue. While most of the current systems use AI driven voice recognition and word matching only, newer systems enable the AI to manage the next steps as well. These tasks initially have manual backups, though over time, the manual processes become less and less necessary or even possible. Earliest adoption will follow industries and functions with high volume, repeatable and low value and/or low risk interactions, such as initial problem diagnosis or requesting customer feedback. These are typical applications of "Level 1" customer support. Diagnosis automation will usually provide for a mechanism to "opt out" and to request a human

representative. Predicting widespread implementation of AI driven retail customer service chatbots is hardly insightful today, it's been underway even before the introduction of GenAI.

However, the nature of these interactions will mature and expand. Initial interactions with these chatbots were typically only web and text based, and whether on the web or even over the phone, customers were generally directed to preconstructed information, such as the relevant section of a manual or "frequently asked questions". With the advent of generative AI, many companies are introducing proper "bots" which create customized responses to customer inquiries on the fly.

These responses are currently mostly text based, but recent progress in both audio and video AI production is rapidly leading us to an environment where a customer will be receiving voice or even a video interaction with an AI customer service agent. That agent will have the advantage of being familiar looking and speaking in a dialect and accent that supports the customer's preference rather than the economics of the company's customer service offering. Over time, the focus of this function will shift to value creation rather than cost minimization, as beyond the initial "training" costs of the AI which are becoming lower than the costs of training a cadre of customer service agents, the AI agents ongoing expense is negligible.

- As technology progresses and customers and companies become more accustomed to interacting directly with AIs, we expect to see full value chains rather than simply individual tasks replaced with AI. This will result in an explosive uptake of the technology. Imagine personalized evaluation and recommendations delivered by AI replacing human sales "robo-calls" or even group seminars. Instead of sales pitches constructed around average needs with an agent listening for an opportunity to transfer the call to a somewhat more experienced agent, AIs will conduct customized interactions based on customer specific information.

  The major difference for the business between single task and full value chain digitization is the level of commitment. While single task automation will likely have either human backup or even be conducted in parallel with human delivery, full value chain digitization will require the eventual abandonment of any pretense of human involvement in the tasks. Most of us have forgotten when banking transactions and securities trading were largely manual tasks. When computers and automation were introduced for these tasks, we had copious manual backup procedures. The introduction of ATMs was accompanied by fears that customers wouldn't trust the machines, fears that were quickly proven unwarranted. Today, if the computer is down, the bank is out of business. By the 1990s Ned Johnson, who owned and ran Fidelity Investments, used to quip that if Fidelity's computers were down for three minutes there would be no more Fidelity and likely a market disruption.

  Taking advantage of this innovative technology will be painful for companies with large, existing, low-level sales and service forces who will be subject to wholesale replacement. In some cases, the natural attrition in these jobs may be sufficient to manage the transition, but for most the shift to a largely automated force will be difficult. CEOs who built successful careers building these teams may struggle to disassemble them in the face of downsizing pain and removing sources of prior competitive advantage.

- The third and most important stage in the business AI evolution will be the introduction of value chains not possible without the new technology. Only consumers older than thirty have been alive at a time when Amazon did not exist; half of the people worldwide have never lived before

or without on-line shopping. Flat screen televisions are a bit younger and smart phones a bit older. The first iPhone was introduced in 2007. Facebook was founded in 2004. VisiCalc, the first computerized spreadsheet was introduced in 1979, Lotus 1-2-3 in 1983 and Excel in 1985. People have forgotten that spreadsheets were initially fourteen column pads of paper and that the original "Facebook" was a forty page, 3" × 6" book of names and postage stamp sized portraits, Zuck's high school's universal student Rolodex. Value chains entirely dependent on Software 1.0 technology have permeated or taken over industry after industry, and humans have shown an indifference to interacting with another human or machine. Rather, we often want to interact with the most efficient, individualized, and expedient solution to our problem or satisfaction of a desire. Software 2.0 based solutions are much more promising.

Each of these phases will introduce both opportunities and risks. The scale increases with dependence on technology, and first mover advantage is by no means guaranteed, though lagging by too much will likely lead to certain failure. Of the ten most valuable companies today, only three existed before the introduction of the personal computer and the internet and all seven of the newer companies are directly involved with and leading the development of AI. But that leadership has been fraught with peril. Apple succeeded where MITS and Commodore did not. Microsoft is the second most valuable company; their initial "big brother partner" IBM barely makes the top one hundred. Alphabet/Google has thrived as has Meta/Facebook, Yahoo and Myspace are gone. One of the authors encouraged the Olim brothers to expand CD-Now past music to be scolded that music was a better business model than books. Amazon is a powerhouse in multiple arenas; the Olin's CD-Now is now gone along with the CDs it sold. As the rollout of AI progresses, we will see a plethora of new entrants, new products from legacy entities, as well as tangential participants. There are already hundreds of recognizable players. With the introduction of all new technologies, most will fail.

To understand what makes for an AI opportunity for your company you need both to understand what the AI does and how it "learns" how to do it, and more importantly, an even more explicit and detailed understanding of your business strategy and competitive and customer landscape. With humans you can take much for granted, with machines, even very smart ones, much less can be simply assumed. Finding the overlap of AI capabilities and your business' opportunities requires the kind of skill set rarely found, expertise in both technology and business strategy.

Behavioral Game Theory

Behavioral game theory is a relatively recent development in mathematics and economics. It is a blend of behavioral economics and game theory. It is perhaps best understood as game theory where there is no assumption that the opponent will necessarily behave in a rational or optimal manner, academically speaking, "relaxing the rational man constraint". Unlike games, in business we may not agree on the rules and even if we do agree, we may not strictly adhere to them. With business AI this leads to two material differences - finding optimal training data and robust evaluation criteria is typically more difficult and the willingness to defer to a counter intuitive recommendation is likely to be rejected or at least limited.

Trust

In what is generally considered one of the breakthrough moments for modern AI, the DeepMind AlphaGo AI defeated world go champion Lee Sedol four games to one in March 2016. While the entire match was quite remarkable on many levels (Sedol offered that while he had previously represented himself, his club or even his country, he had never before thought of representing the human race), we

consider one in particular here. At move 37 in the second game AlphaGo made what was initially considered a mistake, a wasted or "slack" move. Eventually, the move has come to be considered "creative" or "inspired", but at the time the commentators viewed it curiously at best. The initial reaction was that it was a mistake. But AlphaGo had already unexpectedly won the first game, so the experts were circumspect, especially as Lee Sedol initially appeared confused and then concerned.

The move might be considered the equivalent of a AI pricing algorithm suggesting a 50% price reduction for a critical SKU. With only a game of go on the line, this is perhaps an acceptable risk, but if the decision puts a business line or even a business at risk, would managers have the conviction and nerve to follow it, especially given that the AI does not, and frequently cannot, explain why it makes its recommendations? Certainly, a human manager would be more comfortable with a smaller reduction. But if a 50% reduction is optimal, a 20% reduction, or even a 45% reduction, might be damaging or even fatal; graceful degradation is often an oxymoron. The very nature of AI optimization requires a choice between finding an optimal strategy and an optimally acceptable one, but the latter may not lead to anything close to the former. AlphaGo did not have a human minder, most business AIs likely will. How we come to grips with these divergences from human performance may well control our ability to derive benefit from AI and where we will use it.

Training Data

The three key requirements for implementing an AI are having the data that the AI will be trained on, having a reliable evaluation criterion, and having a viable mechanism for evaluating improvement and success.

If you're trying to teach an AI to produce text you need the appropriate material to train it on, a "corpus". If you are trying to teach an AI to write or speak your corpus will be language, if you are teaching it to paint, your training set will be paintings. In either case you will need a large enough set of data for the task at hand, and, depending on the goal, a sufficiently high quality data set. A high-quality data set doesn't necessarily mean just good examples, but will usually include bad ones as well, depending on the evaluation criteria. A common "hallucination" of Generative AI is what is called the "WaLuigi effect" (after the Mario Bros. sidekick's evil twin) where the AI will return the exactly opposite of the indicated answer, which turns out to be a very near neighbor of the correct one. Given that AIs are frequently trained on data found on the internet, we can expect lots of data to train on, but as much of that data may be wrong as right.

Examples aren't the only factors in the quality of the training set. Size is also critical, especially in tuning an AI. Tuning an AI is the process of taking an existing AI, probably trained on a huge body of material, and giving it a much smaller and more relevant set of material to draw from in its primary task. The process can be thought of as the base AI learning to speak, or even to reason, and the tuning data teaching it a specialty. The cost of training the AI can be huge, the tuning set needs to be much smaller as it may vary by application or even by query. In practice, the training set for a base AI might be millions (or more) of pages of text; the tuning set will be a few hundred or thousand at most. But too small a tuning corpus and the output won't be on point or will generate overly repetitive and less helpful outcomes.

For most companies, building this data set, whether taking a training or tuning approach, requires both external and internal data. Having a meaningful process for collecting, cleaning, and qualifying that data is a new process for most companies. This is another area where data truncation, having deleted older data, may prove to have been short sighted. We worked with a company which made a practice of

deleting all client-oriented data older than a month. This has effectively destroyed their ability to build a strong custom corpus. Most companies don't even have good enterprise-wide search capability; they can't easily find anything, let alone everything. And even if they could, there are frequent privacy and security concerns that may compromise those approaches.

More importantly, from a strategic standpoint, the training and tuning corpora represent the past, and the past can get old very quickly. One of Bard's strategic advantages over ChatGPT had been its internet search capability which gave it access to up-to-date data. (That capability also caused security and privacy concerns; there are no simple answers here.) Missing even today's data can make an AI substantially less effective; worse, it will cause a user of that AI to look elsewhere as well, causing a potential disruption point. Current data is also generally the most important. A study we conducted at leading financial services firm found that customer conversations were relevant inversely proportional to their age – the most recent conversations were the most meaningful. Not having access to yesterday's data, let alone today's, can be a handicap. Having a process to incorporate it can be critical.

Business Data

The imperative is properly preparing – training and tuning – your AI platform for the desired task. Typically, we use a mix of three distinct types of data, behavioral, sentimental, and incidental.

Behavioral data is data on *actions* taken by various players – the company, its customers, suppliers, and competitors, as well as third parties such as new entrants, alternatives, and regulators. (If that sounds a lot like *Five* Forces it's not a coincidence.) It's the historical record of what was *done*. Purchases are the primary example. Behavioral data is the most valuable data as it is intrinsically accurate, but it is not always easy or efficient to acquire and often not available until too late to be useful. *Sentimental data* is data on how the players *feel*, or more accurately how they say they feel. It is the historical record of what people said, to be taken more as stated opinion than fact. Sentimental data is often easier to accumulate in advance of action and can be used to supplement understanding and often to predict it. But sentimental data is of far lower quality; people are given to saying how they think they are supposed to feel rather than how they actually do. *Incidental data* is related data, often helpful, maybe even critical, though not directly connected to the business. Economic conditions are the best known example. Traditionally we make judgements on what data we input to construct our algorithms. With AI we have seen that the algorithms often do a better job at selecting and prioritizing the most important data than do humans. But this is a beast with an infinite appetite, and feeding it is an art. Selecting among these data types and the proper blend of inputs form the data sets is of paramount importance.

Evaluation Criteria

Being able to run your AI against a good set of data is the necessary first step, but it is not sufficient. The models need a workable and meaningful set of evaluation criteria. In tasks such as training a base AI, the task is inherently easier, good output can be determined statistically – does the output resemble the input corpus? If you are trying to emulate Shakespeare, can the output, or something like it, be found in Shakespeare's works? In a business simulator, do the decisions look reasonable? In other scenarios we may have a preexisting set of training and testing data that allows us to check if the AI is getting the same answer as a human, which we (sometimes appropriately) consider correct. In some settings we can use reinforcement learning, an approach which rewards the AI for doing a better job. In gaming, as we saw, the AI is "rewarded" for a win and "punished" for a loss. One of AlphaGo's "insights" was that the goal of the game was simply to win, and that improving the odds of winning was more important than winning by a large margin. But in business it's not always as easy to identify what winning means,

and very rarely possible to know if you have won. Games have a final condition, business can, but rarely a positive one.

The most important factors for evaluation criteria are clarity and objectivity. Clarity is often difficult and is an area where the business leader's input is critical; the evaluation criteria that the data scientists will be drawn to are not necessarily the best criteria from a business perspective. Different from AlphaGo's insight about the difference between winning and winning margin, the choice of a standard objective criterion for a business problem is rarely straightforward. For a simple sales task it might seem to be whether the customer bought something. But is that right? Is the goal a short-term purchase or something more sophisticated, such as improvement in long-term customer value? And if long-term, how long? We worked with a salesman who was phenomenal at selling something right now but had very few repeat customers. Good or bad? It's a similar problem to building a sales compensation plan, but in many more simultaneous dimensions.

In most settings, a single interaction will rarely lead to an optimal outcome; those are the product of a series of interactions. In these situations, using an objective, transactional criterion for the interaction may lead to suboptimal results. Here we will need both to evaluate the model's performance over a longer period and may well need to depend on qualitative and sentimental, and often even incidental data. This is not a simple task; transactions and relationships need to be well understood and modelled in a way which is closer to probabilistic scenario planning than simple impact. Each interaction will need to be understood in the context of the long-term goals, which makes defining those goals explicitly even more important. To do this we must get away from the old standby "high", "medium" and "low" cases and develop these scenarios specifically along with the probabilities of achieving each one.

In the well-known Delta Airlines customer service example, the airline discovered in comprehensive testing that passengers who experienced flawless on-time performance were less loyal than those who had experienced a disruption but with exceptional service recovery. Would an AI solving for customer loyalty optimize for on-time performance or would the AI look for opportunities to schedule some disruption with planned service recovery? And if that strategy became known, how would that impact loyalty? How would management react to such a suggestion, or would they even know, as the AI might simply sub-optimize route scheduling while altering staffing levels? How would the regulators react?

On the other hand, AI evaluation criteria also offer complex possibilities less present in other optimization models. Mathematical optimization models require either a single factor to be optimized or an algorithm for managing tradeoffs between multiple criteria. We have already seen that AI models can be "optimized" more broadly; the typical generative AI evaluation criterion is closer to acceptable or commonly found output than correct. Other optimization mechanisms have much more difficulty in distinguishing what makes an acceptable answer. Critically, what this can enable is the ability to optimize over multiple criteria simultaneously without establishing explicit tradeoffs or treating all but one of the criteria as constraints rather than the objective function. AIs are much more capable of handling less explicitly defined tradeoffs between criteria. An AI could be an invaluable companion in managing tradeoffs between various stakeholders. The *Acropolis Advisors 21st Century Stakeholder Optimization Framework and Dashboard* can provide a great starting point for thinking through this more complex optimization.

The second critical factor in the evaluation criterion is objectivity, or at least, repeatability. Like humans, AIs learn from repeated processing. Like humans, we usually want our AIs to learn efficiently and inexpensively, but AIs typically learn over more iterations. If a task involves human feedback, typically as reinforcement learning from human feedback (or RLHF), it will be relatively slow and expensive.

Moreover, humans generally need to be paid and take time off for other activities, like eating and sleeping. AIs requires no such inconveniences nor interruptions. If the evaluation criterion can be automated, the AI will learn much faster and less expensively. This is the major reason that we see a substantial focus on teaching AIs to learn without human intervention, to build their own training and evaluation sets. Moreover, automated objective functions are generally much more consistent than humans and provide more reliable feedback to the AIs. Providing inconsistent evaluation to an AI can produce "interesting" results. The good news here, at least for AIs, is that recent research has shown that RLAIF – reinforcement learning with AI feedback – is showing as useful as RLHF.

The difficulty in a business setting is properly establishing an objective function. Former HBS professor Ken Merchant offered the best starting function for business evaluation.

Merchant's Law:
Cash is truth.
Corollary #1 – More is better than less.
Corollary #2 – Sooner is better than later.

But even this can be problematic, mostly due to uncertainty around timeframes. Public company management has been famously short sighted in the past; quarter-to-quarter performance targeting has long been shown to be sub-optimal, and the pressures of today's "stakeholder" based capitalism is even more complex and onerous on top of what was already a near impossible job. Baking in inconsistent criteria can cripple a model.

Evaluation Mechanism

With the data and good direction on evaluation criteria, we can train and tune the AI. For most executives, the strategic issues are going to be which tasks and projects to undertake, and when and with whom to build them. The critical logical issue is going to be how far to trust the AI, which will be driven by an issue we surfaced above – what to do about intermediate results. In games we know what winning means, and for moves along the way we can generally ask whether the intermediate moves moved us closer to winning. The technique is called "backward reasoning" or "backward chaining" and is best understood as working backward from the goal. But if the ultimate objective is not well understood or defined, or the path to it is too long or arduous, we need another approach.

It may seem odd to focus on intermediate results, but the evaluation criterion takes care of the issues around ultimate results if an ultimate period can be defined, and what defines the mechanism for getting to them is how we measure and manage progress towards that goal. If we are optimizing for five year market value, how should we treat cash flow after each year?

The major factors driving the evaluation process are the business model, the structure of the model, and the nature of the intermediate goals. The model and structure are the core elements of development and productivity in the AI world. These are evolving almost daily. As discussed below, until we see more stability and the emergence of leaders in various tasks and industries, this will be a task best handled with a combination of internal and expert external resources with aligned goals.

For humans who have less capable detailed memory and computational ability than machines we rely on heuristics. Heuristics are shortcuts that we use to solve problems. Essentially, heuristics allow us to

solve an easier problem as a stand-in for a more complex problem that we can't solve. Heuristics are often equally important to AI both for problems that are too complex to solve in a more straightforward manner, but also where the explanation for why the AI is choosing to do something is either absent or too complex to explain (or to understand). In these situations, a heuristic may offer a more workable solution. But using a heuristic means that you are not actually solving the problem but rather addressing a related, and often, but not always, correlated, problem.

Consider the use of relative market share (RMS) as a proxy for gross margin targeting for a business. In forty years of practice across a broad range of industries and tens of thousands of companies, we have found that RMS is significantly correlated with margin. The correlation is so strong that we often believe that pursuing a business strategy to improve RMS will be the optimal strategy for improving both profitability and the value of a company over time. Consider Amazon's strategy in its online shopping business, the business for which it is most well-known. An examination of Amazon's shopping RMS and margin over its lifetime shows that Amazon's strategy had been to eschew profitability in the pursuit of RMS. Effectively, by keeping prices low, Amazon "bought share". This strategy has taken Amazon from an unprofitable, minor player in consumer shopping to a dominant position in a decade. During the same period Walmart was willing to give up share to build additional margin. But RMS does not equal profitability, rather it closely correlates with it over time. Walmart reestablished its industry leadership acknowledging that RMS was as important than margin, and Amazon only achieved positive margins by relaxing its commitment to buying market share. If we train an AI on RMS, we must equally teach it when to break away from the heuristic or it will keep prices low forever despite diminishing gains in strategic position. What constitutes an acceptable algorithm is complex; in this case there will also be input from the courts.

The critical question becomes whether we will impose some general understanding or insight on the problem, or whether we will leave the AI to its own devices. As discussed above, there are strengths and weaknesses to both approaches. If we provide a heuristic, we make the AI more efficient, but we do so at the risk of disabling it from finding its own best way forward.

Strategic Management in the Age of AI

What then will the impact of AI be on the factors that drive business strategy? To best understand this, we go to the source documents of business strategy, the three fundamental frameworks we have used for decades to understand, diagnose, and prescribe strategy. We ask if and how these change with the advent of AI. Fortunately, the authors of two of these frameworks are among our founders (including the authors), so we have the benefit of going directly to the strategy horses' mouths.

The most basic strategy framework is the *Three* (or *Four*) "*Cs*" – customers, costs, competitors (and capabilities) framework developed by Kenichi Ohmae at McKinsey and Company. While obvious in retrospect, it represented a breakthrough when first presented and has stood the test of time. While AI will have a marked impact on all four factors, it doesn't change the paradigm. What AI does do is to alter the tactics associated with each of these dimensions. We don't dwell on the *Three Cs*, largely because it has been fully subsumed and improved on in both Porter's *Five Forces* and Achtmeyer's *Full Potential Paradigm™* analysis.

The second of these fundamental frameworks is Acropolis Advisors co-founder and Harvard Business School Professor Michael E. Porter's ***Five Forces***. Porter's analytic is pervasive in modern strategy conversation. The framework is comprehensive, durable, qualitative, and never trivial to apply, especially when things change.

When we first encountered *Five Forces* in the 1980s and 90s, all the attention was on the horizontal axis – Supplier Power, Industry Competitiveness, and Buyer Power. Most industries were still in some phase of the adoption curve, and we paid short shrift to the notion of likely competition from alternatives or substitutes, let alone the significant hurdles facing a new competitor with a similar product or service. Then we saw information technology driven disruption change everything. In today's business environment, we think first of the threat of substitutes and new entrants. Many industries have reached the point where the adoption curve has plateaued, and the substitution curve has taken over as the primary driver of opportunity. Substitution is most often driven by the introduction of new technology. We saw an article a few years back which claimed that only the threat of substitutes was now a relevant force. Porter's framework is holistic and enduring; omitting forces is a serious mistake and leads to vulnerability.

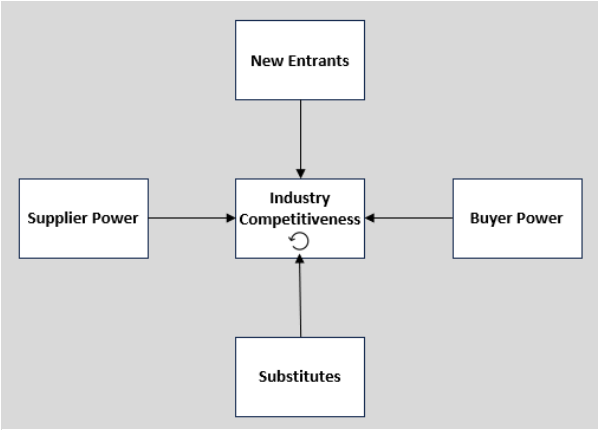

Figure 1 - Five Forces Diagram

While artificial intelligence will impact each of the forces and may, especially in the short run, be a key factor in determining power, we don't see AI creating a new "force". AI will accelerate and strengthen the threats of new entrants and substitutes and will eventually be deterministic in establishing supplier and buyer power as well as upending industry rivalry, but it won't alter the paradigm.

The ***Full Potential Paradigm™*** is Acropolis Advisors co-founder Bill Achtmeyer's strategic analysis framework that uses mathematics and quantitative analysis to understand a company's revenue growth, margin, and valuation multiple in the context of its industry. Achtmeyer is best known as the founder of The Parthenon Group and subsequently leader of EY-Parthenon. *Full Potential* offers the advantage of objectivity and mathematical repeatability, which have become increasingly important in strategic analysis and will be critical in AI driven strategy. Based on concepts extending from "maniacal realism", our shorthand for demonstrable, fact-based analysis, the *Full Potential Paradigm™* delivers a quantitative explanation of profit margin related to competitive position, pricing and cost leadership and a similar decomposition of revenue growth from being in high growth markets versus doing well in the current business mix. Additionally, *Full Potential* delivers an understanding of the drivers of market valuation on both a standalone basis and in the context of competition. Using *Five Forces* and *Full Potential* together provides a powerful combination of qualitative and quantitative platforms for understanding an industry in its competitive state.

The application of AI will dramatically enhance each component of the *Full Potential Paradigm™*, making its understanding even more critical. The *Full Potential Paradigm* equation of

Margin × Revenue × Multiple = Market Value

neither depends on nor is altered by AI, but remains paramount in assessing a company and its strategy.

As with all systemic technologically driven advances, AI will evolve from delivering strategic advantage to becoming a mandatory utility. Companies which fail to adapt will simply fail. This is not a new phenomenon. An analysis of top performing companies on the NYSE over each decade since 1940 found that the probability of a company making the list in three successive decades was lower than that of a company no longer being a going concern over that period. Of the 66 "hottest" company stocks as of 1950, only eighteen still exist as independent companies today, although another twelve still have their name as part of the successor that acquired them. The introduction of AI will be disruptive; it will introduce significant structural disadvantages to entrenched players due to a vastly changed business model causing great difficulty in adapting.

That is the long run, and as John Maynard Keynes wrote, "In the long run we are all dead."

A *Five Forces* and *Value Chain* View

A *Five Forces* analysis immediately reveals that artificial intelligence will serve to intensify competitiveness at every strategic interface point in virtually all sectors of the economy. While much of the initial focus will come over industry competitiveness, AI has the potential to be a game changer in all five forces, between buyers and suppliers as well as a likely accelerator, if not enabler, of new entrants and alternatives. A supplier using AI will certainly gain advantage over a competitor without, and potentially a disruptive advantage, but that supplier may also gain a competitive advantage over the buyer. And the complementary situation is equally possible.

Consider the AI competitive "stack" today. At the base of the AI layer cake are the hardware manufacturers building the key chips required for the matrix algebra at the heart of neural networks. The leading provider of this hardware is Nvidia, which has rocketed into sixth place in the most valuable companies list, with a market cap exceeding $1 trillion, outpacing TSMC, Broadcom, Intel and AMD *combined*. A layer up are the companies doing the fundamental research and development in AI software, a list which includes Alphabet, Amazon, Apple, Microsoft, and Meta, as well as OpenAI, Anthropic and Hugging Face. On top of this layer are the companies building AI applications in sectors and industries, selling to the top layer of firms and consumers implementing AI in their businesses. This third layer includes all the major (and the surviving minor) enterprise software companies as well as a bewildering host of startups. The top layer is the consumers of this technology, mostly businesses at this point but rapidly expanding to individual consumers. This group includes almost every business with significant capabilities in technology, customer service or research and development, and most CEOs not working on their AI strategy are worrying about it.

But how clean is even this stack? Despite, or perhaps because of, Nvidia's commanding 85% share in the enabling hardware, Alphabet, Amazon, Apple, Meta and Microsoft have all announced meaningful forays into AI hardware. Nvidia in turn has announced fundamental AI software technology and applications. Even strong partnerships, such as the relationship between OpenAI and Microsoft are giving new meaning to "coopetition". Though the AI players themselves are perhaps at the bleeding edge of such proactive business redefinition the rest of industry will likely follow. The ability to use technology to understand and serve the needs of a the end customers will likely change the structure of classic supply chains surpassing the fiasco of the supply chain disruption of the Covid pandemic.

Supplier and Buyer Power

The fundamental balance between buyers and suppliers is likely to change with the introduction of AI. Buyers will gain unprecedented access both to information and sourcing alternatives. Suppliers will have

new opportunities to forward integrate into their customer's supply chains. Complex value chains will likely reorient as both buyers and suppliers use the new technology for disintermediation, and we will see the rise of new intermediaries who provide that capability as the value add in the value chain begins to evolve. Dr. Porter detailed how much of this change might occur in his 1985 *Harvard Business Review* article "How Information Gives You Competitive Advantage". While information technology provided the basis for this strategic realignment, it is AI which makes it an imperative. IT tended to scale with existing players, AI can disrupt that scaling.

Buyers will be able to use AI to do more sophisticated searches and analyses of supply alternatives. A sourcing event that might take a few days today could easily be automated using AI to ease cross system integration resulting in a new process that takes seconds. Instead of phone calls and emails understanding product performance, supply availability, and delivered cost, an AI could perform the process effectively instantly changing both the sourcing relationship as well as the breakeven point. We already see products which monitor web purchasing and suggest alternative sites, AI could take this to the next level. While this is clear in a consumer setting, consider a contractor pricing a complex bill of materials. Today, he is likely to put the full list out for a few bids, possibly excluding only a few expensive items which might be bid on their own, and likely to offer the bid to only a few known suppliers. An AI driven approach could make the event much more competitive, balancing costs and availability and offering the package to any potential supplier committing to a contractual schedule. To do that manually might take a small team driving fulfillment costs unnecessarily high, but an AI could do the optimal sourcing as easily, quickly, and efficiently as a human finding a single item.

Suppliers have equally exciting opportunities. There will be corresponding opportunities to integrate their supply chains forward into those of their customers with a much deeper understanding of customer needs. This may provide an opportunity not only for better relationships, but for expansion into related adjacencies, increasing the scope of the relationship as well. Regardless of which side of the relationship the AI progress first emerges, it will result in deep and lasting changes.

As companies think through their AI strategies it will become imperative to focus on these new perspectives and the opportunities that they create. Simple cost reduction efforts will provide a good platform for getting started, but they will quickly suboptimize value creation. If all of your AI projects are cost reduction your team is probably undershooting the mark.

New Entrants and Alternatives

Whether these changes are seen as disruptive or evolutionary will depend on how the incumbent leaders react. The dual risks for the incumbent leaders will appear to be going too slowly so as to be overtaken, and going too quickly taking on unnecessary costs, diverting better spent resources, and adjusting pricing too radically. We believe that the risks of waiting are likely to be far greater than the risks of going too quickly, except on the pricing front where competitive pressure should determine strategy on a case by case basis.

Clearly, the risk from AI will differ by industry. The chart at right depicts one view of the likely impact of AI by industry. The

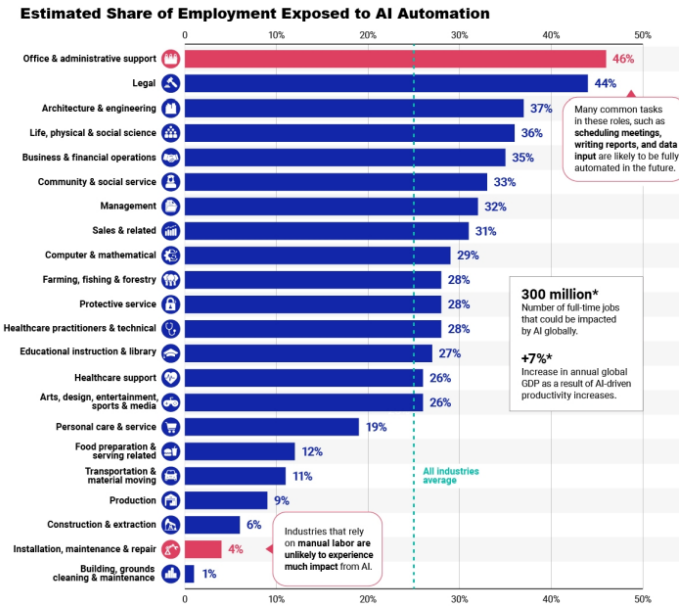

chart shows a range of impact from 46% of jobs exposed to AI automation in office and administrative support to 1% in building, grounds, cleaning & maintenance. While we would disagree with several of the assessments (the analysis ignores the impact of AI on robotics), the more critical issue is that most industries are comprised of multiple value chains each comprised of multiple complex steps with potential AI impact on each step and value chain, as well as the potential for augmentation or full replacement. It is highly unlikely to find a company in the low impact categories without job functions that will be affected, and the less critical these functions in that industry the more heavily impacted. There are office workers in janitorial services; these jobs will not be spared due to their existence in an otherwise low AI-impact industry. In fact, they may be the first to go.

We are currently working with a construction management company who is investing in AI as it offers both the ability to reduce costs in numerous critical functions, but also offers the ability to remake traditional relationships in the industry. The ability to share AI driven estimation and scheduling capability directly with subcontractors looks to remake relationships which have previously been driven by information imbalance.

We believe that this kind of change will become commonplace, especially as we reach later stages of AI adoption. And it is likely that the leaders in the later stages will come from the ranks of the early adopters in the earlier stages. Unlike other technology introductions, such as enterprise ERP, AI offers the ability to rethink value add and remake value chains, not just automate functions and this will lead to much more significant change.

The deeper need and product/service matching capabilities will be provided – either by companies directly in the supply chain or by new entrants offering those capabilities either on a standalone basis or as part of a fuller product space. The ability to meet this new capability demand will, in numerous industries, replace more traditional sources of competitive advantage and value. Whether entrenched leaders lead this change or are disrupted by it remains to be seen.

A *Full Potential Paradigm™* view

Applying a *Full Potential Paradigm™* analysis to AI implementation confirms the disruptive impact of the technology on industries. We believe that unless AI adoption proceeds at an even pace across an industry with all competitors progressing at the same pace, which is extremely unlikely, the technology will effectively split industries in half. Competitors aggressively pursuing AI based strategies, including new entrants and alternatives based on the technology, will have significantly lower operational costs and faster growth resulting in higher PE multiples. These companies will also follow an adoption curve-like path for revenue, while the entrenched competitors will follow the post plateau decline typical in an industry on a substitution curve.

A *Full Potential Paradigm™* view of competitive structure shows that despite relative market share continuing to be the key long-term driver of costs and profitability, AI implementation will likely have an extreme short-term impact on both. This has been all but assured by OpenAI and other competitors' aggressive stance on pricing, making AI training less expensive than comparable human training for most businesses, and AI execution faster, cheaper, and more reliable. We worked with a well-known financial services firm who found that fully training a customer service representative, along side day-to-day experience, took an average of twenty-six months to complete, while the expected tenure of a new CSR was under two years. While we believe that while top attorneys will continue to be highly valued, AI can already produce and review most contracts less sophisticated than a merger and acquisition agreement less expensively and more effectively than most in-house counsel, and do so both immediately and at

negligible cost. These short-term cost advantages will likely lead to aggressive price competition rather than inflated margins as price competition can lead to meaningful share gain.

Beyond margin and revenue, we expect the market to value firms with strong AI delivery at a preferential price earnings multiple. The higher margins enabled by AI times the higher multiples will yield a powerful tool for new leaders to purchase existing competitors which have not made the transition. The resultant entities will be far more likely to commit to the newly enabled, more efficient, lower labor business models. Consider what the impact of a Tesla with a PE ratio in the 70× range *vs* Ford and GM in the 5-6× range if Tesla can extend their labor per vehicle advantage from $1,000 per vehicle to double that with AI and robotics on their side and union labor on the other.

Routine financial analysis will need to be strategically updated, as there will be multiple reasons for potential precipitous changes in profitability. Some firms will follow the example of Amazon's distribution business and OpenAI's pricing model, eschewing short-term profitability for long-term improvement in relative market share (RMS). RMS, which looks at competitive scale rather than simple market share, is a central tenet of the *Full Potential Paradigm™*. Others will find profitability failing as they fight against new efficient business models with rapidly antiquating manual processes. And still others will show declining profitability as they write off those older processes and move to the newly enabled business models. It's only called "disruption" when it's done to the leaders, it's "evolution" when they do it to themselves. It will take strong leadership to survive, let alone direct, the transition, and we believe few leaders will rise sufficiently to the challenge.

AI and the Business of Strategy

In the short run we will see substantial changes in the way that business strategy is undertaken. In some ways these changes will seem incremental as those changes have already permeated tactical decision making, but AI will bring these changes to strategy itself.

Although we expect the characteristics of leadership and vision will continue to define the CEO, the strategy support for those CEOs will change radically in two interrelated ways:

1. Strategic Cycle Time. John Boyd, the father of the "OODA Loop" (observe, orient, decide, and act) and "fifth generation warfare" described the critical military strategic task as "getting inside your opponent's decision loop". The winner in air-to-air combat would not be the faster plane, but rather the pilot with the ability to see and think faster than his opponent. The open cockpit designs of American fighters proved to be more important than the better flight performance characteristics of their Russian made counterparts. Faster AI-enabled execution will fuel this transition for business as well as military strategy.

   Gone will be the days of the "annual strategic planning" exercise. A twelve-month-old strategy will seem aged and irrelevant. Vision will persist, but the strategies to implement those visions will require agile processes that react to competitive changes and enable strategic, as well as tactical response. We expect strategic planning will evolve with shorter and shorter cycles and finally to on demand in response to context. Today's long PowerPoint presentations will become as anachronistic as the inch-thick acetate transparency decks of the 1960s and 1970s.

2. Supporting the faster cycle times will be a move from qualitative strategy to quantitative evaluation. While stories of past successes will not disappear from strategy nor consulting, they will become supported by translation into business models and dynamic simulations such as war

gaming and Monte Carlo analysis. Plans will give way to scenario analysis, fixed targets to responsive models, and single point solutions to more holistic outcomes across a range of measures.

In concert with the above collapse of strategic planning cycles, the required processes and tools required for effective, on-demand, planning will evolve as well. That replacement will need to be as agile as the process that it supports, and that will require complex quantitative models for all but the simplest of businesses.

Looking at single measures, even one as inherently comprehensive as market value, suffers from not being actionable. Point measures undervalue the impact of tradeoffs. What will be necessary is a dashboard of critical measures across constituents and factors, and one which encompasses competitive performance and trends as well as a company score. But these models need be interconnected, and the interconnection needs to be quantitative and explicit. We believe that our Acropolis Advisors stakeholder dashboard will prove one such viable starting point.

At its best, AI is able to distill reams of data and convert it into positive action without human intervention. Initially, this will occur for basic tasks followed by more advanced tasks. However, because there is rarely a "perfect" answer for mission critical decisions, human intervention will always be necessary. AI makes decisions based on stochastic distributions of prior results allowing a considered portion of randomness for creativity. Consequently, AI supports better informed decision making based on probabilistic outcomes but will never be able to replace judgment amongst competing interests. AI will also never replace groundbreaking creativity. The DaVinci's and Einstein's of the world will continue to be the trailblazers doing more than applying randomness.

When companies construct a strategy, they scour the landscape for intelligence on **customer** behavior, needs, and pricing sensitivities; **competitor** capabilities, costs, and value propositions; **supplier** capabilities, costs; and the **societal** requirements for the business. They then compare that intelligence to their own capabilities and costs. This baseline of information will form the basis of developing a sound strategy that can achieve competitive advantage. If the company is a follower in the industry, the strategy will be geared to gain share through cost leadership or differentiation. If the company is a leader, it will be geared to bolster its leadership position by combining both.

Since the AI era is now in full gear and is becoming ubiquitous, it's important to realize that it will impact all of the above inputs to strategy development in some measure. It will require having a "real time" monitoring capability addressing all elements of each aspect of the business. Importantly, companies comprised of people cannot be constantly reoriented. Such continuous turmoil could only result in corporate ruin. When the evidence mounts that the conditions upon which the strategy was developed have been sufficiently altered, the strategy needs to be revisited. Strategy must be fashioned to evolve to consider multiple scenarios becoming more nimble and agile.

Strategic Case Example: Auto OEMs *vs* Unions

Take the current situation that exists between the "Big Three" automakers in the US and the UAW. If you are one of the OEM CEOs, you must be able to process a myriad of factors as well as how those factors will unfold over time. It is a probabilistic linear programming problem that cannot be solved without an objective assessment of the probabilities of the inputs and intermediate steps. If the OEMs were to accept the UAW proposals how would it impact profits, market share, employee satisfaction,

supplier economics, product pricing, customer receptivity, community relations, relative cost position, and on and on? The complement is true for the UAW. If they accept the OEMs proposals, how will that impact member compensation, membership recruitment and retention, dues, benefits, pensions, *etc*.... These analyses must be done in both the short and long run, with respect to their other unionized and non-unionized North American competitors and to global competitors. They must also calculate various scenarios both for EV adoption and the introduction of AI based robotics. Significantly, Chinese EV manufacturer Nio has recently announced plans for a motor manufacturing plant that would produce 400,000 motors annually with a staff of 35. Contrast to the US with 10 million engines made by approximately 60,000 workers. Nio represents a 98.5% reduction in the number of workers, union or otherwise, per vehicle. An "optimal" solution for the UAW that forces the OEMs to follow Nio's lead will decimate or entirely eliminate union membership.

Is there a Pareto optimal solution for both parties? Today, the parties present their proposals and counter proposals after they go back to their corners and painstakingly try to calibrate the outcomes. While better than the old days before computers, it pales next to having the ability to tap into an exquisitely designed AI program which can immediately update the probabilities and outcomes. Even better, a "neutral" and shared AI could present both parties with the logical steps over time in the negotiation and allow them to mutually select and communicate a shared best outcome.

Conclusion

We expect the business strategy firms and internal strategy functions to evolve to support this new AI driven model. Just as with other services industries, the new AI strategy team model will be a small cadre of highly experienced senior practitioners with a small group of talented juniors and analysts with data science, math, and programming skills in support. What takes a team of a dozen three months to do today will become a continuously updated exercise with results available on demand. A CEO considering a new product strategy will expect a model capable of reacting to his vision and direction on an agile basis. This will require an unprecedented synthesis of strategic knowledge with computational ability.

For most executives the AI decision for today comes down to whether, how, and how much to invest.

We think the easiest part of that decision is the "whether". A decision not to invest appropriately is tantamount to a decision to harvest a business. Some competitors in the industry, whether current competitors or new entrants, will be making the investment and will gain an advantage through it. There will be some time to correct for a bad AI strategy, but there may not be enough time to effectively correct for a lack of one; the internal learning will become more and more consequential. Fully relying on external parties, such as vendors and consultants, for parity may seem an attractive option, but will likely relegate competitors to "me too" follower status. In industries where AI is less material this may be acceptable in the short run, but those instances will be rare.

The practical aspects of "how" are also likely to be straightforward. Decisions as to what technology or even type of AI technology are likely to be fluid; the environment is so dynamic that there are no durable right answers, and even otherwise wrong answers are likely to be instructive. Understanding what doesn't work and why will be advantageous, though the usual dictum of "failing fast" will be even more important than usual. The decision whether to use internal or external resources to lead the effort is likely moot; the right answer is going to be both. The number of companies who have the internal talent and experience required is currently negligible and we are yet to see the emergence of stable leaders in application software or consulting; given that practical generative AI is less than a year old, it's

hard to imagine anyone credibly claiming years of experience. We have dozens of years of both computational linguistics and behavioral game theory experience, as we are still feeling our way through this.

We believe that the immediate requirement is a thorough strategic assessment and action plan. The first stage of that assessment is a frank discovery of the potential impact of AI on the company's business model. Every company needs to understand in detail the key components of the business model and how and when these may change with the introduction of AI technology, from cost reduction to new value creation, following both traditional and new alignment models. It's not going to be possible to be accurate here, but the goal must be to identify concrete short-term opportunities, risks, and potential long-term ones.

In concert with this process, we believe that most companies would be well served to prepare what we call a "Differentiation / Impact Matrix", pictured at right.

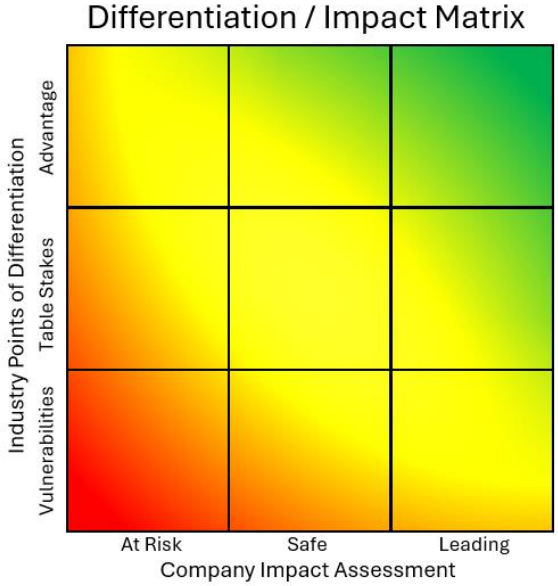

In this matrix, which should be continuously updated as AI impact on your industry becomes clearer, we plot industry differentiation on the vertical axis and company performance on the horizontal much as we look at other sources of growth and risk. We would place each item identified in the business model assessment onto this matrix, with an emphasis on being "maniacally realistic" – frank and data driven – about opportunities and threats. If you find most of the business model items showing up in the green, you are probably being overly optimistic; given public statements none of Google, Microsoft or X (formerly Twitter) would place themselves with much in the red.

A quick note on the classification, as it's perhaps not immediately intuitive. Clearly an industry trend providing competitive advantage where your company has a leadership position is a good thing as industry vulnerability in an area where you are weak and will likely be an early target is a problem. But safety is mostly a middle of the matrix proposition. Being weak at an activity that is likely to provide a strategic advantage for the industry can be as damaging as weakness in an area of vulnerability as the basis for competition and competitive advantage may shift away from you.

Our advice to CEOs follows a uniform path:

1. Create a dynamic model to understand the impact of AI on the industry's business model based on Porter's *Five Forces* and *Value Chain* frameworks.
2. Extend that model to quantitatively describe the value creation opportunity as depicted through Achtmeyer's *Full Potential Paradigm™* framework, delivering revenue, margin, and market value projections.
3. Understand the impact of AI on the company's value proposition to each of the company's stakeholders over time and on the above frameworks.

4. Identify the short- and long-term actions necessary to offset vulnerabilities and enhance strengths in the company's value proposition and model the cost and value of these strategies and actions.
5. Make sure that the model presents the implications of these actions on customers, employees, suppliers, local communities, and society in the relevant frameworks (*e.g.*, standard financial statements, per employee impact statements, the UN Social Development Goals, GHG Protocol).
6. Develop and deploy ongoing measurement and monitoring tools based on the Acropolis Advisors *21st Century Stakeholder Optimization Framework and Dashboard*.

The "how much" resources to invest will be situationally dependent. For most companies this is a true existential risk and needs to be treated as such. Investment should be prioritized as much as would the entrant of a new, well-funded competitor with a competitively attractive product or service. AI will only serve to make the already complex world of the CEO even more so.

Onwards!