

Prologue

How is analytics unique in a corporate context? What have other organizations done right? What have they done wrong? What are the expectations for a new analytics leader?

Building, integrating, and leading effective analytics teams is a business imperative. The organizations that are most successful overall are those that effectively leverage their analytics capabilities to build a sustainable competitive advantage. However, many organizations are simply not getting the return that they expected on their investments in analytics.

Does hiring an engineer cause the surrounding buildings to be more robust? Could hiring five engineers make those buildings even more robust? Would hiring a pharmacist make you healthier? Would hiring an actuary increase your longevity?

In the last 20 years, the once sleepy academic fields of statistics, operations research, and decision support have exploded and been rebranded using terms such as *data science*, *artificial intelligence*, *big data*, and *advanced analytics*, among others. These practices have matured, and they have entered the mainstream. Advanced analytics and AI have moved from

being an investment in the future to a core component of corporate strategy and a key enabler for all areas of the business. Regardless of their size or the industry or sector in which they are used, these technologies are becoming an organizational necessity throughout the world. Organizations use advanced analytics and AI to develop new products, understand their customers, control their costs, and make better-informed decisions.

This new corporate function has been integrated in several support and core functions and has quickly become indispensable. Analytics is expected to add \$16 trillion US to the global economy by 2030 and companies are eager to realize some of that value (PwC, 2017). As a result there has been a surge in demand for practitioners. There are approximately 3.3 million people employed in analytics in North America, and this is projected to grow by 15 percent a year in the United States over the next decade according to the US Bureau of Labor Statistics (2020). Educational institutions are eager to meet this demand.

Essentially every major college and university in the world offers some sort of analytics program or specialization, within multiple faculties such as mathematics, engineering, or business. There are several hundred books published in this space, and it enjoys a highly active online community. These resources are strong, edifying, and comprehensive and cover every new technology, framework, algorithm, and approach. With such an active community, new algorithms and methodologies are packaged and made publicly accessible for tools such as R, Python, and Julia, almost immediately after being developed. The best and brightest are choosing to enter the field, often called the “Sexiest Job of the 21st Century” (Davenport & Patil, 2012).

So, with overwhelming demand and a staggeringly capable pool of talent, why are there so many failures? Why are most organizations struggling to unlock the value in data science and advanced analytics? With so much executive support, so much talent, so much academic focus, why are so few organizations successfully deploying and leveraging analytics? In the 1980s, economist Robert Solow remarked that “you can see the computer age everywhere except in the productivity statistics.” Why now can we see data science transforming organizations without a commensurate improvement in productivity?

The fundamental reason is in the opening questions. Clearly, having five engineers on your payroll will not improve buildings in the vicinity. Even if directly requested, ordered, mandated, or incited to “improve a building,” they will struggle to do so. Replaced with any other profession, it is clear why this cannot work, but it persists as the typical furtive first steps into analytics. This is the essence of what most organizations

have done with their data analytics team. They have hired talented, passionate, ambitious data scientists and asked them to simply “do data science.” Advanced analytics and AI as a practice has been poorly defined in its scope, has been subjected to great overspending, and has been lacking in relevant performance measures. Because of this, it is failing to live up to its potential.

Effectively all organizations realize the benefits of analytics. In a survey by Deloitte in 2020, 43 percent believed their organization would be transformed by analytics within the next 1 to 3 years, and 23 percent within the next year (Ammanath, Jarvis, & Hupfer, 2020). Though most organizations are on board with analytics being a key strategic advantage, they are unaware of how exactly to extract value from the new function.

Short-tenured data scientists, employed in a frothy and competitive market, share stories of unfocused and baffled companies where they have been engaged in operational reporting, confirming executive assumptions, and adding visualizations to legacy reports. Uncertain what to do with the team, and in a final act of surrender, the companies no longer expect the function to “do data science” and transform the team into a disbanded group of de facto technical resources automating onerous spreadsheets in a quasi-IT role.

Contrasted with those organizations who have truly got it right, the differences are stunning. For several companies, well-supported analytical Centers of Excellence are a key team, perpetually hiring and growing, and are solicited for their input and perspectives on all major projects. In others, internal Communities of Practice encourage cross-pollination of ideas and development opportunities for junior data scientists. New products are formed and informed after a thorough analysis by data scientists, who are also supporting human resources with success indicators and spending Fridays pursuing their transformative passion projects. Theories and hypotheses are quickly tested in a cross-functional analytical sandbox. Individuals are sharing their work at conferences and symposia, building eminence, and gaining acclaim for the organization.

The irony is that while most analytics teams exercise great care in deconstructing a problem, modeling each of the elements, and developing robust simulations and marvelously elegant solutions, the business processes that they employ in their project delivery are often at best immature and at worst destructive. The challenge for any leader is to encourage and stimulate logical thinking from a lens of interconnectedness and value creation, and to direct that philosophy to execution as well. As the logician Bertrand Russell said, “What is best in mathematics deserves not merely to be learnt as a task, but to be assimilated as a part

of daily thought, and brought again and again before the mind with ever-renewed encouragement” (Russell, 1902).

The title of this book, *Minding the Machines*, is meant to be an affectionate recursion. Those talented and creative practitioners, craftspeople, data scientists, and machine learning engineers, who create the algorithms that are transforming the way business is done, mind and care for those machines like a shepherd. Those machines need to be trained, informed, given established processes, encouraged to be broadly interoperable, and developed to be applicable to many different situations and problems. Similarly, the practitioners themselves need to be minded, cared for, cultivated, and encouraged for both the team and the individuals to be successful. This concept of recursion occurs throughout the book.

The second key theme is one of parsimony. Parsimony is a philosophy of intentionally expending the minimum amount of energy required in an activity so as to maintain overall efficiency. This is a key part of modeling; it is about keeping things as simple as possible, but no simpler. Similarly, and in the vein of recursion, teams themselves must be parsimonious. For analytics to mature as a practice while still delivering an accelerated time-to-value, teams need to be scrappy and lean.

At a foundational level, the objective of this book is to provide clear insights into how to structure and lead a successful analytics team. This is a deceptively challenging objective since there are no generalized templates from which to work. Establishing a project management office, information services, or human resources department is an understood process and does not vary materially between organizations. Establishing an analytics team, by contrast, requires a significant up-front investment in understanding and contextualizing the initiative. Many organizations have attempted to use operating models and templates from other functions—often IT and operations research. This fundamental misunderstanding of where analytics fits within an organization has led to visible failures and has set back the analytical maturity of many organizations. Business leaders need to hire or develop value-centric talent who can step back from analysis and project management to view their work as existing within a network of individuals and teams with competing priorities and motivations.

Corporations, without a template or a default methodology to benchmark against, have made expensive missteps in building these teams by installing the wrong leaders, copying other functions, positioning the function under the wrong executives, incentivizing destructive behaviors, hiring the wrong people, and committing to the wrong projects. They have contracted expensive consultants to provide roadmaps that do not consider the unique culture or competencies of their organization.

They have embedded disparate data science specialists throughout the organization and incurred enormous technical debt.

These issues are not insurmountable, however. Whether an organization is beginning its first foray into the analytics space or it is rebooting a failed team for the third time, the key is the creation of a carefully considered strategy, the establishment of realistic goals, and the full commitment of executive leadership. The advantages associated with analytics are too great to overlook, and the long-term cumulative impact of interrelated and interdependent models provides a powerful incentive for aggressive adoption.

This book was written for anybody who aspires to lead or be part of an effective analytics team, regardless of managerial experience. Every analytics leader, from a first-time team lead to a seasoned VP, has unique challenges to overcome.

For the Leader from the Business

Every new role is a challenge, regardless of ability, disposition, or motivation. This is particularly the case with a unique subculture of academic technocrats with whom it is difficult to establish credibility without enough time being “hands on keyboard.” Without the respect of your team, it is impossible to get the buy-in required to establish best practices and ensure that the output of the team is not simply self-satisfying experimentation but can bring real value to the organization. As a corollary, every practitioner has experienced a manager who is out of their depth and who has compensated for their lack of self-confidence with authoritarianism and distrust, shifting the focus of the team toward an end they are more confident with, such as reporting.

For all but the most analytically committed organizations, there is a point along the chain of command where a practitioner reports to a non-practitioner. This can be a challenging junction for both parties without clear expectations, transparent communication, and mutual respect. Catching up from a technical perspective isn’t feasible or advisable, but by leveraging your business understanding and domain knowledge to become an intermediary, translating business needs into projects and analytical outputs into operationalizable processes, you can unlock the power of your new team and give them the opportunity to develop into more business-oriented individuals.

For the business leader, I hope that this book helps you to reframe and refine your current leadership abilities, and to use them in an analytics context in order to engage your team and find success together.

For the Career Transitioner

Those who transition to data and analytics mid-career have a key differentiator from those who have entered the field directly from university—breadth and context. The ability to leverage your multidisciplinary background from engineering, finance, sciences, and so on is valuable both to your career and the organization you join.

Though it would be almost impossible to compete with trained data scientists on a technical basis, it is the disciplinary and sector diversity of the team that drives innovation, and those who have worked in multiple industries and functional areas bring a unique perspective to the teams with whom they work. Rather than starting a new career as a new hire and individual contributor, with personal study and intentional self-reflection mid-career transitioners are often able to seamlessly make a lateral move. Having familiarity with the different AutoML and analytics-as-a-service offerings, combined with transferable managerial skills, can make for a powerful combination.

For the career transitioner, I hope that this book helps you to prepare for lateral movement into an analytics role and to use your transferable skills to add value in your new function.

For the Motivated Practitioner

It is an unfortunate truth (and perhaps an unfair generalization) that the skillset that makes a practitioner a competent data scientist is rarely the skillset that makes them a competent manager. Though there are certainly analytically minded people who have the natural inclinations toward leadership and bigger picture thinking, it is rare that in practice those people would have the technical depth to stand out as a candidate for management. Often, those with the natural capabilities required to enter management can appear to be less effective as individual contributors on a purely technical basis.

To make the leap to management is to leave an objective and predictable role with performance metrics such as p-values and ROC curves and exchange them for stakeholder management, workshop facilitation, and inherent subjectivity. Those able to successfully make this transition while maintaining the ability to downshift to provide analytical support establish themselves as leaders in the practice and are in high demand. Exceptional managers who have legitimate technical credentials are the unicorns of data and analytics.

For the practitioner, I hope this book helps you to understand what is required to move up the value chain and to prepare for leadership opportunities.

For the Student

When a student pursues an applied field such as business or engineering, the curriculum is generally developed in a way that seeks to balance between foundational academic elements and applied profession-specific education. The curriculum is updated and maintained such that it remains aligned with the changing needs of the field. For several professions such as accounting, law, and engineering, this takes place within a partnership between the administrative body of the professional practice and the educational institute, and through accreditation it's ensured that graduates of these programs are broadly educated and prepared to work in the field they have studied. This is unfortunately not the case with data and analytics.

Most North American universities have data science or analytics offerings, but having no natural home they are generally provided through multiple faculties such as business, mathematics, engineering, finance, or computer science. These programs provide instruction in highly simulated and well-defined problem solving, focusing on the improvement of a statistical metric. The data is often perfectly presented and accompanied by a well-articulated data dictionary, in great contrast to real life experience. Additionally, most curricula emphasize such topics as computer vision, natural language processing, and reinforcement learning, fairly esoteric topics that have little applied usage in industry. Finally, and most importantly, effectively none have mandatory coursework on the strategic and operational elements of an advanced analytics and AI team. Without this understanding, typical graduates have a thorough mathematic understanding, much in the way of raw horsepower, but require a significant investment in training before they understand how to leverage their education and apply it to a real-life scenario.

With so many new data and analytics graduates competing with mid-career transitioners and a global talent pool, they often seek ways to stand out as a potential hire. With the exception of highly specialized roles in technology companies, the key development opportunity for these new hires is the formation of leadership abilities in an analytical context. Reframing and focusing analytical concepts into a business context is

an immediate and powerful way to differentiate yourself in a new role or in an interview, especially as the profession moves away from long-horizon highly technical solutions toward a focus on immediate value.

For the student, I hope that this book gives you the knowledge to stand out against your peers, to be seen as a strategic thinker, and to be able to add value to whatever organization you choose to work with.

For the Analytics Leader

Compared to other organizational functions, this exciting field has come about abruptly and without a blueprint for how to build or lead these new teams. Often playbooks from other functions have been used, with little success. The lessons that experienced analytics professionals have learned have been hard won. What further complicates the successful deployment of these teams is that they are so sensitive to the state of the organization, its immediate goals, and the technical maturity of its industry. While finance and human resources are largely the same between industries and individual companies, the number of factors impacting analytics are staggering.

Using a generalized, systematic, and sequential approach, adapted to the needs of the individual organization, is the best method to standing up a new team or restructuring an existing team. Once a base template has been established, careful reflection on organizational readiness and analytical maturity combined with regulatory requirements and immediate needs can help with developing a short-term roadmap in collaboration with executive sponsors. Though there is no approach that will work in every situation, these best practices can hopefully help you see a little further over the horizon.

For the current analytics leader, I hope that some parts of this book will challenge your views, other parts will confirm your experience, and the book as a whole will ultimately help you to build out a successful and engaged team.

Structure of This Book

The main body of this book has been organized within three key pillars: strategy, process, and people.

Strategy How to assess organizational readiness, identify gaps, establish an attainable roadmap, engage stakeholders, ensure sponsorship, and properly articulate a value proposition and case for change

Process How to select and manage projects across their life cycle, including design thinking, risk assessment, governance, and operationalization

People How to structure and engage a team, establish productive and parsimonious conventions, and lead a distinct practice with unique requirements

These pillars loosely follow the chronological and logical ordering of priorities with the creation or inheritance of an analytics team, with the understanding that this is an iterative and ongoing effort. The procedural requirements flow naturally from the strategy, and similarly, team structure and convention must be based on the processes that have been created.

Though this has been ordered to facilitate a front-to-back reading, subsections have been intentionally made self-sufficient to allow for ease of referencing, at the cost perhaps of occasional repetition.

Why Is This Book Needed?

It is my personal hope that this book will make creating and leading the function easier and help in some small way to advance the profession. Having been involved with or privy to rebooting these teams in several organizations, I have seen well-intentioned missteps repeated regardless of the maturity and sophistication of the company.

There are several underlying reasons why organizations and individuals struggle to get their hands around analytics.

Communication Gap

The business will rarely, if ever, have the analytics knowledge and vernacular required to clearly articulate its needs and to formulate a problem statement that naturally lends itself to an analytical solution. Whereas Kaggle competitions, hackathons, boot camps, and university assignments present problems with a well-formed data set and a clear desired outcome, business problems are fuzzy, poorly defined, and often posited without a known objective. As practitioners, it is our responsibility to find the underlying issue and present the most situationally appropriate and practical solution.

Advanced analytics and AI practitioners can often have the expectation that their stakeholder group will provide a solution for them. Just as a

doctor cannot expect a patient to diagnose their own health issues and for the doctor's approval an analytics team cannot expect a business unit to suggest an approach, provide a well-formed data set and an objective function, and request a model. What the business unit requests is very often not even what the analytics project lead hears.

Early in the project intake process, an analytics lead will meet with a business lead to discuss an opportunity. The business leader (actuarial, in this example) may say that they want a model that predicts the probability that a policyholder will lapse. The outcome that the leader is hoping for is a way to reduce their lapse rate, but what the analyst hears is, "Ignoring all other considerations, how can I best predict the probability of an individual lapsing?" If the practitioner executes on this misapprehension, the deliverable will have little use for the business; a prediction model of this sort has no operational value. This model would only work on a macro scale, and even if it could be disaggregated, the business would be making expensive concessions in the face of perceived threats.

Empathizing with the underlying needs of the business, understanding what success looks like for the project, and leveraging the domain knowledge of the project sponsor would have highlighted that the value in the analysis was further upstream. The factors driving lapse behavior were where the value to the business was and where an operationalizable change in process was possible.

As with the doctor analogy, it is through deep questioning, structured thinking, and the expert application of professional experience that the ideal path forward is uncovered. That path requires collaboration and the union of deep domain knowledge with analytical expertise.

Troubles with Taylorism

For every decision to be made there is a perception that there must be one optimal choice: a single price point that will maximize profit, a single model that will best predict lapse, or a single classification algorithm that will identify opportunities for upselling. The fact is that in effectively all cases these optima can never be known with certainty and can only be assessed *ex post facto* against true data. In professional practice as well as in university training, the results of a modeling project are typically evaluated against real-world data, giving a concrete measure of performance, whether AUC, or R squared, or another statistical metric.

This has created a professional environment where analysts can confidently point to a single score and have an objective measure of their

performance. They can point with satisfaction to this measurement as an indicator of their success and evidence of the value they bring to the organization. Certainly, performant algorithms are an expectation, but without viewing the work through a lens of true accretive value creation, these statistical metrics are meaningless.

In the 1920s the practice of scientific management led to improvements in the productivity of teams by breaking the process into elements that could be optimized. Through a thorough motion study of workers at Bethlehem Steel, Frederick Taylor created and instituted a process that optimized rest patterns for workers and as a result doubled their productive output (Taylor, 1911). He advocated for all workplace processes to be evaluated in terms of their efficiency and all choice to be removed from the worker. This brutal division of labor and resulting hyperspecialization led to reduced engagement and produced suboptimal outcomes at scale when all factors were considered.

Practitioners need to avoid those actions and policies that create a form of neo-Taylorism within their organizations. Models that fully automate a process and embed simulated human decision making remove the dynamism and innovation that comes from having humans in the loop. It cements a process in place and reduces engagement and stakeholder buy-in. Analytics should support and supplement human endeavor, not supplant it with cold efficiency. It is essential that analytical projects are done within the context of the business and with the goal of maximizing the value to the organization.

Model accuracy needs to be secondary to bigger-picture considerations, including these:

Technical Implementation Is the architecture stable? Does it require intervention?

Political Implementation Does it conflict with other projects? Will implementation create redundancies?

Procedural Implementation Will this fit in with existing processes? Will it require significant changes to current workflows? What are the risks associated with implementation? Will it introduce the potential for human error? Does it have dependencies on processes that are being sunset?

Interoperability Are there downstream processes depending on the results? What are the impacts of a disruption to these processes? Can it be shifted to another system? Does it create dependencies?

Extensibility Can the output be upgraded in the future? Does it require specialized skillsets? Is it generalized enough to be used for other purposes?

Scalability Would this approach work if the volume of data doubled? Tripled?

Stability Has it gone through thorough QA? Has it been tested at the boundary conditions? What happens if data are missing? What happens if it encounters unexpected inputs? How does it handle exceptions?

Interpretability Are the results clearly understandable? Does the process need to be transparent?

Ethics Is it legal? Does it have inherent bias?

Compliance Does it contain personally identifiable information? Does it comply with the laws of the countries in which it will be used? Does it use trusted cloud providers?

Without exception, effort is better spent in discussing and addressing the above considerations than in marginal improvements to model performance. Even a poorly designed model will work with strong phenomena, and a poorly performing model that is in use will outperform a sophisticated model that is sitting idle.

Rinse, Report, Repeat

One of the most memorable scenes from the 1990s classic *Office Space* involves the protagonist being chastised repeatedly by his coworkers and managers about an incorrectly prepared TPS report. The reason this was so memorable is that it reflects the experience of every office worker at some point in their careers. Bureaucracy is a self-propagating danger, and for many larger legacy organizations, entire teams are dedicated to the preparation and dissemination of unread reports. There are several compounding factors that have led to an explosion in the number of reports recently.

With legacy reporting tools, the effort required to create reports created a natural limit on the number that could be produced. This provided a control on the spread of new reports. Unfortunately, the strength and ease of use of reporting and visualization tools that are available today make creating new reports trivial.

In addition, as tenure at the management level falls over time, new managers arrive and request views of the data that are most conducive to their way of thinking. This, combined with a bias to action, can lead

to a flurry of new reports, which are then promptly abandoned at the next reorganization but continue to be maintained.

Finally, a common rationalization for poor decisions by leaders is that they are due to a lack of information. In response to a personal error, and as a cover to their rascalities, it is not unusual for a new report to be requested, ostensibly to prevent the incident from recurring, but almost certainly as presentable evidence of their proactivity.

Despite all of these drivers of new reports, the organization could conceivably reach an isotonic point of stability where new reports are balanced by discontinued reports. Unfortunately, the process to sunset a report in most organizations is more onerous than the process to create a new one, leaving zombie reports that are diligently maintained by an army of analysts but unused.

For decision makers who have a more tenuous understanding of the value offering of analytics, and a desire to have better-informed decision making, the one-off choice to request that the team create a report can quickly become a pattern of distracting asks. At the expense of large transformative projects, more reports are created and maintained. Over time, the value of the analytics team is questioned.

How can this be prevented? Clearly articulating the value proposition of the team and ensuring from the start that it has executive support is essential to standing up a successful team. Respectfully pushing back against reporting and other non-value-add requests, though politically sensitive, must be done. Every member of the team must be empowered and given the autonomy to question the work that is being requested.

Too Fast, Too Slow

Understanding the first-mover advantage associated with data and analytics, many organizations of means yet with limited analytical maturity have hired dozens of practitioners to rapidly scale up and advance their data capabilities. After 18 to 24 months, these large technocratic organizations review the cost/benefit of the new \$10M division and question the value of analytics to the enterprise. Moving too quickly, without the cultural or technical infrastructure to support it, can put an end to analytical ambitions before the first hire has been made.

Alternatively, many organizations have moved too cautiously and hired one or two early-career employees and placed them under a line manager with conflicting priorities and limited analytical understanding. The maturity of the data science team in this case often peaks with gentle scripting and automation, and executives are again left questioning the value of analytics.

It can seem that a Goldilocks approach of creating a mid-sized team could work, but the issue in both extremes has not been the number of people involved, it has been with a lack of a defined strategy. To prevent this situation, organizations need to have a clear objective in mind for the team, an understanding of what the field of advanced analytics and AI is, and the support that is required to make it a success. Further, organizations cannot look within to build out this strategy, as they thoroughly lack the capabilities to define this strategy. It is only by consultation with others and thorough benchmarking to industry and practice standards that this strategy can be developed.

More Data, More Problems

All organizations have some capacity for gathering data, and once the quantity of that data reaches a certain size, a data-centric team (whether business intelligence, data services, self-service insights, or general information technology) is naturally formed. The purview of these data teams in many cases has been confused with analytics and data science, to the point that the intersection of the two is very often considered the scope for both. Though there is certainly a strong mutually supportive relationship between these two functions, they exist in different spaces and with different incentives—one being a cost center and one being a profit center. Executives with medium-term ambitions to develop a next-generation analytical function will frequently prime the pump by engaging the data team and expanding their accountabilities.

The outcome of this is usually expensive, avant-garde data solutions such as data lakes, premature cloud migrations, and unnecessary latency reduction strategies. The organization, still in its analytical infancy, is unable to take advantage of these powerful technologies, the cost of which is allocated to the analytics team. Organizations need to learn to walk before they can run, and the burden of these data-centric initiatives, executed far too early, has weakened many emerging analytics teams who have had to justify their inheritance after the fact.

These issues are the result of a lack of strategic focus. Well-intentioned projects initiated from a mature data team, supported by an executive, and imposed on future analytics teams leads invariably to technical debt and a handicapped team. Analytical excellence needs to be the focus and the function to be optimized and not confused with the activities that enable it. Early holistic strategy development that considers the interactions between these different teams is essential.

Summary

These common issues all share similar root causes across the pillars of strategy, process, and people. Without exception, every failed attempt to build out the analytical function could have been prevented with some forethought, forbearance, and expert advice. Every externality in the equation supports the practice—there are new approaches and new technologies and success stories of analytics teams adding value, optimizing processes, automating, and increasing the bottom line to the organizations in which they work. To be successful, however, standing up these analytics teams needs to be a thoughtful and measured approach that leverages best practices and integrates with the organization while having a mind to the future.

Organizations and leaders need to follow three guiding principles to successfully build and lead advanced analytics and AI teams:

Start Early The best time to have begun your analytics strategy was 10 years ago, but the next best time is today. Discuss with your colleagues in other industries, perform benchmarking exercises, and engage consultants where appropriate.

Go Slow As the team is chopping their way through the jungle of legacy processes, vestigial data pipelines, and change management, ensure that they are taking the time to pave the way behind them by documenting, securing, understanding, automating, and training others in what they are doing. Moving between projects in a frenetic commotion gives the impression of positive activity at the expense of long-term sustainability. Considering the total project life cycle of each project ensures that future people and processes are not entangled in patchy solutions that had only a short-term view.

Commit Fully Assume that the analytics function will not cover its own costs for the first two to three years and evaluate projects over a longer horizon. Integrate analytics into all functions and encourage a cultural change. Furtive deployments of analytics, starting with low-impact functional groups, leads to low-impact results.

It is my sincerest hope that this book helps you in your journey toward achieving analytics excellence in your organization. This practice can do great good in the world—it just needs to be allowed to succeed through planning and organizing and through minding the machines.

References

- Ammanath, B., Jarvis, D., & Hupfer, S. (2020). Thriving in the era of pervasive AI. In *Deloitte's State of AI in the Enterprise*, 3rd Edition. Retrieved from <http://deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html>
- Davenport, T. H., & Patil, D. J. (2012, Oct.). Data scientist: The sexiest job of the 21st century. *Harvard Business Review*. Retrieved from <https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century>
- PwC. (2017). *Sizing the prize: What's the real value of AI for your business and how can you capitalise*. Retrieved from <http://pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf>
- Russell, B. (1902). *Study of mathematics*. Cambridge University Press.
- Taylor, F. W. (1911). *The principles of scientific management*. Harper Bros.
- US Bureau of Labor Statistics. (2020). *Computer and information research scientists: Job outlook*. Retrieved from <http://bls.gov/ooh/computer-and-information-technology/computer-and-information-research-scientists.htm#tab-6>