

Power and Prediction



The Disruptive Economics of
Artificial Intelligence

AJAY
AGRAWAL

JOSHUA
GANS

AVI
GOLDFARB

“This is a book that leaders of all types of organizations should read. It explains the enormous size of the AI opportunity and the challenges in getting there. From banking to manufacturing and from fashion to mining, the impact of AI systems will be ubiquitous, like electricity and the internet.”

—**DOMINIC BARTON**, Chair, Rio Tinto; former Global Managing Partner, McKinsey & Company

“AI may be to the twenty-first century what electricity was to the twentieth. Anyone thinking about our economic future needs to ponder its implications. This is the best book yet that considers what it will mean for all who participate in our economy.”

—**LAWRENCE H. SUMMERS**, Charles W. Eliot University Professor and former president, Harvard University; former secretary, US Treasury; and former chief economist, World Bank

“AI will surely displace jobs and disrupt industries in the decades to come, driven by entrepreneurs who are implementing effectual thinking. The system-level changes that are on the horizon are excitingly discussed in this book, laying the bedrock for the oncoming revolution.”

—**VINOD KHOSLA**, founder, Khosla Ventures; cofounder, Sun Microsystems

“It takes courage to dive in and a willingness to invest time to reap the rewards embedded in these pages. But it is so worth it. The book is a hugely thought-provoking and inspiring primer on how to shape strategy and design organizations in the age of AI.”

—**HEATHER REISMAN**, founder and CEO, Indigo Books and Music

“This book is pretty damn epic. We’re often told that AI will be the most important thing humanity ever works on, yet it feels abstract and niche in its current impact on the world. The authors show us how these two sentiments are not in tension. They provide so many unique and rich examples to really help the reader understand why on a limbic level. It’s a perfect description of

this counterintuitive moment in AI’s history—a must-read for anyone who wants to peek around the corner into AI’s future.”

—**SHIVON ZILIS**, Director of Operations and Special Projects, Neuralink; board member, OpenAI; former project director, Tesla

“Nobody provides more insight into the fundamental economics of AI and what AI truly enables. It’s not just use cases for low-cost prediction—it’s better decision systems. That’s a much bigger step for business and the economy.”

—**TIFF MACKLEM**, governor, Bank of Canada

“Agrawal, Gans, and Goldfarb have done it again! Their new book, *Power and Prediction*, is destined to become the definitive guide to understanding how and why AI is transforming the economy.”

—**ERIK BRYNJOLFSSON**, Jerry Yang and Akiko Yamazaki Professor and Senior Fellow, Stanford Institute for Human-Centered AI; Director, Stanford Digital Economy Lab; coauthor, *The Second Machine Age* and *Machine Platform Crowd*

“Whether we like it or not, artificial intelligence is set to influence every aspect of our lives. How can we make sure that individuals, companies, and organizations benefit from it rather than waste time and resources dealing with unintended consequences? This readable book provides an excellent introduction, emphasizing how AI can improve what we do by providing better predictions and helping reorganize systems.”

—**DARON ACEMOGLU**, Elizabeth and James Killian Professor of Economics, MIT; author, *When Nations Fail*

Power and Prediction

The Disruptive Economics of Artificial Intelligence

AJAY AGRAWAL

JOSHUA GANS

AVI GOLDFARB

HARVARD BUSINESS REVIEW PRESS

BOSTON, MASSACHUSETTS

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel.

800-988-0886, or www.hbr.org/bulksales.

Copyright 2022 Ajay Agrawal, Joshua Gans, and Avi Goldfarb

All rights reserved

No part of this publication may be reproduced, stored in, or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior permission of the publisher. Requests for permission should be directed to permissions@harvardbusiness.org, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

The web addresses referenced in this book were live and correct at the time of the book's publication but may be subject to change.

Cataloging-in-Publication data is forthcoming.

ISBN: 978-1-64782-419-8

eISBN: 978-1-64782-420-4

*To our families, colleagues, students, and startups: you inspired us to think
clearly and deeply about artificial intelligence.*

Thank you.

CONTENTS

1. [Preface: Success from Away?](#)
2. [PART ONE](#)
3. [**The Between Times**](#)
4. [_1. A Parable of Three Entrepreneurs](#)
5. [_2. AI's System Future](#)
6. [_3. AI Is Prediction Technology](#)
7. [PART TWO](#)
8. [**Rules**](#)
9. [_4. To Decide or Not to Decide](#)
10. [_5. Hidden Uncertainty](#)
11. [_6. Rules Are Glue](#)
12. [PART THREE](#)
13. [**Systems**](#)
14. [_7. Glued versus Oiled Systems](#)

15. [8. The System Mindset](#)

16. [9. The Greatest System of All](#)

17. [PART FOUR](#)

18. [**Power**](#)

19. [10. Disruption and Power](#)

20. [11. Do Machines Have Power?](#)

21. [12. Accumulating Power](#)

22. [PART FIVE](#)

23. [**How AI Disrupts**](#)

24. [13. A Great Decoupling](#)

25. [14. Thinking Probabilistically](#)

26. [15. The New Judges](#)

27. [PART SIX](#)

28. [**Envisaging New Systems**](#)

29. [16. Designing Reliable Systems](#)

30. [17. The Blank Slate](#)

31. [18. Anticipating System Change](#)

32. [Epilogue: AI Bias and Systems](#)

33. [Notes](#)

34. [Index](#)

35. [Acknowledgments](#)

36. [About the Authors](#)

PREFACE: SUCCESS FROM AWAY?

When we published *Prediction Machines* in 2018, we thought we had said all we needed to on the economics of AI. We were wrong.

Although we fully realized the technology would continue to evolve—AI was still in its infancy—we knew that the underlying economics would not. That’s the beauty of economics. Technologies change, but economics doesn’t. We laid out a framework for the economics of AI in that book, which remains useful today. However, the *Prediction Machines* framework only told part of the story, the *point solutions* part. In the years since, we discovered that another key part of the AI story had yet to be told, the *systems* part. We tell that story here. How did we miss it in the first place? We wind the tape back to 2017, when we were writing *Prediction Machines*, to explain.

That year, half a decade after Canadian AI pioneers demonstrated the superior performance of deep learning for classifying images, interest was exploding in the new technology. Everyone was talking about AI, and there was speculation it could launch Canada onto the world’s technology stage. It was not a matter of if but when.

We founded a science-oriented startup program, the Creative Destruction Lab, with a stream devoted to AI. Everyone was asking, “Where do you expect to see Canada’s first AI unicorn—the first AI startup to reach a billion-dollar valuation?” Our bet: “Montreal. Or maybe Toronto. Or possibly Edmonton.”

We were not alone. The Canadian government was placing the same bets. On October 26, 2017, we hosted Justin Trudeau, the prime minister of Canada, at our annual conference on AI at the Creative Destruction Lab: Machine Learning and the Market for Intelligence.¹ In his remarks, he emphasized the importance of investing in clusters—geographic regions with diverse industry participants including large enterprises, startups, universities, investors, and talent, where the whole is greater than the sum of the parts, enhancing innovation and creating jobs. The key idea is that colocation matters. A few months later, his government announced

significant funding for five new “superclusters,” including one focused on AI based in Montreal.²

We felt confident in our beliefs about the commercialization of AI. We were, supposedly, world experts on this topic. We had authored a bestselling book on the economics of AI; we had published a swath of scholarly articles and management essays on the topic; we were coediting what would become the primary reference for PhD students in the field, *The Economics of Artificial Intelligence: An Agenda*; we had founded a program for the commercialization of AI that had, to the best of our knowledge, the greatest concentration of AI companies of any program on earth; we were delivering presentations around the world to business and government leaders; and we served on a number of AI-related policy committees, task forces, and roundtables.

Our perspective that AI should be viewed as prediction resonated with practitioners. We were invited to deliver presentations at Google, Netflix, Amazon, Facebook, and Microsoft. Gustav Söderström, head of product, engineering, data, and design at Spotify, one of the world’s largest music-streaming service providers, referenced our book in an interview:

[The authors] put it perfectly in their book *Prediction Machines*. Imagine the prediction accuracy of a machine learning system as a volume knob on a radio.... [W]hen you reach a certain point on that knob—when your predictions are accurate enough—something happens. You cross a threshold where you should actually rethink your whole business model and product based on machine learning.... With Discover Weekly we switched the paradigm from “shopping then shipping,” to “shipping then shopping,” the way [*Prediction Machines*] described. We had reached a level of [prediction] accuracy where we could switch from just giving users even better tools to playlist themselves, to just giving them a weekly playlist and let them save the tracks they really liked. We switched

our vision from “even better tools to playlist yourself” to “you should never have to playlist again.”³

Our approach—to design for a world where quality-adjusted prediction is very cheap—was of practical importance and provided valuable insight into AI strategy.

So, why were we so confident that the first AI unicorn would come from Montreal, Toronto, or Edmonton? We were in communication with two recent Turing Award winners (the equivalent of the Nobel Prize for computer science), recognized for their pioneering work on deep learning, who were based in Montreal and Toronto, as well as one of the primary pioneers of reinforcement learning, who was based in Edmonton. Canada’s government was about to generously fund three new institutions dedicated to advancing research in machine learning—in Montreal, Toronto, and Edmonton. Many global corporations were rushing to set up AI labs in Montreal (e.g., Ericsson, Facebook, Microsoft, Huawei, Samsung), Toronto (e.g., Nvidia, LG Electronics, Johnson & Johnson, Roche, Thomson Reuters, Uber, Adobe), and Edmonton (e.g., Google/DeepMind, Amazon, Mitsubishi, IBM).

It’s fair to say that we thought we knew a lot about the commercialization of AI. Yet, our speculation would have been wrong—a lot wrong. The first Canadian AI unicorn came not from Montreal, Toronto, or Edmonton. It would not even come from our second set of guesses—Vancouver, Calgary, Waterloo, or Halifax. If not from one of these cities—Canada’s technology centers—then where?

On November 19, 2020, the *Wall Street Journal* ran the headline “Nasdaq to Buy Anti-Financial Crime Firm Verafin for \$2.75 Billion.” Verafin is headquartered in St. John’s, Newfoundland. Very few people, certainly not us, would have predicted that the first AI unicorn in Canada would be in this town on the northeast tip of North America.

St. Johns, Newfoundland, is as far away from the action as you can get. Newfoundland is the easternmost province of Canada and has a population

of only about half a million people. It is not on the technology community radar. In fact, even though the United States is Canada's neighbor, many Americans learned about Newfoundland for the first time when the hit Broadway musical *Come from Away* was nominated for Best Musical and in four other categories in the 2017 Tony Awards. The musical is based on the true story of what happened during the week after the September 11 attack, when thirty-eight planes were ordered to land in Newfoundland and the humorously kind residents took in the seven thousand stranded travelers "from Away." But it was there, in Newfoundland, that Brendan Brothers, Jamie King, and Raymond Pretty founded Verafin, which eventually provided fraud detection software to three thousand financial institutions in North America. How could we have missed this? Was it a pure fluke? Random chance? Even the experts get it wrong once in a while. Hindsight is 20/20. There's always some chance of a low-probability event.

What NASDAQ was buying was AI. Verafin had invested heavily and built tools that could predict fraud and validate the identity of bank customers. These were key functions of financial institutions both in terms of their operation and also in terms of their regulatory compliance. To do this requires big data, and bank and credit union data was the biggest of them all.

Upon reflection, it wasn't so random that a business like Verafin might lead the pack. It was inevitable. Our focus on the *possibilities* of prediction machines had blinded us to the *probability* of actual commercial deployments. While we had been focused on the economic properties of AI itself—lowering the cost of prediction—we underestimated the economics of building the new *systems* in which AIs must be embedded.

Had we better understood that then, instead of assessing the landscape of prowess in the production of state-of-the-art machine learning models, we would have instead surveyed the landscape for applications focused on prediction problems where the *systems* in which they would be embedded were already designed for machine prediction and would not require displacing human predictions. We would have looked for enterprises that already employed large teams of data scientists who had integrated

predictive analytics into the organization's workflow. We would have quickly discovered that financial institutions were among the most prevalent, as they employed large teams of data scientists for predicting fraud, money laundering, sanction noncompliance, and other criminal behavior in financial transactions.⁴ Then, we would have looked for small companies that were embracing the recent advances in AI to address these problems. We would have discovered that there was only a handful of such companies in Canada at that time. One was called Verafin, headquartered in St. John's, Newfoundland.

We realized it was time to go back and think more about the economics of AI. Verafin's approach followed the playbook of *Prediction Machines*. There was no surprise there. What was less obvious was why so many other applications were taking so much longer to deploy at scale. We realized that we must consider not only the economics of the technology itself but also the systems in which the technology operates. We must understand the economic forces that led to the *rapid* adoption of AIs for automated fraud detection in banking and product recommendations in e-commerce on the one hand, but the *slow* adoption of AIs for automated underwriting in insurance and drug discovery in pharmaceuticals on the other.

We weren't the only ones who underestimated the challenges with implementing AI in existing organizational designs. Our colleague at the University of Toronto, Geoffrey Hinton, who earned the moniker "the godfather of AI" for his pioneering work on deep learning, also made forecasts that may have underestimated the difficulty in implementation.⁵ He had previously quipped, "[I]f you work as a radiologist, you're like the coyote that's already over the edge of the cliff, but hasn't yet looked down so doesn't realize there's no ground underneath him. People should stop training radiologists now. It's just completely obvious that within five years, deep learning is going to do better than radiologists."⁶ While he was correct about the pace of technical advancements—AI now outperforms radiologists in a broad range of diagnostic tasks—five years after his remarks, the

American College of Radiology reports no decline in the number of new students training in radiology.

It began to dawn on us that we have entered a unique moment in history —The Between Times—*after* witnessing the power of this technology and *before* its widespread adoption. Some implementations are what we call *point solutions*. They are straightforward. For them, adopting AI is just a simple replacement of older machine-generated predictive analytics with newer AI tools (those are happening quickly, like Verafin), whereas other implementations require redesigning the product or service as well as the organization to deliver it in order to fully realize the benefits of AI and justify the cost of adoption. In the latter case, companies and governments are racing to find a profitable pathway for doing so.

We shifted our focus from exploring neural networks to exploring human cognition (how we make decisions), social behavior (why people in some industries are keen to embrace AI quickly while others are resistant), production systems (how some decisions depend on others), and industry structures (how we've hidden certain decisions to shield ourselves from uncertainty).

To explore these phenomena, we met with company leaders, product managers, entrepreneurs, investors, data scientists, and computer scientists implementing AI. We convened workshops and conferences with experts and policymakers, and we examined up close what worked and what didn't with hundreds of VC-funded experiments in the form of AI startups.

Of course, we also returned to first principles in economics as part of a blossoming field of empirical research on the economics of AI that hardly existed only a few years earlier when we wrote *Prediction Machines*. We began to connect the dots and assemble an economic framework that distinguishes between point solutions and system solutions that would not only solve the Verafin puzzle but also provide a forecast for the next wave of AI adoption. By focusing on system solutions rather than point solutions, we could explain how this technology will eventually sweep across industries,

entrenching some incumbents and disrupting others. It was time to write another book. This is that book.

PART ONE

The Between Times

A Parable of Three Entrepreneurs

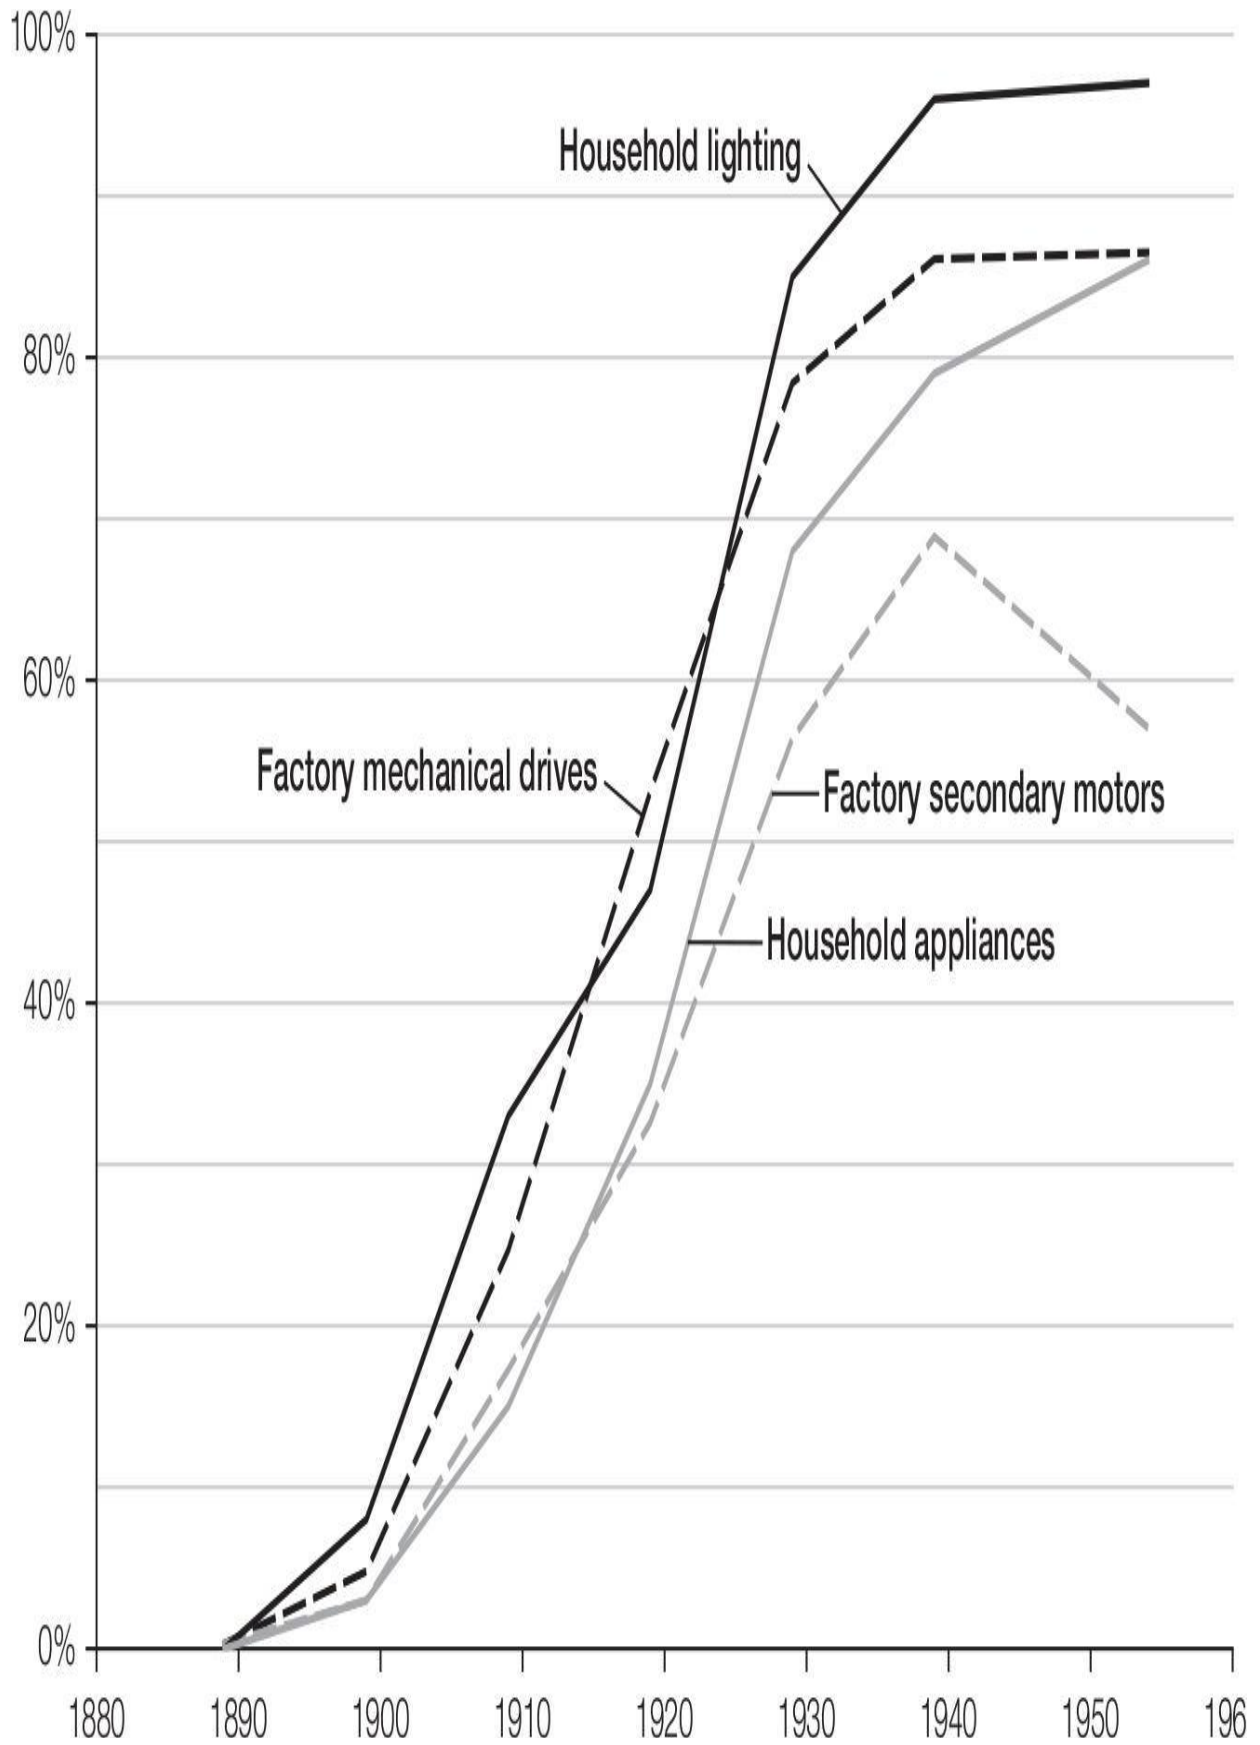
Electricity changed our society. It changed the way we live, providing inexpensive and safe light with the flick of a switch and reducing the burden of housework through consumer goods like refrigerators, washing machines, and vacuum cleaners. It also changed the way we work, powering factories and elevators. What did it take to do all of this? Time.

The ubiquity of electricity makes it difficult to imagine that, at the turn of the twentieth century, two decades after Thomas Edison invented the light bulb, it was pretty much nowhere. In 1879, Edison famously demonstrated the electric light bulb and, just a few years later, turned on Pearl Street Station in Manhattan and lit the streets. Yet twenty years later, only 3 percent of US households had electricity. In factories, there was barely more (see [figure 1-1](#)). However, after another two decades, that number accelerated to half the population. For electricity, these forty years were The Between Times.

There was plenty of enthusiasm for electricity but not much to show for it. We tend to forget this when new radical technologies emerge today. When the light goes on, rather than everything changing, little does. AI's light is on. But we need to do more. We are now in The Between Times for AI—between the demonstration of the technology's capability and the realization of its promise reflected in widespread adoption.

FIGURE 1-1

Adoption of electricity in the United States



Source: Data from Paul A. David, “Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror” (working paper #339, Stanford University, Department of Economics, 1989), [twerp339.pdf](https://www.warwick.ac.uk/twerp339.pdf) (warwick.ac.uk).

For AI, that future is uncertain. But we have seen the pattern for electricity. Thus, to understand the challenges facing the commercialization of AI, put yourself in the minds of entrepreneurs circa the 1880s. Electricity is the future. How would you envisage making it happen?

The Point Solution Entrepreneur

Steam powered the economy in the second half of the nineteenth century. Coal was used to heat water that generated energy, which was then immediately applied to drive levers, pulleys, and belts that in turn allowed for industrial production. By all accounts, steam was the miracle driving the biggest economic revolution since agriculture. So, an entrepreneur wanting to sell electricity would have to encourage would-be customers to take a closer look at steam and identify its warts.

When placed side by side with electricity, those warts were not hard to see. Steam dissipated heat—which was the point—but much of that heat was wasted. Steam power lost between 30 and 85 percent of its potential from condensation, leaky valves, and friction from using a shaft and belts to carry that power to workbenches.¹ The shaft system can be difficult to imagine. Think of a steam power source at one end turning a long three-inch shaft of iron or steel that then allows belts and pulleys to operate along the line. Some shafts may be horizontal, but many factories had multiple stories with shafts in a vertical configuration. One shaft could power hundreds of looms, for example.

The immediate opportunity for electricity was to provide an alternative source at the same point steam power was used—at the end of the shaft. Frank Sprague, a former Edison employee, saw this in 1886 when he developed one of the first electric motors. While Edison was focused on light, Sprague was among those who realized that daytime electric power would be cheap and electric motors could take advantage of that. Sprague used his