Aventine

ARTIFICIAL INTELLIGENCE

One Company's Decision to Put AI Into the Ground

Text
Gideon Lewis-Kraus
April, 2021

Photography
Lyndon French

What does it mean to adopt AI? For John Deere, it meant hundreds of millions of dollars, a new corporate culture and forever changing the relationship between a farmer and the land.

PART 1
THE COMPANY

A visitor in Las Vegas to any recent iteration of CES, the nation's largest consumer technology show, might be forgiven for concluding that the promise of artificial intelligence has at last been fully realized. No product remains immune to the immediate improvements of AI. A pet-food bowl uses advanced recognition technology to forestall overfeeding. A smart mirror promises a more personalized user experience than common mirrors allow. A nursery camera comes equipped with algorithms to monitor a child's sleep patterns, compare them to healthy averages and offer parenting advice. A countertop appliance draws upon a user's lifestyle, diet and exercise habits to distill custom nutritional beverages. TVs adjust their own settings; ovens both identify and then configure themselves to cook the food placed inside. In the span of just the last few years, AI — a shorthand compliment we pay to increasingly clever machines — has made previously unimaginable contributions to everything from tractors to toothbrushes.

Such extravagant orchestrations of AI, especially in the minor key of consumer convenience, seem likely to prompt two broad responses. For the skeptic, the vast hangars of smart tools on exhibit in Las Vegas represent disappointing delivery on excessive hype: The technology that was supposed to render human life all but unrecognizable has instead been enlisted to calibrate the magnesium content in a smoothie. For the believer, however, these are the fitful first steps of wholesale transformation — gimmicks as mere prelude.

But the prelude to what, exactly? The short-term specifics are often glossed over in celebration of the wholesale transformation to come: technological change so great as to overwhelm predictions. A certain inattentiveness to the details can be overlooked given the impact of the envisioned outcomes. As a 2018 McKinsey & Company report put it, AI "may usher in radical — arguably unprecedented — changes in the way people live and work," potentially creating additional economic output of around $13 trillion by 2030 and boosting global GDP by about 1.2 percent a year. A company that adopts AI quickly could see "additional annual net cash flow growth of about 6 percent for more than the next decade." The consultants' message to C.E.O.s is clear: Move forward or be left behind. The pandemic has only increased the sense of opportunity, as well as urgency: When life has to be conducted at arm's length, we look to machines to take up the slack. But irrespective of the headlines and the exhortations, the incorporation of AI, in a single business or across an entire industry, is not a matter of mere code. Omitted from these reports is a description of what any given AI strategy might resemble — the sort of granular account of the time-consuming and resource-intensive requirements that might say something useful about what it looks like to orient a business around artificial intelligence.

Continued on next page…

A John Deere combine driving down a back road in Jackson, Nebraska. Lyndon French for Aventine

This is not a lesson one would likely learn from any of the precocious little widgets on display in Las Vegas, most of which represent only marginal enhancements of existing appliances. In 2019, however, towering over the widgets, was an unusual CES debut: the S-770, a 22-ton John Deere combine harvester. This roving tornado of blades, encased in luminous green, could turn standing corn stalks into clean kernels at the rate of one million plants per hour while cameras monitored the output and a telematics device uploaded 500 variables per second to remote servers. “Why,” a Deere executive asked me at the time, “is John Deere at CES?” I admitted I’d wondered the same thing. “All the big things at CES this year, we’ve been doing them for a long time,” he said. “They aren’t toys, or gadgets that people play with on their countertop.”

John Deere has long been known for its legendary facility with big metal machines. Over the last 20 years, however, and especially over the last five, the firm has devoted vast resources to machine learning — AI, in other words, though the company tends to shy away from using the term. Partly that's because of a Midwestern aversion to showiness and a preference for plain speech. But on a more basic level, it reflects the fact that Deere sees its embrace of artificial intelligence as only the most recent chapter in a story of continuous mechanization that began on an Illinois prairie in 1837. The arc of the company's development, in fact, suggests that we might be better off not seeing AI as a specific goal to be accomplished but instead as a set of tools designed to achieve a variety of certain ends. This makes it easier to talk about what such technologies can and cannot be expected to do, what it could look like to "adopt" them, and where the future of widespread adoption might lead.

To examine AI as a series of concrete processes rather than a mystical goal does not suggest that recent advances in machine learning have not had and will not continue to have dramatic implications. Nor does it suggest that what Deere and other companies have been encouraged to do is easy. There are few, if any, off-the-shelf solutions. No one who reads a consultant's report is likely to imagine that adopting AI will involve putting people on airplanes with folding hand carts to walk thousands of miles of Australian soybean fields, sending data-science Ph.D.s to live with rural Canadian religious sects or retooling a company policy of mandatory drug testing. But those are among the accommodations Deere has had to make to recruit, integrate and retain the human intelligence it needs to support the artificial variety.

The paradox at the heart of "AI adoption" is that it takes an enormous amount of work to refashion a general-purpose technology to best serve specific tasks — and indeed that the most important technological innovations often occur not on the level of the general but in its diligent application to a specific use case.

This paradox is obscured by the view from far above, but is overwhelmingly apparent, as Deere is well aware, on the ground.

Continued on next page…

Deere Labs: The right problems

John Deere Labs takes up the mezzanine floor of a recently renovated building — plate glass, exposed brick — on 2nd Street in San Francisco, two blocks south of Market and a block north of LinkedIn's headquarters. Passersby tend to be constantly ringing the buzzer, most of them wanting a souvenir T-shirt or a green cap, or to get paraphernalia signed for their wistful grandfather. On occasion someone even drops by in search of parts. For efficiency's sake, the Deere Labs staff made up a small card with directions to the closest dealers, in San Jose and Santa Rosa.

Alex Purdy, a large and friendly man with neither the midshipman's posture nor the reserve I would later come to associate with Deere executives, officially opened the outpost in the spring of 2017, a few months after he joined the company. He would stay to oversee what would become a significant shift into AI technologies until he left the company in 2019. The office space could fit a start-up of perhaps a dozen employees, but for some four months he remained the only one. He'd grown up on a farm outside Calgary and for a long time he found the idea that he'd end up somehow back in agriculture about as implausible as the idea that John Deere would end up in Silicon Valley. He'd been working at Boston Consulting Group in Chicago on what he called "the intersection of optimization and industrial goods" — HVAC projects, for example — and John Deere was one of his clients. For two decades, the firm had been investing in what the sector called precision agriculture — the development and deployment of technologies that allow farmers to give industrial-scale acreage the attention and care accorded to small holdings. But like many legacy concerns, Deere had looked to consultants for assistance with its longer-term strategy. A consultant with no relevant domain expertise might be inclined to propose expensive technological solutions to nonexistent problems, but Purdy knew enough about agriculture to see what the adoption of AI techniques could mean for Deere and its customers. In 2015, he told Deere in no uncertain terms that they needed a Silicon Valley presence. It wasn't just that they needed bodies vaguely near Google (where, incidentally, his husband went on to work); it meant the construction of an entirely new digital infrastructure.

Continued on next page…

The John Deere World Headquarters in Moline, Illinois. Ariana McLaughlin for Aventine

As Purdy explained it, two related factors generated the urgency: the market consolidation in Big Ag — the industry's ever-narrowing cluster of major actors — and the market consolidation in farming operations. In 2013, Monsanto paid a billion dollars to acquire Climate Corporation, a venture-funded weather-insurance start-up that gave them a beachhead in information services. Monsanto was not a direct Deere competitor — it sold herbicides and seeds, not tractors — but the acquisition was a signal that Big Ag had begun to recognize the value of Big Data, a shift likely to change the competitive landscape in the long run. The next few years saw a frenzy of M&A activity on the seed-sales side of the business: ChemChina acquired Syngenta; Dow Chemical and DuPont merged (and later unmerged); Bayer bought Monsanto.

Farms, too, had been consolidating for decades. According to U.S.D.A. data, less than a third of farm production in 1991 came from outfits with a million dollars or more in annual sales; by 2015, farms of that size accounted for more than half of the country's agricultural output. There were many reasons for this, but concentrated market power on the part of manufacturers and suppliers was a significant one. As margins have gotten thinner, only the bigger operations have been able to survive. Purdy made the case to Deere that the ongoing amalgamation of farming operations put a new premium on the collection of comprehensive, detailed information — information of the sort that a farmer would once have absorbed from walking his own fields, but on a much larger and more intricate scale. It was an obvious boon for Deere that such information could be gathered by the same machines that exploited it. If, for example, you could develop an herbicide sprayer that sprayed only weeds, not crops, you could shift some of a farmer's cost from a perennial line item for herbicide to a one-time capital expense.

But the ability to do that at reasonable speed over thousands of acres was going to require more than a mechanical upgrade. Purdy may have been slightly ahead of the curve, and he may have had more expertise in his subject matter than is customary among consultants, but by the end of 2015 companies in all sectors were beginning to hear the same from outside strategists. The future of business would be built on artificial intelligence, and artificial intelligence would be built on data.

When Purdy told Deere executives that an investment in AI would be a significant and costly commitment, he meant that the concept in general — and machine learning in particular — only rarely allows for a ready-made fix. It requires a step back to examine the entire business from first principles. The best result stems from a reconsideration of the basic, underlying problems in any system; the worst is a fix in search of something that was broken. An argument could be made that the second outcome describes one of Deere's major recent initiatives, which didn't get a lot of traction.

Since the late 1990s, Deere's largest investments in advanced automation had been in the development of autonomous vehicles. General enthusiasm for the arrival of self-driving cars started in the last five to seven years, but research into how structured environments such as farms could be navigated autonomously had been explored for decades. In Deere's rush to sell fully hands-free tractors, however, the company lost sight of the fact that their customers neither needed nor wanted anything of the sort. While farmers were happy to be freed from driving the long, monotonous straightaways, they never took to the complex, awkward technology that allowed tractors to turn by themselves at the end of each row, and they simply didn't use it. Deere had made a classic mistake:

They'd started with theoretical engineering prowess rather than actual customer need.

In hindsight, the ideal of the fully autonomous tractor had come to look like technology for technology's sake: impressive, but what, exactly, was the point? Farmers weren't going to be getting out of the cab anytime soon. Purdy helped Deere understand that any technological investment made without the user's core objective in mind would likely fail. As he put it: "Nobody buys a tractor to go 'tractoring.' You buy a tractor to do a job. You want to focus on the act of planting. The more you make every seed and corn ear count, the better off the farmer will be."

Farmers were perfectly happy to absorb the trivial cost of hands-on turning at the end of each row. What they really needed, according to Purdy, was the kind of exacting managerial control that would allow for a bigger — and more reliable — yield. The money and talent spent on self-turning tractors would have been much better devoted to collecting and developing information that could be put to productive use.

He explained to me what his proposed pivot meant in practice. "A customer makes a hundred and forty decisions for each seed. I don't know much about seed genetics, and I don't want to. . . . When you take your data off your equipment at the end of the year — how much seed you put down, where you put it down, what you did to it, and your final yield — you can put it into DuPont Pioneer's Granular system, and then your seed adviser will help you determine those seed choices for next year — which hybrid, how much to put down, where — and then you'll work with an agronomist to develop a seed plan, a map, which by the end of the winter goes right back into your equipment."

Purdy stood up at a whiteboard and drew an irregularly shaped field, which he quickly partitioned into what looked like the schematics for a bank heist. "You've got clay soil here, hybrid1 here, hybrid2 there, 20,000 seeds per acre here, 30,000 seeds per acre there, and now you've got ten days to get it into the ground. The equipment is now going to change settings itself as you enter each zone — and do it automatically." With such a system, he explained, a farmer can introduce a planting scheme that is much more carefully planned but also automatically executed. No more defaulting into a one-size-fits-all approach or spending an enormous amount of time and energy trying to make individual plot-by-plot decisions.

To demonstrate, Purdy asked if I wanted to take a turn in their small on-site simulator, an apparatus slightly larger than a typical arcade game. Purdy made all of his engineers work out on the simulator; he wanted them to never forget the labor their software was supposed to replace. The simulator in San Francisco was preset to soybean harvesting, though in Des Moines the engineers can simulate every major crop and a variety of grower conditions in several technological eras. Each software change has to be back-compatible with 30 years of combines. Before I sat down, he pointed to an informal leaderboard scrawled over our heads. I commented that the employee at the top showed a very respectable 47.86 acres per hour. Purdy frowned. “People were plowing through,” he said, “and not actually harvesting correctly.” He said he might have to incorporate a quality metric into the rankings.

I sat in the simulator and Purdy clicked a series of buttons to begin my test. As I began to harvest the pixelated soybeans, Purdy grew alarmed. “You don’t want to have your thresher too low, or you’re going to grind up rocks!” I pulled the joystick to lift the thresher, which reared up into the pixelated air. “But you don’t want it too high, or you’re going to lose too many beans.” I tried to tell him that I was obeying the biblical injunction to leave the corners of your field ungleaned, but he reminded me that since the time of Moses, farming has been a business of slim margins, and that I had to go back and fetch the skipped beans. When I initiated what became an 11-point turn, he said, “Now you’re compacting the soil, which means that everything will be worse next year.” As I continued to torch and salt the earth behind me, he explained that as my combine’s grain tank filled, a second machine, synched to my trajectory, could show up and transfer grain elsewhere for storage, which would cut down on harvesting time by 20 percent.

Were I actually harvesting, my abysmal yield data would be uploaded automatically, allowing me to share it with my seed dealer, for example, and my agronomist. Their recommendations would be incorporated into a feedback loop that would determine an operational plan for the next season. The new challenge was to figure out how Deere’s implements could help *tighten* those feedback loops — adjustments that would take effect not the next season but the next instant. The company’s newest features, introduced in 2017, use machine learning to assess what is going on much closer to the ground, and in real time.

Purdy took me back to the rear of the office and picked up a heavy block of metal. “This is an AI board with Nvidia chips,” he said, referring to hardware designed for edge computing. “Doesn’t look like an Nvidia board, does it?” In fact it looked like the fortified viscera of an amphibious tank. He picked up a plastic toy combine — a child’s version of the hulking apparatus I saw in Las Vegas — to show me where the board would fit into the real thing, and why. “What you want is your grain to end up in the grain truck and your M.O.G. — material other than grain — to be spread back onto the field behind you. If it’s spread nicely, it’ll raise the soil temperature, which means you can plant earlier the next year.” He picked up two samples of possible combine output in sealed petri dishes. One had clean, intact grain; the other was an uneven mixture of cracked grain and M.O.G. He pointed to the second dish. “You take this one to the grain elevator, you’re going to get a significant deduction in price based on the sample. And you —” he paused, and said, almost sadly, “Well, you would’ve done much worse than this sample.”

In the past, he explained, you had to physically climb up on the combine and peer into its maw to monitor the quality of grain you were collecting. Then you would readjust accordingly. The new combines are equipped with an array of ActiveVision cameras that assess the grain coming in and the M.O.G. going out, and suggest settings to increase grain yield and quality. “This is a classic reinforcement learning problem — a multivariate optimization equation,” Purdy said, meaning a procedure to maximize output on the basis of a range of given inputs. Purdy projected at least a 5 percent improvement in yield. In a low-margin business, this was an extraordinary step.

The equations are hard enough, but Deere’s real competitive advantage is elsewhere. The algorithms wouldn’t be worth anything if the equipment weren’t rugged enough for a farm. Deere says there isn’t a digital-first company in the world that can match its ability to make a $15 camera “durable enough to withstand the harsh operating environments of the off-road ag industry.”

Deere Labs: The Culture

If the first step in “AI adoption” is to recognize that you or your customers have problems that lend themselves to novel deep-learning solutions, and the second is to prioritize those particular problems, the third is to build the kind of infrastructure AI requires.

Digital-first companies like Google and Amazon have had the required infrastructure in place since their inception, but legacy firms must build data collection and analysis into their existing practices.

Deere is one of the few major American companies that have been around since before the Civil War, largely because it's very good at making and selling equipment. When Purdy was pitching increased investment in AI, he emphasized that the company's digital investments had to be rooted in that core competency. In an era of profligate AI hype, it was easy to indulge the fantasy that "AI" was a magical profit driver, and Purdy's admonition that one should not put the AI cart before the business horse was almost radical: Deere had to retool its existing products — tractors, planters, sprayers — so that they no longer simply performed jobs but monitored and reported their own performance in ways that could be fed back into the overall system. Harvesters needed cameras to track what was going in and what was coming out; planters needed sensors to record how much pressure was needed to insert each seed at the proper depth. On the tail end of the feedback loop, they needed actuators that would put this data to use — to adjust fan speeds, say, or modify planter pressure. Deere was already competent with hardware. The real challenge was building a data platform that would relay information, process it and ultimately deliver it back to the farmer in useful form.

Deere could have chosen at this point to build a closed data system — one in which a customer would be locked into Deere's infrastructure and forced to use Deere's proprietary services — or an open one that could be accessed and built upon by third parties. Like any company making a transition into data-driven operations, Deere had to decide what business, exactly, it was going to be in. Many hardware companies have elected to sell their tools below cost, collect vast amounts of data from their customers and either use that data to generate ancillary revenue or sell it on to other manufacturers.

Deere, however, had never been in the data business, and knew that issues related to security and, even more important, privacy, could easily undermine the consumer trust the company had cultivated over generations. Executives resisted the popular notion that an "AI company" had to be a data-first company. A future that relied on capturing and using data did not require hoarding that data — in fact, the company believed it was better not to. Deere sold equipment and wanted people to buy more of it. The more easily its platform could be shared with seed dealers, herbicide dealers, agronomists and other specialized businesses that could help make the data useful to a farmer, the more effective that equipment would be and, presumably, the more of it farmers would buy. A primary challenge was thus to design a platform for wide-ranging interoperability, one that would not only adhere to current agtech data standards, but would also help expand those standards to new types of data and new geographical markets, the latter a task that only a firm with Deere's size and standing could effectively do. Deere's machines would collect data at every point of the farm cycle, then facilitate the seamless and instantaneous transfer of that data within its own operations and to trusted third parties. Larger growers with highly centralized outfits could coordinate operations on multiple machines. "Say the manager has only one really good combine operator and two others who aren't as good," Purdy said. "He can lift the settings from the good guy and send them over to the others. The automation allows you to put lower-skilled operators into the cabs and get the same results as the better-skilled people."

Building the data platform, which began around 2015, was a grand undertaking. The company had to develop and maintain global-positioning and telematics systems that would work equally well for the owner of brand-new equipment in the Sacramento Valley and a farmer using third-hand equipment in rural Kazakhstan. It also had to ensure that a cloud-wary farmer could opt out of remote data storage and instead port data around on USB sticks. Virtually nothing about Deere's data collection, data transfer or data management lent itself to a plug-and-play solution. Unlike, say, a recommendation engine for an e-commerce company, you couldn't just buy or scrape existing data and process it in Amazon's cloud.

The application of AI to something as complex as the international agricultural sector might seem like an exceptional case, but in the grand scheme of human industry the real exceptions are the sectors for which off-the-shelf AI solutions already exist.

In a few short years, Deere spent hundreds of millions of dollars creating its data platform, called Operations Center. For a company with a market cap of about $50 billion at the time and annual revenue of $30 billion to $40 billion, this was not an enormous sum in relative terms, but it was a striking investment for a project that had only an indirect relationship to Deere's bottom line, which was still driven by sales of colossal metal contraptions. "Building out this vision of how data can connect across operators has meant hiring people in data analytics, edge computing, cloud computing, A.P.I.

and S.D.K. development, and one main reason to be here is to get new talent," Purdy said of his emphasis on Deere having a Silicon Valley presence.

In the past, Deere had a reputation for, as Purdy put it, "hiring people from college and growing them up." The company had already spent more than 20 years building out its I.S.G., or Intelligent Solutions Group, a precision-ag unit of more than 800 people based in a suburb of Des Moines. Purdy understood the increasingly competitive labor market for machine-learning talent, and told Deere that it would soon lag too far behind without a dedicated outpost in one of the few AI hubs. "There's an institutional community around AI that it's beneficial to be a part of," he said. "It moves so fast, and if you're not around the people doing it, you can't keep up."

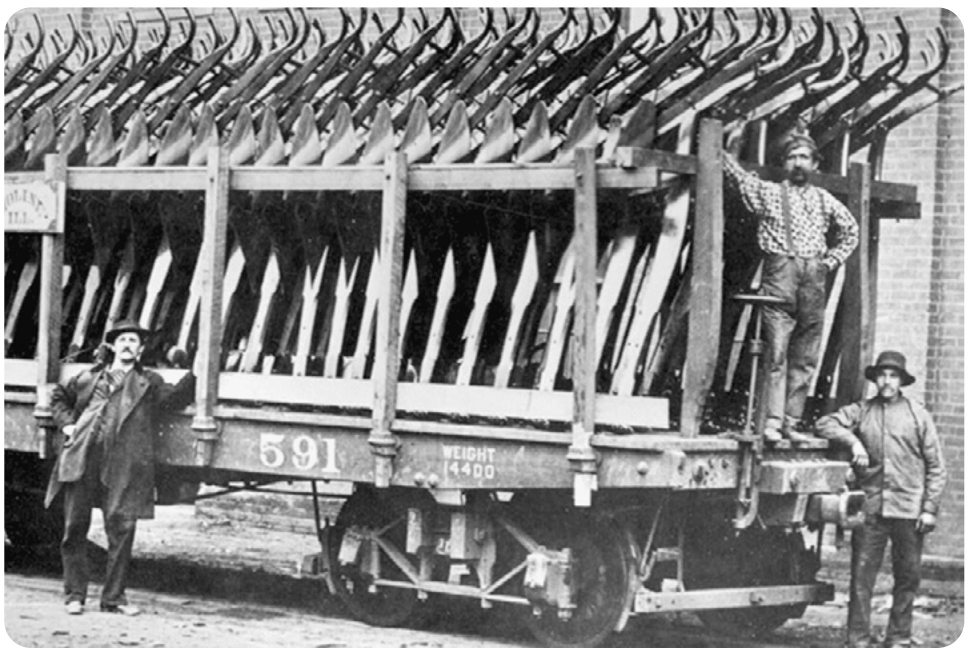

Deere plows, circa the 1880s. Photo: Library of Congress

But Purdy was picky about his team because he knew how common it was for companies to assume that a data platform could be imposed on the customer by fiat. On the contrary, it had to be constructed by engineers who understood how and why it would be used. Even with standard-issue consumer products, it's easy for engineers to mistake their own experience and desires with those of an average user. In something like agriculture, the gap between an engineer in San Francisco and a farmer in Iowa is even more considerable.

One reason Purdy was drawn to Deere was that it recognized that the most important factor in a successful tenure there is empathy with the customer, a quality the company went to some lengths to cultivate. Over the past few years, Deere has developed a yearlong program called Through Their Eyes, which pairs engineers with a dealer and a customer over the course of five or six farm visits that coincide with the stages of the crop cycle. Shortly after joining up with the company, Purdy himself had been placed in Boone, Iowa. Other participants had been sent to rural Mississippi and even to remote Canada to break bread with Hutterites, Amish-like communities that have no televisions in their homes but might buy a dozen brand-new tractors in a year.

Purdy understood enough about company culture to look for recruits willing to put in the time to go out and meet the farmers.

To his relief, Purdy found that these mucky commitments didn't deter as many recruits as he'd feared. If anything, the nature of the task helped his cause. The pitch was simple: You can either go work on an app that puts dog ears on selfies, as I heard one executive put it, or you can help feed the 1.8 billion people who will be born in the next 50 years. The clarity of the problem and the goal — to sell farm implements that would increase productivity — and the obvious definition of the problems — "Optimize combine operations to improve the quality of the grain yield" — presented an appealing challenge to the engineering mind. "It's not AI for AI's sake," one of Purdy's recent recruits told me. "The problems are both well defined and also extremely complicated. Like, 'weather' is just one input in the whole thing. We're coming at it from the perspective of, This is what I'm trying to achieve, how do I use technology to get there?"

But if AI recruits were expected to adapt themselves to the company culture, the company also had to adapt to the new recruits. Deere, for example, has long had a zero-tolerance policy on drugs, as many Midwestern manufacturing corporations do, and that discouraged some prospective employees. "Working on that with HR," Purdy said, "is a little uncomfortable but sorely needed." There was also the matter of the dress code. Deere's Intelligent Solutions Group used to have the kind of dress code that was part of doing business in Des Moines. "You could wear jeans but the list of rules was a page long — no shorts, shoulders covered. A year later, our dress code has been shortened to 'dress appropriately,' and as it turns out the world has not fallen apart," said Purdy, pausing for a moment. "But we recently had the C.E.O.'s staff here, and there was one dude with a ripped T-shirt. I said, ‘Can you please put a coat on for today?’” Purdy had recently gone to a Lesbians Who Tech summit and set up a booth there. “We can’t have only mechanical engineers from Iowa State.”

The Acquisition

Even once a company (a) determines that it has the sort of problems that can be solved by machine learning; (b) decides to prioritize said problems; and (c) sets off to create the infrastructure to support a machine-learning solution, it still has to evaluate what can best be done in-house and what can be better sourced elsewhere. The idea behind Deere’s investment in precision agriculture in the late 1990s was that greater productivity would come not from scale alone, but from technologies that worked at scale with the intimacy and flexibility one would devote to a smaller operation. While farmers have traditionally made decisions at the farm level, Deere saw a future in which decisions would be made field by field, then zone by zone, and ultimately plant by plant. Deere was already at work on these techniques, but the company couldn’t help noticing that this sort of fidelity — actual plant-level attention — was the aim of a small start-up called Blue River Technology.

Jorge Heraud, whose only exposure to agriculture was the time he spent pulling weeds in his grandfather’s tomato patch in northern Peru, is a co-founder. He had studied electrical engineering and engineering management and worked at a company designing self-steering tractors. In 2011, when he came out of a one-year executive M.B.A. program at Stanford, he and a graduate student in robotics started Blue River Technology. They had decided that precision agriculture seemed like a green field, so to speak, for entrepreneurial activity, and elected to begin with the problem of weeding.

They looked first at lettuce. Silicon Valley wasn’t far from Salinas, known as the lettuce capital of the world, where it was grown year-round. While some variants of corn and soybeans are genetically modified to be resistant to herbicides, and can thus survive blanket applications, spraying a lettuce field will kill both the weeds and the lettuce. As a result, crews have to go through and weed the lettuce by hand. “It’s super labor-intensive. They go plant-by-plant, using a machine called a hoe,” he explained, making it clear that his time in

Silicon Valley had given him approximately zero ambient familiarity with even the simplest aspects of agriculture. Furthermore, growers complained of a shortage of seasonal laborers willing to wield those hoes — possibly reflecting the unattractiveness of wages. As far as cash crops went, lettuce was small potatoes, but it seemed like a good opportunity to refine a labor-saving technology.

One of their initial plans was to develop a tractor-pulled sprayer that would use machine learning to differentiate weeds from crops: They called it See & Spray. But the machine-learning methods then in vogue, just at the cusp of the machine-learning revolution, only gave them an accuracy rate of around 80 percent, and no farmer was going to buy a machine that relieved him of a fifth of his crop. Heraud and his team knew that they needed an accuracy rate that approached 100 percent: "If you're only right 90 percent of the time, in fraud detection that's pretty good. But in ag, the farmer fires you."

By early 2017, however, the advances in deep learning, as well as major steps forward in edge computing, enabled them to improve their product dramatically. They found themselves in the position to pivot away from lettuce and into the cash cow, so to speak, of the major row crops — corn, soybeans, cotton — that account for the majority of agricultural production in the United States. The broadcast application of weed killer to acres of land planted with herbicide-resistant crops is tremendously wasteful; an implement that sprayed herbicide only on weeds could help large commodity farmers diminish their herbicide use by an order of magnitude. This meant not only cost savings but a welcome de-escalation of consumers' exposure to harmful chemicals. This would have an array of salutary knock-on effects: If a farmer no longer had to rely on the broadcast application of herbicides, for example, he would no longer need to buy expensive herbicide-resistant seeds. And if such a discriminating approach worked for herbicides, there was no reason it couldn't someday work, conversely, for nutrients.

At the time Blue River set about training its machine-learning systems, Google's object recognition technology could differentiate between a baboon and a spider monkey with ease, but none of the off-the-shelf neural networks could spot the difference between a weed and a plant. Extensive datasets of specific, labeled plants aren't exactly lying around, and the work Blue River had done with lettuce wasn't going to help them when it came to cotton. The company wanted to get started with data collection as quickly as possible, but it was winter in the Northern Hemisphere. The solution was to send a team to Australia for three months of antipodal summer to walk miles of cotton and soybean fields and come back with hundreds of thousands of images of the young plants. They used a combination of an internal team and Mechanical Turk labor to label the images for machine consumption. But even this data was only of middling use. It turned out, for example, that the system they'd trained on Australian cotton didn't recognize the varieties grown in Arkansas, which meant gathering additional imagery from the American South. Also, lighting conditions varied, and identification had to be accurate for wet plants as well as dry ones; a hailstorm in Texas informed them that their neural network was mistaking battered cotton for weeds.

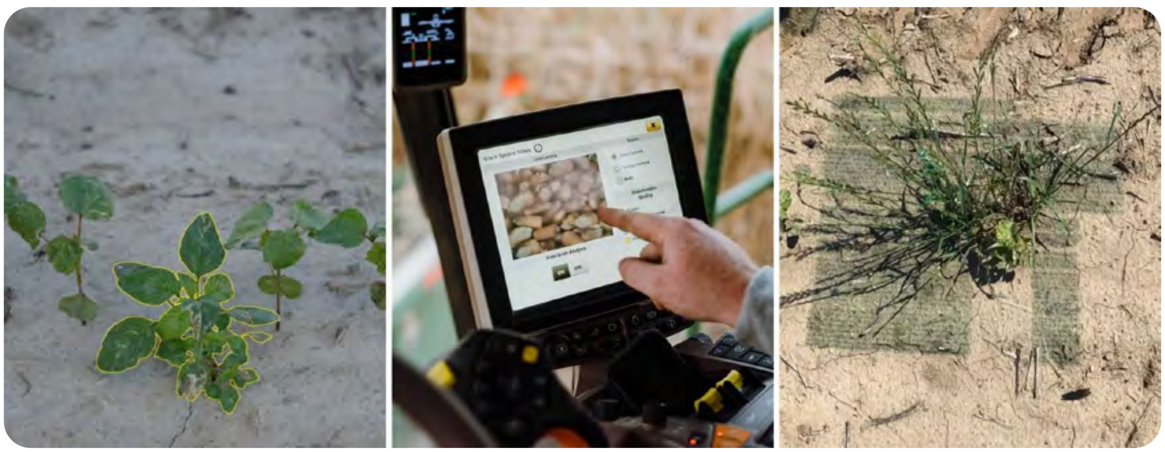

FROM LEFT TO RIGHT: See & Spray technology recognizes a plant with computer vision; a screen allows for a real-time scan of grain quality; dye shows how See & Spray can apply herbicide on and around weeds. Center image by Lyndon French for Aventine; Outside images from Blue River

The rigors of the farm environment made hardware a challenge, too. Cameras had to withstand pressurized washing, for example, and Blue River had to deal with the effects of age on mechanical parts like the sprayers' solenoid valves. One of the company's earliest employees, Matt Colgan, was a Stanford aeronautics and astronautics student who, as an undergraduate, had helped design the New Horizons spacecraft, which traveled to Pluto. He was good at harsh environments. "When you write software," he told me, "you usually don't have to worry about the radiation belts around Jupiter, which can switch your 0s to 1s." (This is why, to take a terrestrial example, the robots brought in to clean up Chernobyl and Fukushima almost immediately died, and why the claim that automated agents will be useful in nuclear disasters is highly dubious.) Still, Colgan hadn't expected that his deep-learning job would include maintaining 36 mobile greenhouses (shipped from Amazon) for the field-test site, to help with soil temperature. One day workers forgot to replace the sandbags anchoring the greenhouses to the ground. The wind picked up, and the greenhouses flew off. "I didn't think that when I left Stanford with an aerospace Ph.D. I'd be picking greenhouses out of ditches," he said.

Heraud was initially interested in John Deere as a potential distributor for Blue River's own sprayer. The value of the sprayer was obvious to Deere, but equally obvious was the fact that Blue River did not need to be in the manufacturing business. What Deere already had — production experience with rugged machinery, as well as an educated dealer network that didn't just sell to customers but knew how to train them — was what Blue River lacked. As the talks continued, the possibility of a distribution agreement became the possibility of an investment, which ultimately became talk of an acquisition.

The team at Blue River knew that the acquisition would ultimately depend on an in-field demonstration of its technology. So on a summer day in 2017 just after dawn, there they were in a cotton field in Yuma, Ariz., connecting over video chat with the Deere executives in Moline. As the sprayer went along, it deposited jets of blue ink on weeds the size of postage stamps. John May, who became Deere's C.E.O. in November 2019, was among those watching on a giant movie screen in the executive boardroom. At the end of the demonstration, he turned to a Blue River manager in the room, took his glasses off, and asked about the rear cameras the team had installed on the planter. He had one question. "Can you do that on a harvester?"

Deere is a conservative company, but no eyebrows were raised when the engineer said, "Fuck, yeah."

The Deere execs passed around green hats to the Blue River team. The September 2017 acquisition — for $305 million — was one of the biggest bets Deere's agricultural division had ever made.

Blue River, like most acquired companies, did not get everything it wanted. While its employees got to keep their own email addresses and open-plan bring-your-dog start-up culture — or at least as much of it as could be preserved — the company had to abandon one of its major programs, an effort to use drone-mounted cameras for fast plant phenotyping. Multispectral cameras measured leaf area and health, and thermal cameras measured water stress. Deere's focus was on collecting such information on its own equipment and in real time, so Blue River shifted everything it had mounted on drones to equipment in the field. That meant some accommodation. "Take their boom sprayer. It's a hundred and twenty feet wide," one Blue River engineer told me. "Blue River was afforded exactly so much mass and so much space on the boom to fit our cameras."

What was most important, Blue River people reiterated again and again, was Deere's focus on the challenges that farmers face, so many of which could be formulated as problems deep learning could solve. While one can't specify with any precision exactly how to improve a farm's yield, for example, one could nevertheless make good inferences based on a lot of data. "Pretend this table is a farm," one Blue River engineer said to me, as he began to assemble a small army of plastic toys in a neat row. Selecting a dwarf model of a tractor with a plow attached, he said, "OK, so you've just harvested, and your grain is in an elevator. The tillage machines move slowly and turn over the soil, and you want clods of an approximate size. If they're too big there's not enough airflow, and if they're too small you get sandy, compact soil. Right now you set and forget — one size fits all — but it's much better if you perceive what's going in and coming out and optimize." You couldn't necessarily explain the steps to get there, but you could recognize and specify what success looked like and let the machine figure it out.

He put the tillage toy back in the neat row and reached for the planter toy. "You don't want to plant a corn seed in a leaf, first of all. You're digging a V-shaped trench, and the planter needs to shoot the seed down a tube backwards to hit the trench and bounce vertically. You don't want them too close together or they compete for resources. If they're too far apart you waste land — and we're talking about millimeters scaled across thousands and thousands of acres. After you plant, your closing wheel closes the trench up. If you do it too hard, the seed has to waste energy pushing up through soil. If you do it too loosely, there are air pockets. These are all mini robotics and computer-vision problems, and they have to take into account soil humidity, uphill versus downhill, et cetera, to get the most out of every single seed in real time — which is hard when the planter moves 20 miles per hour and is dropping seeds every few inches."

I had gotten the point, but we were only halfway through the line of toys, which were showing their own signs of asset depreciation. The engineer clearly didn't want to miss the chance to show me the sprayer — "What does this crop need? Is it nitrogen-stressed? Water-stressed? Bugs on it? We only want to give nitrogen to the nitrogen-stressed" — and, finally, the combine harvester, which had to separate the pure, high-quality grain from the M.O.G. it spread out the back. He finished with a final description of the problems of stalk-width variance and accurate biomass perception, and took a deep breath. "You see? Machine learning to the rescue."

He was making an important distinction here, but one that wasn't necessarily obvious. The concept of "machine learning to the rescue" is very different from the concept of "the S-770 Combine to the rescue"; the emphasis is not on the products but on the extended processes in which the products are embedded. The effects of this would only become clear over time. All the data that was collected from the earlier generations of connected combines would be used to further train the learning algorithms. Better data produced better algorithms, and then better algorithms produced better data. In the summer of 2020, Deere released a newer version of the combine I'd seen in Las Vegas. Jahmy Hindman, the company's newly appointed C.T.O., told me that the new machine's physical volume was almost identical to its predecessor, but that improvements in the software — in only one year — had achieved a 45 percent improvement in total harvesting capacity. See & Spray was on a similar trajectory. Over the summer of 2020, Deere had worked with a network of trusted growers to pilot the machines on different crop types and in different growing conditions. Over the course of only one season they had become confident enough in the product's performance that they would soon be put into production at scale. These experiments with feedback cycles had inspired, over the first half of 2020, a dramatic reconfiguration of the company itself. "We reorganized the business away from a product-centric approach — planters and combines — into a business that supports the production systems in which our customers operate — corn, soy, grains," Hindman said, and paused. "It's been a complete shift in our mentality."

PART 2
THE CUSTOMER

Taylor Nelson represents the fifth generation of Nelsons to have worked his family's land, but it was not a foregone conclusion that he would take up the mantle. No one in his family had ever gone off to college and decided to come back to the farm. Taylor had gone off to school at the University of Nebraska, Lincoln, some two hours due south of his home in Jackson, and returned, in 2012, with a degree in agricultural economics. His younger brother followed him — he became the farm's agronomist — but few members of his generation plan to have anything to do with farming. According to U.S.D.A. data, the average age of a "farm producer" is 57.5; it hasn't been below 50 since Eisenhower was in office. About a third of farmers are over 65.

When Nelson was a kid, there were still little farmsteads — with livestock and children — on the corner of each little parcel of land, but agricultural consolidation had left the softly rolling landscape increasingly depleted of family outposts. Nelson had a business to inherit only because his family had been in a position to scale up. In the 1990s, his father, Doug, and uncle had embarked upon an expansionary era for Nelson Farms, taking them out of livestock and devoting all of their acreage to corn and soybeans. So while his grandfather and great-uncle had worked 1,200 acres of diversified cropland and raised hogs and cattle, Nelson, along with his family and a single full-time employee, now farms 11,000 acres of only corn and soybeans.

In the spring of 2019, we drove over the river bluffs and down to the flats; his little farmstead stood in the distance. He told me how nice it was at his home in the summertime, when the crops grew up all around and obscured the view in all directions. At night, all you could hear were the irrigation pivots running. It felt like you were in your own little world, and you could almost convince yourself that not much had changed since your grandfather's day. But of course it all had. Back then, his family farm was not only smaller but much more diversified. His grandfather's hog and cattle facilities were long gone, swept aside in the streamlining of operations that technology at scale had both warranted and demanded.

Consolidation of farm operations wasn't something that happened once, with the transition from horses to the tractor. It has been a never-ending process, as scale in farming everywhere has fallen victim to a runaway feedback loop. Margins are increasingly thin: The consolidation of manufacturers and suppliers drives up input costs, while oversupply keeps prices down — all this in a largely monocultural environment that is increasingly vulnerable to macro forces like climate change and international trade. A single presidential tweet that happens to drive the price of soybeans down a dime can mean the difference between a profitable year and a loss. Volume is the only way to survive, and technology is the only way to achieve volume.

Taylor Nelson, left, with his father, Doug Nelson, right. Lyndon French for Aventine

The paradox, of course, is that technology is at once the solution to and an aggravator of these problems. You buy a fancy seeder to plant more efficiently, a fancy sprayer to protect your crops more efficiently, and a fancy harvester to cull every last kernel of yield. These costs can only be justified when the machines are put to use across more acres, but the more acreage you have under cultivation, the more technology you need to get a handle on it. In the new era of Deere's Operations Center, all the data a farmer needs now comes from sensors newly embedded in machines. Not availing yourself of this new technology means either limiting the size of your operation or reverting to one-size-fits-all guesswork that's bad for the farm, bad for the crops, bad for the consumer and, for a variety of reasons, bad for the planet. "You're operating on a big enough scale that you can't wrap your arms around every field every day," Nelson told me. And consolidation isn't just about total acreage but about field shape and location. To take advantage of the machines' capabilities, you need to consolidate your own holdings such that your fields are contiguous. Over the last two decades, Nelson Farms has been constantly enlarging and contracting its acreage to reduce a sprawling archipelago to something a little more tractable; even now, their 12,000 acres are spread over 61 parcels of land.

"You look at the last 50 years of farming evolution, and it just doesn't seem possible," Nelson said. "You put my grandpa's tractor next to our tractors, and it looks like a little toy. It's fun to tell my grandpa about what we're doing now. He just shakes his head. He was looking at my sprayer the other day and said, 'You gotta get some wheels on the end of the boom so it doesn't hit the ground,' and I said, 'Grandpa, it has radar! It's not all levers anymore!'"

And all of the new technology is expensive, and risky. A combine his grandfather bought in 1979 cost $50,000. One of the new machines Nelson just purchased cost $400,000, more than double the price in adjusted dollars, and is both vastly more productive and vastly more complicated. In 2019 Nelson Farms purchased three new S780 combine harvesters. While the same model, which Deere calls a "factory on wheels," had looked impossibly colossal and alien next to the smart pet-food bowls in Las Vegas, the three machines side by side looked merely enormous in Nelson's hangar. They were outfitted with the company's new ActiveVision camera system, which uses computer vision and machine learning to optimize the harvest settings in an algorithm called AutoMaintain, which works as part of a system called CombineAdvisor™; this was more or less the mechanism that Purdy had outlined with his child's toy and his petri dishes. Initially, Nelson had wanted to start with only two ActiveVision-driven systems, but he'd been talked into buying a third combine. Uniformity of output is crucial, and the machines work best when they are all the same. "There are only so many things I as an operator can do at any one time," Nelson told me. "The systems here can take over stuff you'd otherwise neglect. With these, you can optimize the settings based on the quality of the final grain sample, and that's something new."

But technology can only get a farmer so far; there are so many things beyond human control. It was already the second half of May and he'd barely begun his planting for the season. Ideally, everything would have been in the ground two weeks earlier — the rule of thumb was early May — but the Midwest had seen unrelenting rain throughout the spring, not to mention a bomb cyclone. Portions of Nelson Farms' holdings were in the flood plain of the Missouri, and when their fields weren't underwater they were ribboned with flotsam. It had also been cold. The first day they were able to put soybean seeds into the ground was the 24th of April, and since then they hadn't had even two consecutive good days to string together. When I visited him, it was an hour after dawn on a beautiful morning, warm enough at 7 a.m. that the mist had mostly burned off the hilly, naked and dun-colored fields, and Nelson was optimistic that he'd get a good 14 or 16 hours in his planter.

It's easy, with the inaugural use of any technology, to assume and hope that the novelty introduced is simply a better solution to an existing problem. But the use of new machinery changes all sorts of other considerations in unpredictable and nonlinear ways. Take the example of large planters. Increased acreage changes how you think about what you're putting into the ground. Nelson needed a seed hybrid "reliable enough to take care of itself." This choice involved its own set of trade-offs. You might want to make a slight sacrifice on total yield in order to prioritize hybrids that were particularly drought- or wind-tolerant. If you have areas of sandy soil, you would choose a seed hybrid engineered for greater leaf area to shade the ground and prevent evaporation. Of course, you could choose a different hybrid for each area of each field, but that would become a lot to keep track of, and threaten the whole point of economies of scale. Nelson worked with only six or eight hybrids for all of his fields, which meant he'd selected varieties suited for a wide range of conditions. Cultivar selection and breeding has been around for thousands of years, of course, but from Nelson's perspective it's another pressing, information-driven expense paid to another price-setting, information-driven conglomerate — and one more set of variables to be fed into the vast analytics engine that is the data-driven farm.

All of the technology that went into seed development — special treatments, seed genetics, special trait packages that mitigate environmental factors — didn't come cheap. In the 1990s, Nelson could get $2 per bushel of corn. The price has almost doubled since then, but the farm's input costs have risen much more quickly, cutting deeply into margins. In the 1990s, Nelson's family put about $200 per acre into the ground; now they spend four times that. Movements of input prices aren't always rational. Ten years ago, corn prices shot up to $8, and the expenses followed — especially the value of land. Commodity prices have retreated considerably since 2014, but it was a slow drag waiting for costs to follow suit, and in the meantime smaller farmers were driven out of the business.

Finally, there's been the effect of technology on labor, not only who labors and how much they're paid but the nature of the work they do. When farms were diverse, there was year-round work for utility laborers; with livestock around, there was always something to do. On the homogenized, one- or two-crop farm of today the demand is seasonal, and difficult to meet. This has required more technology to compensate, which then further dampens the labor market — another runaway feedback loop. For the vanishing number of people who still work the land, the job itself has become almost unrecognizable: The more complicated the technology, the more training required. "Since I came back into the operation," Nelson said, "I've had to become a tech specialist. You have to have somebody to oversee successful implementation, to train people, to troubleshoot. When you're out there at midnight and have a glitch, you have to be self-sufficient and figure it out."

Continued on next page…

Doug Nelson starting up the combine at the Nelson Farm in Jackson, Nebraska.
Lyndon French for Aventine

Outside his farmstead, Nelson took me into a shed he'd converted into a small-scale seed treatment facility. He had a machine of Rube Goldberg convolution that applied little red dots — combinations of fungicides, insecticides and inoculants that, among other things, might allow seeds planted early to withstand adverse soil conditions. At the end of the process, a large tumbler rolled the seeds around until their coatings were even. The setup would have been $100,000 new, but he'd recently gotten it at auction for only $30,000. Treated seeds purchased from a distributor could cost $12-14 per unit — about an acre's worth — but the new contraption allowed Nelson to do it himself for only $5 per unit. At six thousand units per year, the machine saved them $42,000 annually. But there was no way to account for the machine's cost in increased complexity — the time dedicated not only to its construction, operation and upkeep but to the continuous adjustment of its settings, not to mention the additional labor that went into Nelson's self-training. "There's no way my dad or uncle would even attempt to operate a treater with all the electronics involved," Nelson said sympathetically, "but I figured I could figure it out." The unit economics of farm investment are thus far from straightforward: Increased efficiency in one dimension can lead to a host of new complications in others.

Nelson ushered me onto the back of a four-wheeler and drove us at breakneck speed along a rutted track out to the morning's field. He slowed as we approached his planter, which sat behind his tractor like a 120-foot-wide rake. He stopped and calmly said, "Well, this is a disaster."

FROM LEFT TO RIGHT: A cob of corn in the combine; a camera that monitors corn as it comes through the grain elevator; load cells (the silver-colored discs) record the quality of the yield as it is being processed. Lyndon French for Aventine

The field's enormous irrigation pivot — the long radii of wheeled scaffolding that, rotating around a central pivot, create the lush circles one sees from the air — had driven up onto the planter apparatus. The two huge pieces of metal were entangled in a smashup of broken-off press wheels, busted hoses and exposed wiring. The pivot was controlled by a cellphone. Someone had turned it on the previous night to move it out of the planter's way, but whoever had done so had forgotten to set the auto-stop for the correct position, and the pivot had continued on its automatic rounds until it collided with the planter.

Even when everything can be arranged at the press of a button, someone still has to remember to press the right button. And there is an ever-increasing number of buttons to push.

Nelson made a phone call. (Later, the Deere PR rep said, "I bet you heard some cuss words!" I did not.) They were going to have to get a winch out to the field, lift the pivot off the planter, and get the planter into the shop for repairs. (Deere has been criticized for the fact that its proprietary hardware and software makes it almost impossible for a farmer to repair his machines himself.) The whole thing might only take a morning, if they were lucky, but the planting window, already tight, was closing.

The Plow and the Tractor

Nelson, who has seen technologies change his livelihood on a season-by-season basis, has a particularly acute sense of innovation in all of its flowering unpredictability. But even while he regards his workflow with greater dizziness than he suspects his grandfather did, the feedback loops between new machinery and the consequences of its implementation characterize the adoption of any tool. The history of Deere itself makes a vivid case for automation as a social process. As Doug Sauder, Purdy's replacement at Deere Labs in San Francisco, said of their emphasis on AI, "This is no different than the type of innovation that we've been doing for 180 years."

The 1830s were a vicious time in Vermont. A mania for sheep, and the consolidation of land and power required to support them, had forced upon Vermont's citizens a fragile economic monoculture. In the fall of 1836, a young Vermont blacksmith named John Deere left behind his family and his remaining property to flee his debts and possible imprisonment, setting out, with a legendary $73.73 in his pocket, for the western frontier.

At the time, farming in the prairie was severely limited by the intractability of the soil, the consistency of which was so thick that farmers went out to till their fields with cumbersome wooden paddles, pausing every few yards to scrape gobs of sticky humus from their plows' cast-iron blades. One day in 1837, Deere came across a broken mill saw with a steel blade well polished from use. Thinking it might serve as a blessedly sleek surface, he got to work fashioning it into a plow. A local farmer came into his workshop and spied the singular plow; after a short free trial, he returned to pay for the implement and bellowed, “Now get a move on you, and make me two more plows just like the other one!” Suddenly, the prairie was governable. And with that, the well-worn myth took hold: One obsessive blacksmith single-handedly opened up the West to rapid expansion.

The promise of the autonomous mechanism has always been a natural branding hook, and Deere instinctively knew to call his invention the “self-scouring” plow. The famous implement did not, of course, scour itself; it simply did not require hand-scouring. As it turns out, Deere’s genius — despite the legend — ultimately had much less to do with technological innovation than it did with his talent for marketing and sales. “There were many builders of good machinery,” writes Wayne G. Broehl Jr., in 1984’s anvil-sized “John Deere’s Company,” “and so what really counted was persuading the farmer to adopt certain practices and buy the equipment. Deere’s reputation in these early years rested on his manufactured product, but it was largely his ability to dramatize these products and get them into the hands of his customers, scattered out over the wide prairies, that made him a business success.”

This required both the establishment of new distribution channels and the extension of increasingly sophisticated financing mechanisms to cash-poor farmers. As the agricultural historian Wayne D. Rasmussen pointed out, the story of agriculture after the proliferation of Deere’s plow and other increasingly “essential” capital expenditures is a story about the shift in the presiding balance of power: “The capital needed to establish a farm increased, making it harder for laborers, tenants and young people to become operators of farms. Farmers became more dependent on bankers and merchants.”

Classical economic models of the so-called adoption curve for new technologies trace the arc of a tool — from fringe indulgence to functional necessity — on the basis of sheer productivity. The story of John Deere’s self-scouring plow, however, demonstrates that a tool’s increasing returns at scale are wont to alter the market structure itself. The very idea of a simple “adoption curve” for a new technology — by which an inventor with an obviously superior product is gradually able to convince buyers to overcome the forces of inertia — is a simplification to the point of inaccuracy, and never more so than in the case of labor-reducing tools. In other words, “pure” technological efficiency is inseparable from market prowess and, ultimately, market power.”

And the more powerful the tool, the more accelerated the process. Manufacturers had experimented with steam-tractor designs in the 19th century, but trouble with leaks, faulty boilers and soil contamination proved insurmountable. Gasoline tractors, which were prevalent enough by 1910 to show up for the first time in that year's census statistics for national horsepower, were an improvement, but when they didn't inadvertently set crops on fire they tended to break or get bogged down in the mud, and were frequently returned to the dealer as unusable. After an initial boom, the bottom fell out of the market. Manufacturers were so entranced with automation per se that they made a critical mistake: They neglected the questions of customer need, habit and expertise. As one 1916 book put it, "The human factor had been too lightly considered. The big gas tractor had been far from foolproof. Its intricacies demanded the control of experts and there were not enough experts to go around. The automobile had not yet begun to teach mechanics to hundreds of thousands of farmers. There was no permanent employment for high-priced mechanics on the farm. Just as is the case with harvest hands — when they were needed most they were scarcest."

It wasn't until 1918 that Deere moved seriously into the tractor business, with the historic acquisition of the Waterloo Gasoline Engine Company. These tractors could be used for plowing, and in limited cases for some harvesting, or for driving other pieces of machinery. "But these tasks," explains Broehl, "were only part of a farmer's spectrum of horsepower needs, for with all of the row crops — corn, hay, cotton, kafir, potatoes, tobacco, peas, peanuts, beans, sugar beets — not only the cultivation but generally the planting was done by horse. In the way tractors were constructed, it was not very easy to drive them down a row to plant or cultivate." Even two decades after their commercial introduction, gasoline tractors were neither more reliable nor more productive than horses.

The decisive factor, in the end, was less a matter of simple unit economics than exogenous shock. "World War II was the impetus," Rasmussen wrote, "for the virtually complete transition to mechanization." Farm labor was scarce and prices were high. By the time these temporary conditions subsided, the tractor's dominance had become irreversible, largely because we alter our environment to accommodate our tools as much as we alter our tools to accommodate our environment. Part of this is basic standardization — of row-spacing to suit wheel placement, for example — but there's no bright line between merely "standardizing" an existing landscape to better serve a new tool and wholly reorganizing it around a particular technology's strengths and weaknesses. Once early mechanical harvesting equipment was brought to the tomato in the early 1960s, botanists set themselves to the development of a new tomato strain that regularized the plant's maturation so the whole field could be harvested at once. This new breed also had toughened skins that machine-picking wouldn't bruise.

The many accommodations made for these new technologies, incremental as they might have felt to the actors themselves, irrevocably altered the national economy. Tractors and combines didn't just benefit from economies of scale, they necessitated them. Factor costs changed. Even if it made more sense to keep using horses — which were almost always more cost-efficient for small farms than tractors were — the increasing scarcity of breeders drove the price of horses up. Farmers with the resources and the financing to expand their holdings and purchase machines were able to survive; smaller operations were forced to fold.

And because the impatience to adopt a new technology is often driven not by a sense of how a tool should *ultimately* be used but by how it can be used right now, short-term incentives for productivity enhancements can overshadow any sense of actual need or proportion. Over the course of only a decade or two, for example, the excessive use of state-of-the-art plowing contributed to immiserating overproduction and dangerous topsoil erosion — the Dust Bowl. And as once-diversified farms gave way to single- or dual-crop production, the labor market grew increasingly contorted. While a farmer might once have been able to amortize the cost of full-time help over a range of ongoing tasks required by livestock and vegetables, a wholesale shift to row crops made labor needs seasonal. This, in turn, contributed to labor shortages — i.e., the inability to find help at a depressed wage — that could be construed as a "problem" solved by ever more mechanization. We tell ourselves that we first mechanized the jobs that were dirty, repetitive or dangerous, but another way to look at it is that we first mechanized the jobs that lent themselves to mechanization, and then we mechanized the jobs that it was all of a sudden harder to find people to do.

Looking back at earlier ages of automation, it can seem as though everything worked out just fine. After all, sheer productivity has seen explosive growth. At the beginning of the 20th century, about a third of Americans worked in agriculture; now that number is well under 1 percent and farm output has tripled since the end of World War II. And while the lack of massive structural unemployment shows that the labor market managed to adjust, we're still dealing with a profound cultural upheaval. Rasmussen quotes a 1971 column by an Iowa politician that could easily have been written by Taylor Nelson today: "The productive capacity of power machinery has greatly reduced the farm population. Occupied farming units have become fewer and fewer, and farther and farther apart, as producers with power machinery reach out for more and more land to justify their investment. Country churches, country schools, country society and small country towns have suffered. In fact, many of them have completely disappeared."

Such an assessment, it should be said, fails to consider the reality of cheap, plentiful food enabled by automation, and the relief from punishing tasks that new machinery has brought to the lives of farmers. Nobody is proposing that we return to a more primitive system; as Rasmussen put it, whatever solution we improvise for the problem of the rural unemployed, "it would hardly be humane to return them to dawn-to-dusk labor chopping cotton, thinning beets or flailing grain." But this is not to accept a purely economic analysis of technological change. From the perspective of a historian, any smooth adoption or productivity curve neglects a host of externalities, and in its single-minded emphasis on the dimension of "cost" it conceals a larger story: the collapse of a series of utterly different economic — and social — arrangements.

"It seems," Rasmussen concludes, that our agricultural bounty has been "secured at the expense of the farmer."

Thus, while the kinds of artificial intelligence Nelson has adopted are very good at solving particular problems, we cannot turn to them to help us figure out which problems — and, most important, whose problems — should rank as priorities.

The Land

Nelson called his dad and his uncle to help with the pivot/planter situation and, without wasting a minute for imprecations or counterfactuals, zoomed us across the field to his other planter. We got up into the cab of the tractor, a brand-new John Deere 9470 RT, with a John Deere 1775 ExactEmerge planter (16 rows on 30 inches per row) behind it. The planter, Nelson explained, was Deere's "latest and greatest," able to plant at twice the speed of the previous generation while maintaining near-perfect performance. Nelson revved up to about 9 m.p.h. and gestured toward one of his in-cab monitor panels, which displayed a 99.2 percent rate of "singulation," or planting accuracy. The monitor informed us that at this rate he could plant 60.6 acres per hour. In theory, with two planters — assuming the other one could be repaired quickly — two men could plant all of their corn in less than a week of 14-hour days. But there were big storms coming in over the weekend, and the areas of lower elevation could be back underwater in a few days' time. Also, he had a 2-year-old, and his wife, who manages the town's convenience store, was 9 months pregnant.

These planters, which Deere brought to market in 2015, had a variety of new features. Each row's individual piston asserted its own variable hydraulic downforce; sensors took hundreds of readings per second to ensure that each seed was planted exactly two inches deep. When we drove over a sash of flood-borne detritus, for example, the machine auto-adjusted the pressure to bore through it. All of the information the sensors were compiling — not only the location of each seed but the applied downforce necessary to plant it — was not only being shared with Nelson's other machines, but was streaming over the cloud to his iPad and his office computer for future evaluation and use. The data was also automatically shared with his partners, allowing his sales rep at the seed company to help him make finer-grained decisions about the right hybrid for next year. As we crossed an invisible line in the field, the dashboard map charting our progress changed color, indicating that we'd gotten beyond the reach of the irrigation pivot. His seed-per-acre settings switched automatically to a preprogrammed setting for non-irrigated land, one based on the script he'd written over the winter to match their field maps.

Continued on next page…

The Nelson Farm in Jackson Nebraska as the sun sets over the horizon. Lyndon French for Aventine

Nelson used the tractor's dead-reckoning system to navigate the straightaways, which freed him to spend most of his time looking backward at his progress — or sitting with his son, whom he so rarely saw this time of year. The added spatial capabilities of his new planter went much further than mere dead reckoning. Computer-vision-equipped cameras allowed for precise row-tracking, which he'd need if the field weren't perfectly straight. Once the leaves started to grow in, however, the cameras could no longer see beneath the canopy, so later in the summer a third system would kick in, a network of "feelers" that determined row location from the sensory input of vegetal brush. The real benefits of Deere's investment in machine learning, however, were most apparent in Nelson's new combine harvesters, which would use similar computer-vision systems to optimize his harvest in the fall, just as Purdy had described at Deere Labs in San Francisco. And, in the very near future, probably as soon as this coming summer, there would be innovations like Blue River's

See & Spray, which would automate precision weeding and nutrient delivery. Nelson was sparing in his use of the term AI — he much preferred the humbler and less mystical language of automation — but in these cases he was perfectly content to concede the term.

The pathways of mechanization since the self-scouring plow are littered with prophetic promotions of new dawns. Recently, we have been encouraged to believe that something called “artificial intelligence” is a special tool, one wholly and magically different from all of our previous gadgets — a change so great, in fact, that it renders our previous language inadequate. We have been hearing this for 60 years now: Early breakthroughs in the field promised to inaugurate nothing less than a new epoch in human history insofar as they secured the mechanization of cognitive tasks (chess) rather than physical ones (assembly lines). If we are still living in a recognizable world, advocates of modern machine learning suggest, it’s because it’s only very recently that AI is poised to deliver on its longstanding promise. Soon everything really will be different. But even a cursory familiarity with technological history suggests that that rupture — what, precisely, is different, and how, and for whom — is in the eye of the beholder. Still, even the skeptical are likely to allow that there’s something about recent advances in AI that feels intuitively like qualitative development.

Part of it is the data. Older generations of tractors were clearly artificially intelligent in the sense that they accomplished a human task without the same kind of human effort. But they certainly weren’t using the collection and analysis of data as the basis for the automation of something we would plausibly call a “decision.” As enthusiastic as Nelson is about the prospects for ever-increasing machine capability, he clearly has some reservations.

There was, for one thing, the ownership and security of that data. On an obvious level, he uses Deere products because he has come to trust that the company will never put anything in the field until it works, and he knows that if he needs technical support at 10 p.m. on a Friday night he can get someone on the phone. But Nelson has a sophisticated grasp of surveillance issues, and he sees that the need for trust extends well beyond technical assistance. His fields are twice as productive as they were a generation ago in large part because of the ever-increasing amount of data they can collect and exploit, but that is as dangerous as it is auspicious. “All this data could be used against us, and we've got to be concerned about data privacy. It could affect grain prices or input costs. You have to take the big companies at their word that they've got your best interests at heart. Deere has real credibility with us, which is why we allow our data to pass through their system — as opposed to some of the others who offer stuff like this." For customers like Nelson, the fears have thus largely been allayed by Deere's reputation; the company's decision not to repossess any farm equipment during the Depression has given it a purchase on generations of good will.

But there is something more, something that goes beyond what Deere as one firm can or can't do. For now, Nelson still feels in control of his own decisions, at least on an executive level. He told me that every once in a while over the course of the day he likes to get out of the cab to check on the depth and uniformity of the planting for himself. On his monitors, there are four or five different major settings that he likes to watch, and occasionally he has made ad hoc adjustments based on his visual observations. In these cases, however, he couldn't necessarily articulate a rationale for his actions. "There are so many intangibles to wrap your mind around," he said, "to give you a holistic sense of what's going on." This intuition couldn't easily be taught, at least not via explicit instruction. "When you hire somebody new, it's hard for them to understand why you're making this or that decision," Nelson said, "or even that there's a decision point there at all. All the hours I spent out here as a kid, those are irreplaceable. These decisions I make, all these formulations, they go back to just growing up around it." The tractor beeped and he adjusted something with unconscious grace. "You used to have a farm stand on every little corner of land, and they all had kids, and today you see that less and less. In my grandpa's generation, it was all just hard work — throwing bales, getting it done — and now it's all so much more complex, it takes a lot of skilled labor to run all this. The more autonomous the machines are, the easier it is to fill that void." At the same time, however, the more autonomous the machines are, the more autonomous the machines will have to be. "You automate one process," he conceded, "and there's another component that needs to be automated to go along."

Continued on next page…

Farmers working into the night during harvest. Lyndon French for Aventine

In his 1984 book “Forces of Production,” an account of the automation of machine tools, the historian David F. Noble cited an executive’s quip that you could no longer be seen in your country club if you didn’t have an automated system; if you couldn’t talk about it on the third tee, you’d feel embarrassed. While the often sensational salesmanship for the adoption of AI promises that it will add some calculated percentage to any firm’s bottom line, the longer view of the economic historian suggests that it can take a long time for the productivity savings to appear, if they do at all. Salesmen aren’t always up front, for one thing, about the full extent of the investment required to support automation — the folding hand carts in the antipodal cotton fields. And even in cases in which hindsight has made such costs clear, conscientious accountants can disagree: For decades, economists argued about the exact moment when the tractor became as cost-efficient as the horse. Such disagreements can be chalked up in part to professional methodology, but the real challenge is how to account for the consequences of path dependence. As the environment is rebuilt to suit new technological orders, certain developments are rendered inevitable and others impossible; there's no stable basis for comparison. We can only make assumptions. Are we certain that a fully automated agricultural sector is better than one that, say, encouraged the coexistence of horses and tractors? We simply have no way to know. We tend to leave these counterfactuals alone and draw the lesson that the only options are to join the technological march or to doom ourselves to irrelevance. It remains unclear whether the adoption of AI represents a choice masquerading as an imperative, or an imperative masquerading as a choice.

Any leap into the technological unknown, however, has consequences, and once you've custom-fit your environment to ensure that your technology "works," it's very hard to reverse course.

As Nelson and I made our way in even switchbacks across the field, Nelson told me that, in theory, Deere's remote-management system would allow him to put an unskilled operator in the cab and choose all of his settings from elsewhere. I asked him how his son would feel, 15 years from now, if an alert popped up on his screen to inform him that his father had changed his settings. "It's a conundrum," he said. Nelson wants the technology; he needs it. "But sitting in an office controlling a bunch of machines — that's not farming to me. I spend all this time making decisions, and the real reward is to get to go out and do it. It's the highlight of my year, in the spring and in the fall to get to be out here. Every year you get a fresh start — you've researched all these new technologies and ideas — and every year you get to go out and do a new thing. You get another chance at the cycle. It's very refreshing to me, and exciting."