

NBER WORKING PAPER SERIES

ARTIFICIAL INTELLIGENCE:
THE AMBIGUOUS LABOR MARKET IMPACT OF AUTOMATING PREDICTION

Ajay Agrawal
Joshua S. Gans
Avi Goldfarb

Working Paper 25619
<http://www.nber.org/papers/w25619>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019

Ajay Agrawal is the Geoffrey Taber Chair in Entrepreneurship and Innovation and Professor of Strategic Management, Joshua S. Gans is the Jeffrey Skoll Chair in Technical Innovation and Entrepreneurship and Professor of Strategic Management, and Avi Goldfarb is the Rotman Chair in Artificial Intelligence and Healthcare and Professor of Marketing, all at the Rotman School of Management, University of Toronto, Toronto, Canada, and the National Bureau of Economic Research. Their email addresses are ajay@agrawal.ca, joshua.gans@utoronto.ca, and agoldfarb@rotman.utoronto.ca. We thank Tim Taylor for comprehensive comments on an earlier version of this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w25619.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Ajay Agrawal, Joshua S. Gans, and Avi Goldfarb. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction
Ajay Agrawal, Joshua S. Gans, and Avi Goldfarb
NBER Working Paper No. 25619
February 2019
JEL No. J20,O33

ABSTRACT

Recent advances in artificial intelligence are primarily driven by machine learning, a prediction technology. Prediction is useful because it is an input into decision-making. In order to appreciate the impact of artificial intelligence on jobs, it is important to understand the relative roles of prediction and decision tasks. We describe and provide examples of how artificial intelligence will affect labor, emphasizing differences between when automating prediction leads to automating decisions versus enhancing decision-making by humans.

Ajay Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
ajay.agrawal@rotman.utoronto.ca

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
agoldfarb@rotman.utoronto.ca

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@gmail.com

Much of the public attention paid to artificial intelligence concerns its impact on jobs. Understanding this impact requires comprehending the capabilities of this technology. The majority of recent achievements in artificial intelligence are the result of advances in machine learning, a branch of computational statistics. Most of the concepts in standard machine learning textbooks (like Alpaydin 2010 and Hastie, Tibshirani, and Friedman 2009) are familiar to economists, like regression, maximum likelihood estimation, clustering, and nonparametric regression. Other techniques are just entering the econometrician's toolkit: regression trees, neural networks, and reinforcement learning (for discussions, see Varian 2014; Mullainathan and Spiess 2017; Athey and Imbens 2017). Over the past decade or so, advances in computer speed, data collection, data storage, and algorithms have led to substantial improvements in these techniques such that their use for commercial applications is proceeding rapidly.

Machine learning does not represent an increase in artificial *general* intelligence of the kind that could substitute machines for all aspects of human cognition, but rather one particular aspect of intelligence: *prediction* (Agrawal, Gans, and Goldfarb 2018). We define prediction in the statistical sense of using existing data to fill in missing information. As deep-learning pioneer Geoffrey Hinton (2016) said, "Take any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to make it work better than the existing techniques."

Prediction is useful because it is an input into decision-making. Prediction has no value in the absence of a decision. In this sense, each prediction task is a perfect complement to a decision task. A prediction specifies the confidence of a probability associated with an outcome under conditions of uncertainty. As an input into decision-making under uncertainty, prediction is essential to many occupations, including service industries: teachers decide how to educate students, managers decide who to recruit and reward, and janitors decide how to deal with a given mess. This wide breadth of application means that developments in artificial intelligence represent what Bresnahan and Trajtenberg (1995) called a "general purpose technology."

Prediction, however, is not the only element of a decision. Effective decision-making also requires collection and organization of data, the ability to take an action based on a decision, and the judgment to evaluate the payoffs associated with different outcomes. We characterize the decision task as distinct from the prediction task (Agrawal, Gans, and Goldfarb 2018, 2019a).

We examine four direct effects through which advances in prediction technology may affect labor in a task-based framework: 1) substituting capital for labor in prediction tasks, 2) automating decision tasks when automating prediction increases the relative returns to capital versus labor, 3) enhancing labor when automating the prediction task increases labor productivity in related decision tasks and thereby increases the relative returns to labor versus capital in those tasks, and 4) creating new decision tasks when automating prediction sufficiently reduces uncertainty as to enable new decisions that were not feasible before.

First, artificial intelligence may directly substitute capital for labor in prediction tasks. Some tasks, like demand forecasting, are already prediction tasks. Where humans currently perform these prediction tasks, they are increasingly replaced by artificial intelligence. At the same time, other tasks that were not historically viewed as prediction tasks are being transformed into prediction-oriented tasks as machine learning improves and the quality-adjusted cost of prediction decreases. For example, many parts of the workflow in human resources are being broken down into prediction tasks so that they can then be performed by artificial intelligence tools. For example, recruiting is the task of predicting, based on resumes, cover letters, LinkedIn profiles, and interview transcripts, which subset of applicants will perform best in the job. Promotion is the task of predicting which existing employees will perform best in a higher-level position. And retention is the task of predicting which star employees are most likely to leave and which of the available incentive options could most effectively be employed to encourage them to stay.

Second, when automated prediction can increase the relative returns to capital versus labor in complementary decision tasks, it can lead to the complete automation of a complementary decision task. For example, human reaction times are slower than those for machines. The returns to a machine predicting a potential car accident a few seconds or even a fraction of a second before a human would predict the accident is higher when the response time is faster. Thus, automating the prediction task increases the returns to also automating certain decision tasks associated with vehicle control. Sometimes, the artificial intelligence is able to make better predictions than a human could because it has access to different data, such as feeds from cameras, RADAR, and

LIDAR around a car.¹ Once the prediction task is automated, it increases the returns to automating some of the complementary tasks, such as those associated with vehicle control.

Third, automating the prediction task, in some cases, may have no impact on the productivity of capital performing a complementary task but may increase the productivity of labor. For example, ODS Medical developed a way of transforming brain surgery for cancer patients. Previously, a surgeon would remove a tumor and surrounding tissue based on previous imaging (say, an MRI). However, to be certain all cancerous tissue is removed, surgeons frequently end up removing more brain matter than necessary. The ODS Medical device, which resembles a connected pen-like camera, uses artificial intelligence to predict whether an area of brain tissue has cancer cells or not. Thus, while the operation is taking place, the surgeon can obtain an immediate recommendation as to whether a particular area should be removed. By predicting with more than 90% accuracy whether a cell is cancerous, the device enables the surgeon to reduce both type I errors (removing non-cancerous tissue) and type II errors (leaving cancerous tissue). The effect is to augment the labor of brain surgeons. Put simply, given a prediction, human decision-makers can in some cases make more nuanced and improved choices.

The fourth and final type of direct impact of artificial intelligence on labor happens when automated prediction sufficiently reduces uncertainty as to enable new decision tasks that did not exist before. The new tasks can be performed by capital or labor, depending on the relative costs of each (Agrawal, Gans, and Goldfarb, 2019a, 2019b). Some tasks that are not economically viable when uncertainty is high become viable as prediction technology reduces the level of uncertainty. This relates to the reinstatement force in Acemoglu and Restrepo (2019) where a freeing up of labor as a result of automation increases the returns to technologies that use labor for new tasks. At this early stage in the development and use of machine learning, there are few tangible examples of new tasks that have already arisen because of recent advances in prediction technology.

The interaction of these four forces determines the net *direct* effect of cheaper quality-adjusted predictions on labor demand. There are also *indirect* effects: As some tasks become more efficient, demand for upstream and downstream tasks might change. For example, an artificial intelligence that automates translation on an online trading platform significantly enhances international trade (Brynjolfsson, Hui, and Liu, 2018). The application of

¹For an example of artificial intelligence detecting a crash two cars ahead via RADAR, before it was humanly possible to predict, see https://www.youtube.com/watch?v=FadR7ETT_1k (accessed on October 3, 2018).

this technology not only affects translators, but also the labor involved upstream and downstream on both sides of the trade.

For individual workers, the relative importance of these forces will depend on the degree to which the core skill they bring to their job is predicated on prediction. Workers, whose core skill is something other than prediction, such as the brain surgeon described above, may find that automated prediction enhances the value of their occupation. On the other hand, workers whose core skill is prediction, such as human resource workers who screen resumes, may find the value of their occupation diminished.

In our work with the Creative Destruction Lab at the University of Toronto, looking systematically at several hundred artificial intelligence startups in the last few years, we have found that these firms often discuss how their technology will affect labor markets in specific occupations through substitution, complementarity, *and* demand expansion. We have seen very few companies building unambiguously labor-replacing technologies.

Overall, we cannot assess the net effect of artificial intelligence on labor, even in the short run. Instead, most applications of artificial intelligence have multiple forces that impact jobs, both increasing and decreasing the demand for labor. The net effect is an empirical question and will vary across applications and industries.

Automating Prediction Tasks

In this section, we describe examples that highlight substitution of human prediction by machine prediction in real-world applications. The subsequent two sections describe how such substitutions could increase either the relative returns to capital or to labor in the decision task and therefore result in automating or increasing the demand for labor in the decision task, respectively.

Prediction in Legal Services

A number of artificial intelligence applications substitute capital for labor by automating prediction tasks in legal work, while still leaving the decision tasks to the human lawyer. We describe two examples.

Kira Systems uses artificial intelligence technology to scan contracts and summarize relevant content. This may involve predicting which party in a particular lease agreement is liable for what actions or expenses or it may involve scanning all of the contracts signed by a firm to

predict which ones would be impacted if that firm were involved in a merger or acquisition. It is still up to human lawyers to make the decisions (as regulation requires), but Kira's technology predicts the relevance of clauses and information in a fraction of the time it would take a lawyer or paralegal.

In addition, artificial intelligence technology is being used to predict likely judicial outcomes based on earlier legal judgments. Blue J Legal's artificial intelligence scans tax law and decisions to provide firms with predictions of their tax liability. As one example, tax law is often ambiguous on how income should be classified (as discussed in Alarie 2018). At one extreme, if someone trades securities multiple times per day and holds securities for a short time period, then the profits are likely to be classified as business income. In contrast, if trades are rare and assets are held for decades, then profits are likely to be classified by the courts as capital gains. Currently, a lawyer who takes on a case collects the specific facts, conducts research on past judicial decisions in similar cases, and makes predictions about the case at hand. Blue J Legal uses machine learning to predict the outcome of new fact scenarios in tax and employment law cases. In addition to a prediction, the software provides a "case finder" that identifies the most relevant cases that help generate the prediction.

The end result of this process is not certainty. In the securities trading example above, the artificial intelligence predicts the likelihood of particular case facts being classified as business income or capital gains. As Blue J Legal founder Benjamin Alarie (2018) describes it, judges take the input of facts found at trial and output a judgment, using legal reasoning as a mapping function from inputs to outputs. In contrast, the Blue J Legal artificial intelligence utilizes test-facts that are assumed and entered into the system by the user, rather than the case facts found by trial. Instead of legal reasoning, the mapping function is a prediction generated by machine learning that is based on training data of past cases. Blue J Legal claims 90 percent accuracy. Given the uncertainty, a lawyer still makes the ultimate decision.

These examples show how a machine can save time and improve accuracy in generating predictions. In legal work, lawyers still make the ultimate decision. Thus, it is hard to forecast how the effect of these artificial intelligence technologies will show up in the aggregate labor statistics for legal work. They substitute for lawyers' prediction tasks but may create opportunities at the decision-task level because better prediction might affect prices and quantities in a way that increases demand for legal decision-making overall.

Prediction in Driving

The potential for mass adoption of fully autonomous vehicles generates headlines, but prediction technology is already changing driving in a number of ways that do not replace human drivers with machines.

For example, vehicle manufacturers use artificial intelligence to warn drivers about imminent risks like “there is probably a car in your blind spot” or “there is likely a pedestrian behind your car” in the form of a beep or blinking light. The machine provides the prediction, but the driver is still responsible for the decision of whether to stop, turn, or proceed.

Vehicle maintenance scheduling is another a prediction problem. Decades ago, Rust (1987) developed an empirical model of Harold Zurcher, who was the superintendent of maintenance at the Madison Wisconsin Metropolitan Bus Company. Using statistical predictions of Zurcher’s decisions, the model could be used to substitute for his predictions about when buses would break down. Today, advances in sensors and prediction algorithms have led to many new products that predict when a vehicle will break down and thus inform the decision of whether to bring a vehicle in for maintenance.

Finally, prediction is changing commercial driving by providing effective predictions of the most efficient route between two locations at any given time. Perhaps the most dramatic example is the case of London taxicabs. For decades, earning a taxi license in London meant acquiring “The Knowledge,” which involved learning the location of every address in London as well as the shortest route between any two addresses. To pass the resulting test took two to four years of study with the help of specialist training schools. But now, best-route prediction apps like Waze deliver “The Knowledge” to any driver with a smartphone, which is part of what enables ride-sharing services such as Uber to compete with London taxis. Although the skill of London cabbies did not diminish, their competitive advantage was seriously eroded by artificial intelligence.

The end result on employment is unclear. While it is surely negative for the incomes of London cabbies, overall it may be positive if more drivers (are allowed by regulators to) enter the market. This provides some insight into the types of jobs likely to be most negatively affected by artificial intelligence: jobs in which the core skill involves a prediction task.

Predictions in Email Responses

Composing an email response can be formulated as a prediction problem. Google developed Smart Reply for its email service, Gmail, using artificial intelligence to scan incoming emails and predict possible responses. Smart Reply doesn't automate sending the email response but rather predicts possible responses and provides the user with three suggestions. In 2018, within weeks of Google rolling out Smart Reply as a default setting for all of its 1.4 billion active Gmail accounts, 10% of all Gmail responses sent were generated by Smart Reply (as reported by Marcelis and MacMillan 2018). This saves the user the time of composing a response in cases where one of the three predicted replies are sufficient. However, the user must still decide whether to send a predicted response or to compose one directly.

In some cases, this kind of artificial intelligence implementation might lead to a setting where a worker must still apply judgment about the benefits and costs of a particular decision before deciding or taking an action; in others, it might automate the full decision.

To understand how drafting email might affect different types of jobs differently, we turn to the O*NET database. Sponsored by the U.S. Department of Labor through a grant to the North Carolina Department of Commerce, O*NET offers detailed descriptions of the tasks involved in almost 1,000 occupations (<https://www.onetcenter.org>).

This data includes a task described as “Prepare responses to correspondence containing routine inquiries.” The job of Executive Assistants includes this task, along with eight other occupations: Correspondence Clerks, Tellers, Receptionists and Information Clerks, License Clerks, Legal Secretaries, Insurance Policy Processing Clerks, Medical Secretaries, and Loan Interviewers and Clerks. Executive Assistants would typically draft possible responses for someone else to decide whether or not to send, and so a system like Gmail's Smart Reply fully automates the Executive Assistant's decision. In the other jobs, the worker might make use of this technology but still retain the decision task of what to ultimately send. So, in the former case the artificial intelligence replaces labor and in the latter case it enhances labor.

In this section, we provided examples in which machine prediction displaces human prediction. However, we can say little about the overall effect on jobs, which depends on the impact of better prediction on the decision tasks that they inform.

Automating Decision Tasks

Under certain conditions, automating the prediction task increases the relative returns to automating the decision task, compared to performing the decision task with human labor. In other words, when artificial intelligence is used to automate prediction, it can enhance the usefulness of implementing other technologies to automate the decision. If both the prediction and the decision are automated, it must be possible to specify the desired action to be taken for each realization of uncertainty (that is, for each realization of a prediction). For reasons of simplicity, the most common type of machine-based decision is binary – say, to reject or accept a credit application or to recommend or reject a candidate for a job interview. As artificial intelligence improves, it will provide better predictions in more complex environments.

We have started to see this type of automation in environments where machine-learning techniques are applied to mimic human decision-making. For example, a machine fitted with sensors is trained by observing the choices made by a human operator. With sufficient observations, the machine learns to predict what action a human would take given different sensory inputs. For instance, the autonomous operation of vehicles on public roads has been advanced by humans driving millions of kilometers in vehicles that are able to collect both the perception data regarding environmental conditions on the road (input) and the action data regarding the decisions made by human drivers behind the wheel (output) in response to the perception data. In many cases, the response time of an automated control system is sufficiently faster than that of a human so machines are better able to take advantage of the higher fidelity predictions generated by artificial intelligence compared to predictions generated by humans. Thus, the returns to automating the action decision (control of the vehicle) increase upon automating the prediction task. So, the application of artificial intelligence to automate the prediction task leads to automating the entire driving task resulting in a full substitution of labor for capital.

Autonomous driving gets a great deal of attention because so many people spend so much time driving. The labor-saving time from automation is therefore potentially large. However, this is also an area where removing human operation completely involves substantial risks, because the cost of failure can be so high. At present, a measure of human supervision is still required due to the probability of edge cases arising for which the machine has not been appropriately trained.

Commercial cleaning illustrates a more pedestrian attempt at automation. A&K Robotics takes existing, human-operated cleaning devices, retrofits them with sensors and a motor, and then trains a machine learning-based model using human operator data so the machine can eventually

be operated autonomously. Artificial intelligence enables prediction of the correct path for the cleaning robot to take and also can adjust for unexpected surprises that appear in that path. Given these predictions, it is possible to pre-specify what the cleaning robot should do in a wide range of predicted scenarios, and so the decisions and actions can be automated. If successful, the human operators will no longer be necessary. The company emphasizes how this will increase workplace productivity, reduce workplace injuries, and reduce costs.

Artificial intelligence also has enabled the automation of vehicles which move items from the storage part of a warehouse to the packing and shipping department. Much of this automation occurred without machine learning, by simply using dedicated tracks for delivery vehicles. However, recent applications of artificial intelligence enable swarms of robots to predict optimal routes and avoid collisions, eliminating the need for human controllers to decide on route planning. Under these conditions, warehouse vehicles can be fully automated.

Similarly, vehicle automation is growing in the mining industry, in particular for remote operations. In Australia's Pilbara region, the iron-ore mining sites are over 1000 miles from the nearest major city. Given the remoteness of the region and the extremely hot temperatures, human truck drivers are unusually expensive. Mining giant Rio Tinto initially addressed this problem by driving the trucks remotely from the offices in Perth, but in 2016 the company went a step further and deployed dozens of self-driving trucks. Artificial intelligence made this automation of the steering decision task possible by predicting hazards in the roads and by coordinating the trucks with each other. As with robot cleaning and robots in warehouses, with better prediction other technologies made automation of the decision tasks possible. Automated prediction was the last step in removing humans from the decisions involved.

Unlike autonomous vehicles on public roadways, in these controlled environments with far fewer "edge" cases, cheap prediction has already led to widespread automation of the decisions. In this way, jobs in which the core bottleneck to automation is prediction become more likely candidates for elimination.

Augmenting Labor on Decision Tasks

Discussions of artificial intelligence often envision a future of full automation. Although this may happen in some situations, in this section we discuss examples where the automation of prediction through artificial intelligence can improve decision-making by humans and

consequently the productivity of labor, specifically by allowing workers to make state-contingent decisions that reduce errors, enhancing payoffs.

Bail Decisions

Judges make decisions about whether to grant bail and thus to allow the temporary release of an accused person awaiting trial, sometimes on the condition that a sum of money is lodged to guarantee their appearance in court. Kleinberg et al. (2017) study the predictions that inform this decision:

Soon after arrest, a judge *decides* where defendants will await trial, at home or in jail. By law, this decision should be based solely on a prediction: What will the defendant do if released? Will they flee or commit a new crime? ... Currently the predictions on which these decisions are based are, in most jurisdictions, formed by some judge processing available case information in their head. ... A judge must trade off these risks [flee or commit a new crime] against the cost of incarceration. This is a consequential decision for defendants since jail spells typically last several months (or longer); recent research documents large costs of detention even over the long term. It is also costly to society: at any point in time the US has over 750,000 people in jail, disproportionately drawn from disadvantaged and minority populations.

Bail risk-prediction software will not replace people. Judges will continue to weigh the relative costs of errors, and in fact the US legal system requires human judges to decide. But artificial intelligence could enhance the productivity of judges. The main social gains here may not be in hours saved for judges as a group but rather from the improvement in prediction accuracy. Police arrest more than 10 million people per year in the United States. Based on AIs trained on a large historical dataset to predict decisions and outcomes, the authors report simulations that show enhanced prediction quality could enable crime reductions up to 24.7 percent with no change in jailing rates or jailing rate reductions up to 41.9 percent with no increase in crime rates. In other words, if judicial output were measured in a quality-adjusted way, output and hence labor productivity could rise significantly.

Emergency Medicine

Prediction technology additionally has the potential to make medicine more efficient and effective through the personalization of treatment. In particular, machine learning can identify those patients for which a given treatment will be most effective.

In the United States in particular, many prominent policymakers, economists, and medical researchers have argued that doctors test too much. For example, doctors appear to test too much

for heart attacks among patients who arrive at the emergency department of hospitals, in the sense that the average return to testing appears to be less than the average cost. Mullainathan and Obermeyer (2018, on who we draw heavily in this discussion) emphasize that prediction tools enable doctors to use the theoretically relevant object: marginal benefits and marginal costs. Using machine learning, Mullainathan and Obermeyer demonstrate that not only does the system suffer from over testing (many low-risk patients are tested), but it also suffers from *under* testing (many high-risk patients go untested). Using better prediction models when patients arrive at the emergency department could substantially increase health outcomes for the same spending (or substantially reduce spending for equal health outcomes).

Automation of the prediction of who to test does not change the workflow. The machine prediction saves only a little time for the doctors, nurses, and staff working in the emergency department, who would otherwise make the testing decision quickly, but it improves efficiency and outcomes, both reducing the number of unnecessary tests and increasing necessary ones, thus saving lives. In other words, automating the prediction task improves the productivity of emergency medicine in the context of heart attacks without much substitution for human work. Furthermore, by helping to overcome problems of under-testing, it increases demand for labor in complementary downstream tasks such as surgery, the subject of the next section.

In this section, we highlighted how automation of the prediction task can enhance labor in the decision task. The prediction tasks in these examples do not require much human labor, and so the net effect could be labor-enhancing with better and sometimes more decision tasks done by labor.

Indirect Effects: Augmenting Labor on Other Tasks

When automated prediction leads to better decisions, labor can also be augmented through tasks that are upstream or downstream of the improved decision task.

Drug Discovery

A company called Atomwise uses artificial intelligence to enhance the drug discovery process. Traditionally, identifying molecules that could most efficiently bind with proteins for a given therapeutic target was largely based on educated guesses and, given the number of potential

combinations, it was highly inefficient. Downstream experiments to test whether a molecule could be of use in a treatment often had to deal with a number of poor-quality candidate molecules.

Atomwise automates the task of predicting which molecules have the most potential for exploration. Their software classifies foundational building blocks of organic chemistry and predicts the outcomes of real-world physical experiments. This makes the decision of which molecules to test more efficient. This increased efficiency, specifically enabling lower cost and higher accuracy decisions on which molecules to test, increases the returns to the downstream lab testing procedure that is conducted by humans. As a consequence, the demand for labor to conduct such testing is likely to increase. Furthermore, higher yield due to better prediction of which chemicals might work increases the number of humans needed in the downstream tasks of bringing these chemicals to market. In other words, automated prediction in drug discovery is leading to increased use of already-existing complementary tasks, performed by humans in downstream occupations.

Language Translation

Machine language translation offers another example of machine prediction affecting a wide variety of downstream decisions. In this setting, artificial intelligence predicts how a human would translate a string of characters from one language into another. In one of the first attempts to estimate the economic impact of a commercial deployment of artificial intelligence, Brynjolfsson, Hui, and Liu (2018) measure the effect of an improvement in the quality of translation by an artificial intelligence on the volume of trade conducted on the online platform eBay. The authors find that moving to a translation using artificial intelligence resulted in a 17.5% increase in the volume of trade. This improvement in the prediction task results in a significant increase in downstream trade activity, much of which we can assume is performed by human labor. Of course, increased trade has many forms of impact on economic activity, and so we cannot draw any conclusions on the overall impact of this implementation of artificial intelligence on labor in equilibrium.

This section provided examples of the effects of improved prediction on workers in other parts of the production chain, beyond the focal prediction and decision tasks.

The Case of Radiology

Radiology offers an example of how artificial intelligence is leading to the automation of an occupation. At an artificial intelligence conference several years ago, deep learning pioneer Geoffrey Hinton (2016) publicly asserted, “We should stop training radiologists now,” comparing the profession to Wile E. Coyote from the Road Runner cartoon who has run off the cliff but hasn’t yet looked down. However, the effect of artificial intelligence on the number of workers in radiology turns out, on closer examination, to be ambiguous and nuanced. This occupation offers a useful case study that employs several of the themes developed so far.

Hinton’s remark was motivated by the progress of artificial intelligence tools that are increasingly applied to identify abnormalities in medical images. IBM and GE commercialized artificial intelligence tools that identify breast, lung, and other cancers from medical images. Smaller companies and startups have similar products. For example, Zebra Medical Vision received approval from the Food and Drug Administration to predict whether coronary heart disease is present in a CT scan. Zebra also develops tools to predict the presence of various medical issues, including bone, liver, and lung disease (as discussed at <http://www.zebra-med.com>).

A common practice is to embed image-recognition technology using artificial intelligence into the software that radiologists use to read scans. The software highlights areas predicted to be abnormalities. Radiologists examine the highlighted image when interpreting and reporting on the results. This approach uses artificial intelligence to augment the diagnosis decisions of humans rather than replace them altogether (Wang et al 2016). In these cases, a human radiologist remains in the loop for each scan, but the readings become faster and more accurate. If the number of scans stays fixed, then the demand for radiologists declines. On the other hand, if readings are faster, more accurate, and cheaper, then the number of scans could increase enough to counteract the increased number of scans read per radiologist. This scenario belongs in the earlier section where artificial intelligence automates the prediction task but not the decision.

However, some recent research suggests that machine prediction can meet or even surpass human diagnostic accuracy in detecting some types of disease (for example, Lee et al. 2017). While the current level of technology suggests a human should remain in the loop, it is plausible that over time artificial intelligence will lead to full automation of the image interpretation task. In this scenario, if the “interpret imaging results” task is done by machine, and to the extent that this task

takes up a significant fraction of the overall time, then automating this task could reduce the demand for radiologists.

But even in this situation, many tasks in the workflow of diagnostic radiologists would remain: choosing the exam, directing the technologists, reporting on the results, and deciding on an action given the probabilities reported by the machine. Many radiologists serve as the “doctor’s doctor,” communicating the meaning of images to other patient-facing doctors (Hall 2009). The interpretation of scans is often probabilistic, and radiologists have expertise in interpreting probabilities to help the patient-facing doctor recommend a course of action. Thus, reporting on the results may require a human intermediary between the machine prediction and the doctor who requested a test. For example, a human is needed to consider payoffs in order to recommend a course of action. What is the cost of conducting a biopsy if no disease is present? What is the cost of failing to conduct a biopsy if disease is present? In other words, what is the probability of disease threshold over which a biopsy (or some other further action) should be conducted? How does that vary based on patient characteristics, whether fully codifiable (such as age and medical history) or not (such as the doctor’s sense of the patient’s personality and preferences)? As the prediction task becomes better, faster, and cheaper, the demand for these related, complementary tasks may increase. In other words, it is plausible that automating the image prediction task, while reducing the demand for labor to perform that specific task, may increase the overall demand for labor due to an increased demand for complementary tasks.

In Table 1, we list the 29 different tasks that comprise the radiologists’ workflow according to the occupational classification database O*NET. Only two of these tasks are directly affected by an image recognition AI: #3 and #25. Overall, the 29 tasks reveal that even if image interpretation becomes fully automated, plenty of tasks for humans remain. The key open question is whether those tasks are best conducted by a radiologist. Perhaps some of these tasks might be better performed by medical practitioners with different expertise? For example, judgment on the best course of action for a patient might be best decided by a primary care physician or perhaps even a social worker. The supervision of radiology technologists might be better managed by more experienced radiology technologists.

<p>Table 1 Twenty-nine Tasks Associated with the Occupation of Radiologist</p>
--

1. Obtain patients' histories from electronic records, patient interviews, dictated reports, or by communicating with referring clinicians.
2. Prepare comprehensive interpretive reports of findings.
3. Perform or interpret the outcomes of diagnostic imaging procedures including magnetic resonance imaging (MRI), computer tomography (CT), positron emission tomography (PET), nuclear cardiology treadmill studies, mammography, or ultrasound.
4. Review or transmit images and information using picture archiving or communications systems.
5. Communicate examination results or diagnostic information to referring physicians, patients, or families.
6. Evaluate medical information to determine patients' risk factors, such as allergies to contrast agents, or to make decisions regarding the appropriateness of procedures.
7. Provide counseling to radiologic patients to explain the processes, risks, benefits, or alternative treatments.
8. Instruct radiologic staff in desired techniques, positions, or projections.
9. Confer with medical professionals regarding image-based diagnoses.
10. Coordinate radiological services with other medical activities.
11. Document the performance, interpretation, or outcomes of all procedures performed.
12. Establish or enforce standards for protection of patients or personnel.
13. Develop or monitor procedures to ensure adequate quality control of images.
14. Recognize or treat complications during and after procedures, including blood pressure problems, pain, oversedation, or bleeding.
15. Administer radiopaque substances by injection, orally, or as enemas to render internal structures and organs visible on x-ray films or fluoroscopic screens.
16. Participate in continuing education activities to maintain and develop expertise.
17. Participate in quality improvement activities including discussions of areas where risk of error is high.
18. Supervise and teach residents or medical students.
19. Implement protocols in areas such as drugs, resuscitation, emergencies, power failures, or infection control.
20. Schedule examinations and assign radiologic personnel.

21. Provide advice on types or quantities of radiology equipment needed to maintain facilities.
22. Participate in research projects involving radiology.
23. Perform interventional procedures such as image-guided biopsy, percutaneous transluminal angioplasty, transhepatic biliary drainage, or nephrostomy catheter placement.
24. Administer or maintain conscious sedation during and after procedures.
25. Interpret images using computer-aided detection or diagnosis systems.
26. Serve as an offsite teleradiologist for facilities that do not have on-site radiologists.
27. Develop treatment plans for radiology patients.
28. Treat malignant internal or external growths by exposure to radiation from radiographs (x-rays), high energy sources, or natural or synthetic radioisotopes.
29. Conduct physical examinations to inform decisions about appropriate procedures

Source: O*NET, <https://www.onetonline.org/link/summary/29-1069.10>.

Technology using artificial intelligence will also affect radiology in a variety of other ways, apart from predicting abnormalities in scans. For example, radiologists often dictate their reports. Past practice was that the recorded reports were sent to a (human) transcription service (as in Task #4 in Table 1). But many radiology departments already use artificial-intelligence-based transcription services to automate the transcription task. While this step can reduce costs and reduce wait times for radiologists and patients, the direct effect is the elimination of transcription-related jobs (as discussed in Thrall et al 2018), not radiologists.

The overall message here is that even when considering what may seem at first to be a clear-cut case—automation of the prediction related to reading medical image scans—the overall effects on jobs can be complex. Humans working in radiology who are not radiologists and do not work on scans—such as those providing transcription services—may have their jobs automated completely. Radiologists perform many other non-prediction tasks, and so artificial intelligence is unlikely to automate these tasks; however, it is not clear that radiologists will be the humans who perform these tasks if reading scans becomes automated. It is ultimately not obvious even whether the number of radiologists will rise or fall, since that will depend on whether radiologists perform the non-prediction tasks and whether overall demand for radiology services rises as radiology becomes more efficient.

New Tasks Through New Decisions

As artificial intelligence improves prediction, it may allow for new decisions to be made where previously it was impossible or too costly to do so. As Herbert Simon (1972) emphasized, when rationality is bounded – for example, in terms of being able to distinguish adequately between important outcomes in a complex environment – economic agents will instead resort to rules. Those rules can take various forms: they may be followed by individuals, be part of operational procedures in companies, or be embedded in machines. However, when uncertainty is reduced, generic rules may be replaced with probability-driven decisions. This is important because state-contingent choices can be consequential for companies. For example, Google deployed artificial intelligence developed by its DeepMind unit to optimize the use of air conditioners in its data centers. The artificial intelligence enables new decisions on energy usage. The end result was a 40% reduction in energy used in a highly energy intensive operation (Evans and Gao 2016).

New tasks may also be performed by humans. As highlighted above in the context of drug discovery, artificial intelligence is already having an impact on scientific research. Uncertainty is pervasive in many aspects of research, and so prediction technology is likely to have a large effect on the production of science. Cockburn, Henderson, and Stern (2019) show that machine learning is used by scientists in a wide variety of fields. Figure 1 shows the number of publications in computer science and applications journals by artificial intelligence field. The blue, dark red, and dark green lines show publications in three different artificial intelligence subfields of computer science: machine learning, robotics, and symbolic logic. The results show a slow and steady increase in publications in all three, with the largest increase in machine learning. The most striking result in the figure, however, is the yellow line. It shows the increase in publications outside of computer science that mention machine learning. In other words, it demonstrates that, since 2012, the biggest change in artificial intelligence publications did not occur in computer science. It occurred in other fields of science that use machine learning. The same is not true for robotics and symbolic logic (light blue and pink lines). This use of machine learning as an input into innovation is apparent in our own field of economics, in which researchers are increasingly using machine learning to improve our statistical models and advance our knowledge of economics (Athey 2019).




Figure 1: Publications in Computer Science versus Application Journals, by artificial intelligence Field
 Source: Cockburn, Henderson, and Stern (2019, Figure 4)

In this way, the recent advances in machine learning can be seen as an invention in the method of inventing, as highlighted in Griliches (1957) for hybrid corn. Cockburn, Henderson, and Stern (2019) describe this insight: “The challenge presented by advances in artificial intelligence is that they appear to be research tools that not only have the potential to change the method of innovation itself, but also have implications across a wide range of fields.” Agrawal, McHale, and Oettl explain how artificial intelligence may influence the knowledge production function (2019a) and model the implications of using artificial intelligence to produce a map of the complex combinatorial search space of ideas for the purpose of reducing the cost of predicting which combinations of ideas offer the greatest promise (2019b). In the context of Atomwise, new tasks may arise if drug discovery becomes more efficient and drugs can be better-targeted to narrower populations. In other words, in addition to increasing demand for existing tasks, artificial intelligence is likely to create innovations that lead to new industries and new types of jobs with new tasks in those industries.

At this early stage in the diffusion of machine learning technology, examples of new tasks created by artificial intelligence are scarce and speculative. However, we provide three examples that apply artificial intelligence in unique areas that suggest the potential for new industries, jobs, and tasks.

At the University of Toronto, Alan Aspuru-Guzik's research group is developing a "self-driving chemistry lab" that enables the discovery of chemicals and materials at a fraction of the price of a current lab. Using advances in robotics and machine learning, the lab could be deployed in thousands of locations around the world, without the need for a local workforce with deep expertise in chemistry. This would enable industries in rural areas and developing countries to have access to a wide variety of materials. The inventors emphasize that this tool could "provide the scientific community with an easy-to-use package to facilitate novel discovery at a faster pace" (Roch et al. 2018, p. 1) and "democratize autonomous discovery" (p. 12). One can imagine many new tasks associated with the arrival of an autonomous chemistry lab with the capability of on-site discovery.

The commercialization of space offers another example of how machine learning could generate a new industry at a commercial scale. Uncertainty is a key challenge to operating assets in space. For example, the risk of destruction from debris is a well-established deterrent for deploying commercial satellites (Liou and Johnson 2006). The company Seer Tracking has built an artificial intelligence to predict the trajectory of space debris. Most directly, this could create a set of (human and machine) tasks focused on moving space assets out of the way of incoming debris. Perhaps more importantly, it could enable more commercial opportunities in space by reducing the risk that a space asset will be destroyed by debris. In other words, the uncertainty associated with space debris may mean that some decision tasks are never undertaken. Resolving this uncertainty could enable new commercial opportunities in space.

Another example is the management of chronic disease. One such disease, diabetes, leads to hundreds of thousands of deaths annually. Key to preventing hospital admissions and severe complications is the control of blood glucose levels; however, many diabetes patients have difficulty maintaining a relatively safe level of glucose control. Better prediction of current glucose levels could substantially reduce complications by enabling response (Ismail 2017). This improved prediction could generate a new set of tasks that benefit from lower cost and more accurate monitoring. Potential new tasks in managing chronic diseases like diabetes include managing the sensors, interpreting the data, and real-time advising on dietary and exercise habits. Inherent in these new tasks is the ability to tailor medical decisions and treatment to individual patients

(Contreras and Vehi 2018). Improved management of chronic disease could arise because of the new individual-level decisions enabled by better prediction.

These examples are speculative. The technology is too early and the diffusion is too limited to offer definitive examples of new tasks arising from recently automated predictions. At the time of this writing, the most likely consequences of artificial intelligence on labor come from *existing* tasks that are affected from better, faster, and cheaper prediction. We already observe real examples of the reduction of labor due to automating existing prediction tasks, and we also see examples of increased demand for labor due to enhanced demand for certain existing tasks that are complements to prediction. Our broader theme is that uncertainty can render certain activities economically infeasible and so reduced uncertainty can enable new opportunities and new tasks to be implemented by some mixture of capital and labor.

Conclusion

Our contribution to the task-based model of technology and labor (as discussed in Autor and Acemoglu 2011; Acemoglu and Restrepo 2019) is to highlight the usefulness of thinking in terms of prediction tasks and decision tasks, where decision tasks are perfect complements to prediction tasks, in the sense that prediction has no value without a decision. This structure describes how artificial intelligence directly substitutes capital for labor in the case of prediction tasks and may indirectly effect decision tasks by increasing or decreasing the relative returns to labor versus capital for decision tasks. It may also lead to increases in labor tasks upstream or downstream.

For any given worker, a key predictor of whether artificial intelligence will substitute for their job is the degree to which the core skill they bring to the job involves prediction. Transcription jobs are being automated as the core skill of that labor is predicting which words to type upon hearing a recording. For London taxi drivers, when artificial intelligence was employed to predict the optimal route through the City's streets, their jobs were put at risk (though other drivers' labor became augmented).

Artificial intelligence does not fit easily into existing analyses of the effect of automation on labor markets. The reasons are threefold. First, prediction is always strictly complementary to other tasks – namely decision-related tasks. Those tasks can be existing or newly possible because of better prediction. Second, better prediction improves decisions—whether taken by labor or

capital—by enabling more nuanced decisions through the reduction of uncertainty. Finally, it is not yet possible to say whether the net impact on decision tasks – whether existing or new – is likely to favor labor or capital. We have found important examples of both, and there is no obvious reason for a particular bias to emerge. Thus, we caution on drawing broad inferences from the research on factory automation (for example, Acemoglu and Restrepo 2017; Autor and Salomons 2018) in forecasting the net near-term consequences of artificial intelligence for labor markets.

References

- Acemoglu, Daron, and David Autor. 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. Chapter 12. In *Handbook of Labor Economics*. Eds. David Card and Orley Ashenfelter, Elsevier, Volume 4, Part B, 1043-1171,
- Acemoglu, Daron, and Pascual Restrepo. 2017. Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper 23285.
- Acemoglu, Daron, and Pascual Restrepo. 2019. Artificial Intelligence, Automation and Work. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Acemoglu, Daron, and Pascual Restrepo. 2019. Automation and New Tasks: The Implications of the Task Content of Technology for Labor Demand. *Journal of Economic Perspectives* (forthcoming).
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019a. Prediction Judgment and Complexity: A Theory of Decision Making and Artificial Intelligence. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019b. Dumb machines, smart machines, and humans: The impact of artificial intelligence on task allocation. Working paper, University of Toronto.
- Agrawal, Ajay, John McHale, and Alexander Oettl. 2019a. Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Agrawal, Ajay, John McHale, and Alexander Oettl. 2019b. A Multi-task Combinatorial Model of AI-Aided Scientific Discovery and Innovation. Working paper, University of Toronto.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015 The Skill Complementarity of Broadband Internet. *Quarterly Journal of Economics* 130(4), 1781–1824.
- Alarie, Benjamin. 2018. “Predicting Judicial Decision-Making with Machine Learning.” *Blue J Legal*, February 28, <https://www.youtube.com/watch?v=-pqbQejGVpc&submissionGuid=c9852990-2f9e-46e6-8cd2-b7507bebcfe9>. Accessed February 20, 2019.

- Alpaydin, Ethem. 2010. *Introduction to Machine Learning*. Second Edition. MIT Press, Cambridge MA.
- Athey, Susan. 2019. The impact of machine learning on economics. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Autor, David, and Ana Salomons. 2018. Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*.
- Bresnahan, T. and M. Trajtenberg. 1995. General Purpose Technologies ‘Engines of Growth’? *Journal of Econometrics*. 65. 83-108.
- Brynjolfsson, Erik, Xiang Hui, and Meng Liu. 2018. Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform. NBER Working Paper 24917.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2019. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Cockburn, Iain, Rebecca Henderson, and Scott Stern. 2019. The Impact of Artificial Intelligence on Innovation. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Contreras, Ivan, and Josep Vehi. 2018. Artificial Intelligence for Diabetes Management and Decision Support: Literature Review. *Journal of Medical Research* 20(5): e10775.
- Evans, Richard, and Jim Gao. 2016. “DeepMind AI Reduces Google Data Centre Cooling Bill by 40%.” DeepMind blog, July 20, <https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>.
- Griliches, Z. 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica*. 25(4). 501-522.
- Hall, Ferris M. 2009. “The Radiology Report of the Future.” *Radiology*, 251: 2 (May), 313-316,
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data mining, Inference, and Prediction*. Second Edition. Springer, New York.

- Hinton, Geoff. 2016. "On Radiology." Creative Destruction Lab: Machine Learning and the Market for Intelligence," <https://www.youtube.com/watch?v=2HMPRXstSvQ>. Accessed February 20, 2019.
- Ismail, Nick. 2017. "How AI can improve patient outcomes for type 1 and 2 diabetes." *Information Age*. <https://www.information-age.com/ai-can-improve-patient-outcomes-type-1-2-diabetes-123467505/>. July 25. Accessed February 20, 2019.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. Human Decisions and Machine Predictions, *Quarterly Journal of Economics*, 133(1), 237–293.
- Lee, Hyunkwang, Shahein Tajmir, Jenny Lee, Maurice Zissen, Bethel Ayele Yesiwas, Tarik K. Alkasab, Garry Choy, and Synho Do. 2017. "Fully Automated Deep Learning System for Bone Age Assessment." *Journal of Digital Imaging*, 30:4 (August), 427-441.
- Liou, J.-C., and N.L. Johnson. 2006. Risks in Space from Orbiting Debris. *Science*, 311, 340-341.
- Marcelis, David, and Douglas MacMillan. 2018. "Is This Article Worth Reading? Gmail's Suggested Reply: 'Haha, Thanks!'" *Wall Street Journal*, September 18, <https://www.wsj.com/articles/very-interesting-awesome-love-it-gmail-users-confront-chipper-smart-reply-1537282569>. Accessed February 20, 2019.
- Mullainathan and Obermeyer. 2018. Using Machine Learning to Understand Human Decision-Making: Application to Health Care. Working paper, University of Chicago.
- Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives*, 31 (2): 87-106.
- Roch, Loic, Florian Hase, Christoph Kreisbeck, Teresa Tomayo-Mendoza, Lars Yunker, Jason Hein, and Alan Aspuru-Guzik. 2018. ChemOS: An orchestration software to democratize autonomous discover. ChemRxiv. https://chemrxiv.org/articles/ChemOS_An_Orchestration_Software_to_Democratize_Autonomous_Discovery/5953606.
- Rust, J. 1987. Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 55(5), 999-1033.
- Simon, Herbert. 1972. Theories of Bounded Rationality. In McGuire and Radner Eds. *Decision and Organization*. Chapter 8. North-Holland, Amsterdam.
- Thrall, James H., Xiang Li, Quanzheng Li, Cinthia Cruz, Synho Do, Keith Dreyer, and James Brink. 2018. "Artificial Intelligence and Machine Learning in Radiology: Opportunities,

Challenges, Pitfalls, and Criteria for Success. *Journal of the American College of Radiology*, 15:3 (Part B, March), 504–508.

Varian, Hal R. 2014. “Big Data: New Tricks for Econometrics.” *Journal of Economic Perspectives*, 28 (2): 3-28.

Wang, D., A. Khosla, R. Gargeya, H. Irshad, and A. H. Beck. 2016. Deep Learning for Identifying Metastatic Breast Cancer. arXiv preprint arXiv:1606.05718.