Oxford Review of Economic Policy, Volume 34, Number 3, 2018, pp. 393–417

Computers and populism: artificial intelligence, jobs, and politics in the near term

Frank Levy*

Abstract: I project the near-term future of work to ask whether job losses induced by artificial intelligence will increase the appeal of populist politics. The paper first explains how computers and machine learning automate workplace tasks. Automated tasks help to both create and eliminate jobs and I show why job elimination centres in blue-collar and clerical work—impacts similar to those of manufactured imports and offshored services. I sketch the near-term evolution of three technologies aimed at blue-collar and clerical occupations: autonomous long-distance trucks, automated customer service responses, and industrial robotics. I estimate that in the next 5–7 years, the jobs lost to each of these technologies will be modest but visible. I then outline the structure of populist politics. Populist surges are rare but a populist candidate who pits 'the people' (truck drivers, call centre operators, factory operatives) against 'the elite' (software developers, etc.) will be mining many of the US regional and education fault lines that were part of the 2016 presidential election.

Keywords: populism, artificial intelligence, computers, future of work

JEL classification: J23, J24, M51, O33

I. Introduction

In this article, I start to explore how artificial intelligence (AI) will change the economy in the next 5–7 years. At first glance, the short horizon is small beer: many articles now predict how AI will change the economy in a decade or two (Frey and Osborne, 2013). I believe these long-run predictions suffer from a common weakness. As Giorgio Presidente (2017) writes:

*MIT, Harvard Medical School, and Duke Robotics; e-mail: flevy@mit.edu

Thanks go to Ryan Hill and Eric Huntly for research assistance, the MIT CSAIL AI-PI Friday lunch for extensive conversation, and Daron Acemoglu, David Autor, Roy Bahat, Gordon Berlin, Jack Citrin, Henrik Christensen, Ron Friedmann, Jed Kolko, John Leonard, Becca Levy, Mac McCorkle, John Markoff, Eric Patashnik, Giorgio Presidente, Andrea Salvatori, Aaron Smith, Kathy Swartz, Moshe Vardi, Fred Yang, and several software developers and call centre managers who requested anonymity. This work was supported by a grant from the Russell Sage Foundation.
doi:10.1093/oxrep/gry004

> The current debate on 'the future of work' or 'jobs at risk of automation' seems to implicitly adopt a pure science-push view, which assumes a path for technology driven by what science makes achievable, rather than what is needed by firms.

The science-push view also has no role for institutions, politics, or policy and so it risks oversimplified conclusions. Over time, some technologies will deploy faster than others and some occupations will be disrupted faster than others. The sequence and speed of developments and people's reactions to the developments will jointly determine how the economy evolves. A description of international trade's impact on the US economy would be misleading if it omitted trade's role in reviving US populism, an important force in the 2016 presidential election (Autor *et al.*, 2017).

I develop the argument in four parts. In section II, I give a basic explanation of how AI replaces, modifies, and creates and replaces jobs, with an emphasis on the role of machine learning. In section III, I apply this theory to today's economy. Using four examples of existing jobs, I show how current AI is helping to slowly polarize the occupational structure, displacing people from blue-collar and working-class jobs into lower-wage work—the same displacement caused by manufactured goods and off-shored services. In section IV, I project the likely near-term job losses from three 'hot' AI applications: autonomous trucks, automated customer service responses, and industrial robotics. In section V, I discuss why these near-term job losses may or may not increase the appeal of populist politics. Section VI concludes.

II. How artificial intelligence changes human work[1]

To understand how AI disrupts the job market, note first that computers often automate part of a job rather than an entire job—take as an example an automated teller machine (ATM) and a bank teller's job. For this reason, it is useful to think of a job as a set of tasks (Autor *et al.*, 2003). Our focus will be on how AI automates a task: how AI uses digital technology to achieve the end result of a task, though not necessarily as a human would achieve it. An e-mailed message transmits text very differently from a postman delivering a letter.

The theory of task automation begins with two observations:

- all human work involves the processing of information. A financial analyst reading a report, a chef tasting a sauce, a farmer looking to the sky for signs of rain: each is an example of processing information to understand what to do next or to update a picture of the world;
- a computer processes information by executing instructions.

It follows that for AI to automate a task, it must be possible to model the required information processing by applying a set of instructions. To perform the task without error, the instructions must specify an action for every possible contingency (though we will see that, in many cases, this high bar cannot be met).

[1] Some parts of this section draw on Remus and Levy (2017).

Software models to automate tasks are built using two kinds of instructions—*deductive* instructions and *data-driven* instructions. Deductive instructions, sometimes called *rules*, are used when we can articulate the information-processing structure. An example is the self-service airline check-in kiosk that processes information from a credit card and the airline's reservation database into a boarding pass. A simplified set of deductive instructions might read, in part:

- read the name on the credit card;
- check whether the name on the credit card matches a name in the reservation database:
  - if yes, check that the customer has a seat assignment,
  - if no, instruct the customer to see desk agent.

Note that the software can handle all contingencies because it has the option of referring a customer to a human desk agent. Without that option, an unanticipated contingency would cause the software to grind to a halt.

Data-driven instructions are used when we are not conscious of the information-processing structure—for example, the visual information processing by which a driver sees and makes sense of a traffic light. In some cases, it is possible to approximate unconscious information processing by estimating a statistical model that directly relates the information output to the information inputs with no attempt to model the intervening steps. Data-driven instructions are the estimated equations of such a statistical model.

Consider an information-processing problem that is of interest to lawyers: the mental process of a particular judge in reaching a verdict in a non-jury case. A lawyer who understands a judge's mental process may be able to predict whether the judge will rule for the plaintiff or the defendant in, say, an upcoming medical malpractice case. In a statistical model of the judge's mental process, the information inputs include the facts of the case and the elements of the cause of action. The information output is the judge's verdict. The judge's decision process may be opaque but it can be approximated by a statistical (linear regression) model that is estimated using a set of the judge's prior verdicts in similar cases. The model can be sketched as follows:

$$Y_i = \beta_1 X_{1i} + \beta_2 X_{2i} \ldots\ldots\ldots + \mu_i \tag{1}$$

where: $Y_i = 1$ if the judge decides in favour of the plaintiff in the ith case; $Y_i = 0$ if the judge decides in favour of the defendant in the ith prior case; X_{1i}, X_{2i}…. are case characteristics drawn from the record of the ith prior case, including the facts of the case and elements of the cause of action; and $\beta_{1,}$ β_2…. are the estimated coefficients of the case characteristics; μ_i is a stochastic error term for the ith judicial decision.

In this estimation, the judge's prior cases are called the training sample and the estimation process is called training or 'supervised (machine) learning'—supervised because the estimated parameters are forced to align as much as possible with the judge's prior verdicts; learning because the estimation process can be seen as learning the relationship (summarized in βs) between the case characteristics and the judge's verdicts.[2] Once estimated, equation (1) becomes a data-driven instruction that can be

[2] The estimation process is also described as pattern recognition as the algorithm searches for the pattern of case characteristics that best predict the judge's decision.

applied to characteristics of an upcoming case to estimate the *ex ante* probabilities that the judge decides for the plaintiff or for the defendant.

The model in equation (1) uses a linear regression for ease of exposition. Linear regressions sharply restrict the mathematical form of relationships between case characteristics and the judge's verdict. For this reason, a researcher might use a more complex statistical estimator—a probit, a neural network—to capture non-linear relationships including threshold values and complex interactions among the case characteristics. But the underlying idea remains unchanged: estimate a model that uses characteristics of prior cases to predict the judge's verdict.

Note the word 'predict' in the last sentence. While the airport kiosk creates a boarding pass with certainty, the machine-learning model creates a prediction of the judge's decision with the possibility of error (Agrawal *et al.*, 2016).

Machine-learning predictions lie at the heart of other aspects of AI, including computer vision. Computer vision refers to a computer's ability to scan, for example, the digital image in Figure 1 and identify it as a kitten as opposed to a puppy, a small child, a bicycle, a Pontiac, or some other object.

From a machine-learning perspective, the image of the kitten is a collection of data. The ability to analyse these data rests on the fact that the image is digitized. Viewed at the level of pixels as in Figure 2, the digital image has many specific features—edges where adjacent pixels differ sharply in their colour or intensity, corners where two edges meet, and so on. Roughly speaking, these features play the role of the case characteristics (the Xs) in equation 1.[3]

In modelling the judge's decision process, there are two outputs to consider—a decision for the plaintiff and a decision for the defendant. In modelling vision, an image might represent any of thousands of different objects. Nonetheless, both models use a similar predictive logic. In the example of the judge, the statistical model is trained (estimated) using the judge's past cases. In the vision example, the statistical model is

Figure 1: An image of a kitten to be classified by computer vision

[3] This description would have been accurate 5 years ago. Today, neural net models move directly from the pixels of the digital image to estimation without the user explicitly identifying edges, corners, and other features.

Figure 2: An enlarged section of the kitten image to be classified

trained using a large collection of images of various objects. In the example of the judge, the outputs are the estimated probabilities that the judge decides for the plaintiff and for the defendant. In the vision example, the outputs are the estimated probabilities that the image is, respectively, a kitten, a puppy, a small child, a bicycle, a Pontiac, a refrigerator, and so on. In the example of the judge, the statistical procedure estimates the model's coefficients (the βs) to maximize the probability of correctly predicting the judge's past verdicts. In the vision example, the statistical procedure[4] estimates the model's coefficients to maximize the probability that the image is identified as a kitten.

Machine-learning prediction models similar in spirit (but not detail) are used to recognize spoken words, to predict meaning from spoken or written words, to predict whether a particular credit card transaction is fraudulent, to predict which persons in a call centre database are most likely to make a purchase, to predict an individual's disease based on an individual's symptoms and medical history, and so on.

Many of the algorithms used to estimate these models—e.g. neural networks—were well into development in the 1980s but they required what were then impractically large computational resources. What made the algorithms practical were big gains in computer power and the development of large, digitized data sets that together allow the design and training of highly refined models.[5]

A caveat to this progress is the way that a model's complexity can obscure why it arrives at a particular prediction. Many models can estimate both a best prediction and the probability that the 'best' prediction is correct: 'the judge will decide this case in favour of the defendant with probability equal to 0.73'. Some models can also display the proximate statistical factors that drive the model's prediction, but listing these factors usually falls short of an overall logic.

[4] In the vision case, the statistical procedure is quite complicated and will typically involve multiple iterations as the estimation systematically adjusts the coefficients (e.g. backwards propagation) to improve the model's predictive ability.

[5] For a brief historical discussion of neural networks, see the Stanford University website 'Neural Networks': https://cs.stanford.edu/people/eroberts/courses/soco/projects/neural-networks/index.html

Caveat aside, today's machine-learning models have significant implications for the labour market. Recall the condition for automating a task described earlier: for a computer to automate a task, it must be possible to model the required information processing using a set of instructions.

If modelling were limited to deductive instructions—modelling tasks where we can articulate the information-processing structure—automation would have a relatively small reach. Machine learning allows computers to model tasks where part or all of the information processing is unconscious and so opens a much wider set of tasks to potential automation.

A perspective on this development comes from the writing of scientist and philosopher Michael Polanyi. A half-century ago, Polanyi wrote: '[W]e know more than we can tell' (Polanyi, 1967), an idea that became known as Polanyi's paradox. As an example, we can know how to ride a bicycle but we can't explain to a child how to ride a bicycle in a way that keeps her from falling as she learns. By discovering information-processing instructions that we cannot articulate, machine learning allows us to unravel at least a part of Polanyi's paradox.

Despite this progress, there are, at least for the present, limits to the kinds of tasks that can be automated. To understand these limits, let us return to the problem of modelling the judge's decision-making process, in particular the characteristics of that problem that made modelling feasible.

At the outset, the judge's decision-making process has a constant 'structure'. Consider our reasoning. We assumed it was possible to model the judge's information processing by applying an unchanging set of instructions to case characteristics. This means that if the judge had, in the past, decided five cases with identical characteristics, he must have reached the same verdict in each case. If the judge had reached different verdicts in some of the identical cases, the model (equation 1) would not have been a good statistical fit for the data in the judge's past cases and it would do a poor job predicting the judge's future verdicts (as befits an unpredictable judge).[6]

As a second limitation, the estimated model could only predict future verdicts in cases that were generally similar to the training sample of past cases on which the model was estimated. For example, if the judge's past cases all involved female plaintiffs, the model might not accurately predict the judge's decisions in a future case with a male plaintiff.

This limit on predictions is part of a general problem in which machine-learning models, like all statistical models, have potential difficulties in predicting outcomes for cases that lie 'outside' the data on which they were estimated. For this reason, the development of autonomous vehicles is being slowed by the need to collect training data on almost any situation a vehicle might face.

In the absence of complete training data, a machine-learning model has several options. One option is to fudge: the model gives its 'best' answer without warning the user that its best answer may not be very good. On 9 March 2013, an iPhone Siri was asked: 'Can a dog jump over a house?' Many 4-year-olds can answer that question but Siri had not been trained on the question. It produced a set of telephone directory listings and said, 'Ok. One of these kennels looks fairly close to you.'

[6] Autor *et al.* (2003) define a 'routine' task as one that can be modelled using only deductive rules. By this definition, many repetitive tasks are not 'routine'. For example, classifying movie reviews as positive or negative is not routine because it requires a machine-learning model to interpret (process) the written language.

A different option is to continue to train the model on the job—'reinforcement learning'. Consider the Siri example above. Beneath the user interface, Siri's software has retrieved a list of candidate answers to the dog/house question. Based on prior training, the software estimated a probability that each candidate answer was the correct answer and the software displayed the answer with the highest estimated probability.[7] If the Siri user could have rated the 'kennels' answer as unsatisfactory, the negative rating would have been a signal to adjust (retrain) the algorithm that selects the 'best' answer. Note, however, that this on-the-job training requires that the user knows the correct answer.[8] It also assumes that a software error on the job will not result in catastrophe—an incorrect reading of a red light by an autonomous school bus.

As the Siri example suggests, computers cannot yet participate in sustained, unstructured human interaction. Such interaction often depends on formulating responses to unanticipated questions and statements. This, in turn, requires recognizing the broader context in which words are being used—not only the surrounding words, but the identity and motivation of the speaker and the purpose of the communication—hard information to ascertain.[9]

The history of AI includes many problems that were solved much more quickly than predicted. At this point, however, the best guess is that AI is likely to make most near-term progress in automating narrow, structured tasks—structured because machine learning is predicated on identifying an underlying pattern of information processing; narrow because it is desirable to estimate the model on a training sample that contains most of the situations the software is likely to confront.

III. Applying the theory: which jobs are vulnerable today?

In thinking about the next 5–7 years, it is important to remember that AI and digitization create, as well as destroy, jobs. Virtually all computer applications involve jobs in programming, data science, and related fields. Networks, robotics, and other physical carriers of technology create jobs for equipment installation and maintenance. Most important are job-producing products and services that would not exist without AI: the driver who works for Amazon's 2-day delivery service, the manufacturer of specialized hunting gear who reaches a large market over the internet, a pharmaceutical geneticist whose work would hardly exist without computerized gene sequencing.

But if near-term AI stimulates a political reaction, the reaction will likely begin with job losses rather than new jobs. We begin to characterize the jobs lost to AI by examining four current occupations. The details that follow come from the Bureau of Labor Statistics *Occupational Outlook Handbook*[10] with detailed citations omitted for brevity.

[7] In all likelihood, the kennel answer did not have a high probability of being correct—just a higher probability than the alternatives. Because Siri does not report the probabilities, the user did not know this.

[8] Alternatively, the model may be performing a task where the software itself can determine whether the performance was successful. The software that beat grandmasters at the game Go trained by playing thousands of games against itself and adjusting strategies depending on whether they were successful.

[9] Examples of the state of unstructured computer conversation appear in the 2016 Loebner Prize Competition: http://www.loebner.net/Prizef/loebner-prize.html.

[10] https://www.bls.gov/ooh/

(i) Bank tellers (2016 median pay: $27,260)

The Bureau of Labor Statistics (BLS) describes a teller's job as follows: 'Tellers are responsible for accurately processing routine transactions at a bank. These transactions include cashing checks, depositing money, and collecting loan payments.'

To which they might have added: marketing the bank's certificates of deposit, student loans, and financial advising services; helping customers access safe deposit boxes; etc. Cashing cheques and accepting deposits are structured tasks, but automation was delayed until software could process the required information inputs: recognizing a security code that identified the person making the transaction and reading the dollar amount on a cheque or a piece of currency. By 1969, the cash-dispensing task had been solved with a magnetically coded plastic card to input the security code. Sometime later, improved character-recognition software allowed ATMs to accept cheque and paper money deposits.

Because ATMs were automating part of a teller's job, it was reasonable to predict that banks would only need tellers for the unstructured parts of the job and so fewer people would work as tellers. The prediction had some validity: the ATM reduced the number of tellers required for a typical urban bank from 21 to 13 (Bessen, 2015). But as ATMs spread through the US economy the number of tellers continued to grow for a time because the number of branch banks grew. The reduced number of tellers made it cheaper to operate a branch bank and, driven by a combination of competitive strategy and deregulation, banks opened more branches.[11]

(ii) Medical transcriptionists (2016 median pay: $35,720)

The BLS description of a medical transcriptionist's job is:

> Medical transcriptionists, sometimes referred to as healthcare documentation specialists, listen to voice recordings that physicians and other healthcare workers make and convert them into written reports. They may also review and edit medical documents created using speech recognition technology. Transcriptionists interpret medical terminology and abbreviations in preparing patients' medical histories, discharge summaries, and other documents.

The job description does not mention 'structured', but the work clearly is structured with a goal of always interpreting 'medical terminology and abbreviations' in the same way.

Because the work is highly structured, a growing fraction of medical transcription is already computerized, a fact noted in the job description's 'review and edit medical documents created using speech recognition technology'. Speech recognition does not solve everything: a human still has to remove the 'ahs' and 'ums'. But in the case of radiology, for example, the transcribed speech appears at the bottom of radiologist's screen allowing the radiologist to correct her own report with no need for outside editors.[12]

[11] The 'Jevons effect' describes a sequence where technical change reduces the need for, say, production workers per car, but demand for production workers increases because of increased demand for the now cheaper cars.

[12] The effect is similar to the way personal computers allowed people to type their own documents and so reduced the demand for secretaries.

Because the software is transcribing as the report is being dictated, some versions of the software can search the text for key words indicating a critical medical problem and flag the problem for the radiologist and the patient's physician.[13] Ultimately, transcription software should be able to translate the radiologist's report into current procedural terminology (CPT) codes that are sent to the insurer for reimbursement. Extracting CPT codes from a report would automate part of another occupation—the medical records and health information technician (2016 median pay: $38,040).

Where the ATM automates a part of what bank tellers do, transcription software will eventually automate virtually all of what a medical transcriptionist does. Even if cheaper transcription increases the demand for transcription, the demand for human transcriptionists is unlikely to increase.

(iii) Janitors and building cleaners (median annual pay: $24,190)

The BLS description of the job of janitor/building cleaner is as follows:

> **What Janitors and Building Cleaners Do**
> Janitors and building cleaners keep many types of buildings clean, orderly, and in good condition.
>
> **Work Environment**
> Most janitors and building cleaners work indoors. However, some work outdoors part of the time, sweeping walkways, mowing lawns, and removing snow. Because office buildings often are cleaned while they are empty, many cleaners work evening hours. The work can be physically demanding and sometimes dirty and unpleasant.

Janitorial work is often described as simple because most people can meet its cognitive requirements. Nonetheless, the job is extremely difficult to automate—one of a set of occupations whose tasks are easy for humans and very hard for computers (Autor *et al.*, 2003, Table 1). The janitor's job and many other low-wage jobs—e.g. home health aide—require unstructured conversation and extensive, unstructured physical movement. Structured (repetitive) physical movement has long been automated—witness the success of manufacturing robots—but unstructured physical movement, like walking across an unanticipated, uneven surface, represents a challenge (IEEE Spectrum, 2015). In the near term, automation is unlikely to affect the janitor's job.

(iv) Lawyer (2016 median annual pay: $118,160)

The BLS describes the job of a lawyer as follows: 'Lawyers advise and represent individuals, businesses, and government agencies on legal issues and disputes.'

[13] See for example, http://info.nuance.com/powerscribe360-reporting-demo?gclid=EAIaIQobChMIsZmVi9j81QIVC1gNCh1sEA2WEAAYASAAEgI_6vD_BwE

In 2011, a *New York Times* article attracted substantial attention with the headline 'Armies of Expensive Lawyers, Replaced by Cheaper Software' (Markoff, 2011). In reality, the software replaced some lawyers in performing one highly structured task: classifying documents in the 'discovery' phase of a trial in which each party can require the other parties to turn over all correspondence, memos, etc. relevant to a specific issue. A party receiving the request must apply a lawyer's judgment to a potentially large set of documents to determine which documents are responsive to the request.

Traditionally, document review was performed by junior lawyers. The 1990s move to digitized documents enormously increased the volume of documents to be reviewed, the cost of discovery, and pressure for an automated solution. The solution ultimately[14] involved a machine-learning application—'predictive coding'. A group of lawyers reads through a sample of the documents to be reviewed and classifies each sample document as responsive or unresponsive. The classified sample of documents becomes the training sample for a machine-learning model. The model is similar in spirit to equation (1)—the judge's decision-making process—where the dependent variable is a (0,1) variable depending on whether the sample document was 'not responsive' or 'responsive' to the discovery request, and the independent variables are words, word sequences, etc. in the document. The typical model reports a prediction and a probability for each document—e.g. 'This document is responsive with probability 0.87.'[15]

Analysis of lawyers' billing records suggests that if predictive coding and several related applications were all adopted immediately—a very unlikely scenario—software could substitute for roughly 13 per cent of corporate lawyers' time. The rest of a lawyer's time is spent on unstructured cognitive work: developing legal arguments, plotting courtroom strategy, etc. (Remus and Levy, 2017, p. 533 ff). Note that 13 per cent substitution is significantly larger than software's effect on the low-wage janitor's job and much less than software's effect on the middle-wage job of medical transcriptionist.

Taken together, these occupational descriptions help to explain automation's role in the gradual polarization of the occupational distribution (Levy and Murnane, 2004; Buyst *et al.*, 2018, this issue). Automation of the lowest wage jobs is hard because these jobs require unstructured physical activity and unstructured social interaction. Automation of the highest wage jobs is hard because the jobs (including many of the jobs AI creates) require unstructured cognitive activity and unstructured social interaction. Automation's impact is relatively larger in mid-wage, mid-skilled occupations because they involve relatively higher levels of structured physical and/or cognitive tasks. These are relative distinctions. We have seen that part of a lawyer's job can be automated and camera-enabled robots can automate part of the job of a security guard. But on balance, near-term AI will have the greatest effect on blue-collar work, clerical work, and other mid-skilled occupations.

Gradual occupational polarization, driven by AI, is reinforced by globalization. Many of the low-wage occupations that are hard to automate (janitor/building cleaner,

[14] An earlier solution used deductive instructions to search documents for keywords. This approach is still used in some situations but can be problematic since an idea can be expressed in a variety of words while a single word can have multiple meanings (chocolate chip, computer chip, paint chip).

[15] The probabilities permit a triage where documents with low probabilities of being responsive (e.g. $p < 0.4$) are discarded, documents with a high probabilities of being responsive ($p > 0.9$) are automatically given to the opposing party, and documents with ($0.4 < p < 0.9$) are reviewed again, this time by human lawyers.

home health aide,[16] etc.) must be performed in person and so cannot be offshored. High-wage occupations with substantial unstructured cognitive work (dealing with 'special cases') are hard both to automate and to explain to hand over to others—particularly persons not fluent in English. Compared to many low-wage and high-wage occupations, a greater fraction of mid-skilled occupations consists of structured tasks that do not have to be performed on site and can be described in instructions that can be automated or explained to foreign producers (Levy and Murnane, 2004). Thus, call centre work is performed by both Philippine operators using scripts (instructions) and by natural language-processing software. Domestic manufacturing jobs move to offshore producers and are replaced by domestic robots. The US tax code is a set of rules, and basic US tax returns are offshored to junior accountants in India and carried out by software such as TurboTax and TaxCut. While some displaced workers move to higher-wage jobs, more are moving to lower-wage jobs.[17]

To show this, Figure 3 displays the 2000 and 2016 distribution of all employed persons across occupational categories. Categories are listed left to right by increasing median hourly earnings. The left bar in each pair refers to the occupation's employment share in 2000. Category bars in black are those that score high in routine cognitive content and/or routine physical content using measures developed by Acemoglu and Autor (2011) and so are potentially most vulnerable to computer substitution and globalization. The fraction of employed persons in these categories declined from 39.2 per cent in 2000 to 33.3 per cent in 2016. Simultaneously, fractions of persons in higher- and lower-wage categories both increased in a pattern of gradual occupational polarization.

This polarization comes with a caveat. When tracking all employed persons over 16 years (Figure 3), occupational shifts over time are caused by technology and globalization, but also by a changing mix of workers, including shifts in levels of education and the proportions of men and women. To remove the influence of demographic factors, Figure 4 compares the 2000 and 2016 occupational categories for employed white men who do not have a BA (but who may have an Associates' Degree). Again, pairs of black markers refer to categories that are most vulnerable to automation or offshoring. Taken as a group, the proportion of the men in mid-skilled categories fell from 42.5 per cent in 2000 to 39.3 per cent in 2016, an employment loss of 1.3m mid-skilled jobs. The corresponding employment gain was largely in lower-wage jobs including food preparation and serving, building and grounds cleaning and maintenance, etc.

To be clear, the economy still has significant numbers of mid-skilled jobs, including some jobs driven by AI (Mandel, 2017)—but the near-term trend is negative.

IV. The near-term future: how fast will AI applications advance?

Long-term projections of AI's impact on jobs frequently end in mass unemployment. Whether or not these projections are correct, AI's near-term effect is not mass unemployment but occupational polarization and a slowly growing fraction of persons

[16] 2016 Median Annual Pay: $22,600.

[17] In reality, there is no clean distinction between automation/digitization and globalization since much of today's globalization would be impossible without the internet.

Figure 3: US occupational distributions for 2000 and 2016, and 2016 median hourly earnings (right-hand axis)

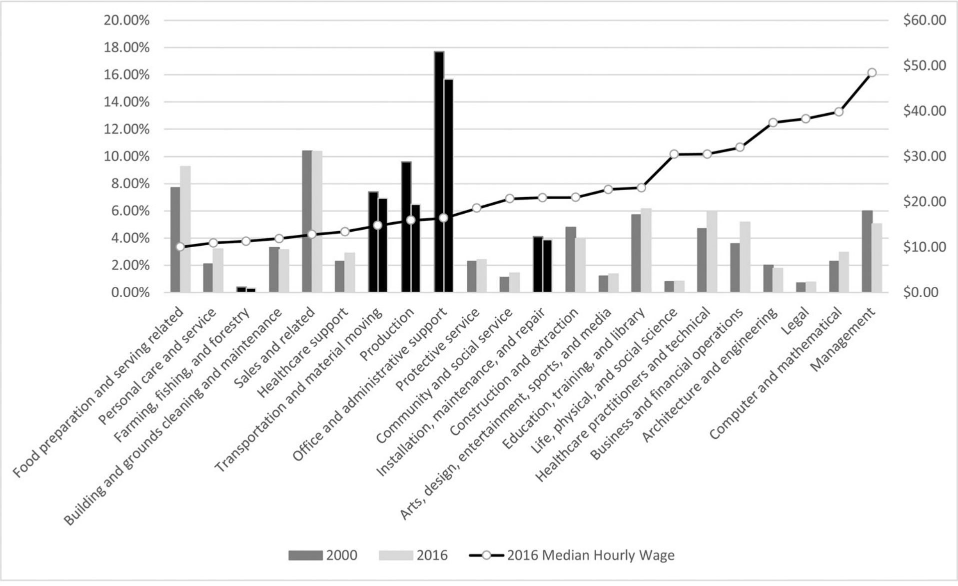

Note: Pairs of black bars denote occupations vulnerable to automation and offshoring.
Source: Author's tabulations of 2000 Decennial Census and 2016 American Community Survey.

moving from mid-skilled jobs into lower-wage work. If AI job displacement has near-term political repercussions, this will be the reason.

The previous section identified current occupational categories most vulnerable to automation (Figure 3). This section projects job losses in the near-term future by sketching the potential diffusion of three 'hot' technologies through 2024. 'Autonomous long-distance trucks' replaces 'Heavy and tractor-trailer truck drivers' (in the category of 'Transportation and material moving occupations'). 'Automated customer responses' replaces 'Customer service representatives' (in 'Office and administrative support occupations'). 'Industrial robots' replaces assembly line workers which we approximate here with the BLS occupational title of 'Assemblers and fabricators' (in 'Production occupations').[18] To provide some perspective, I compare my assumptions about near-term job losses with BLS 10-year occupational projections for 2014–24 and the recently released 2016–26.

(i) Autonomous long-distance trucks

A new technology enters the economy through three stages: performance in the ideal conditions of the laboratory, performance in the workplace with all its complications, and widespread workplace adoption. Many technologies occupy more than one stage at

[18] The BLS lists the occupational category of 'Production occupations' that includes multiple job titles but does not have a detailed job title of 'Production worker'.

Figure 4: US occupational distribution for white men with education lower than BA level, 2000 and 2016

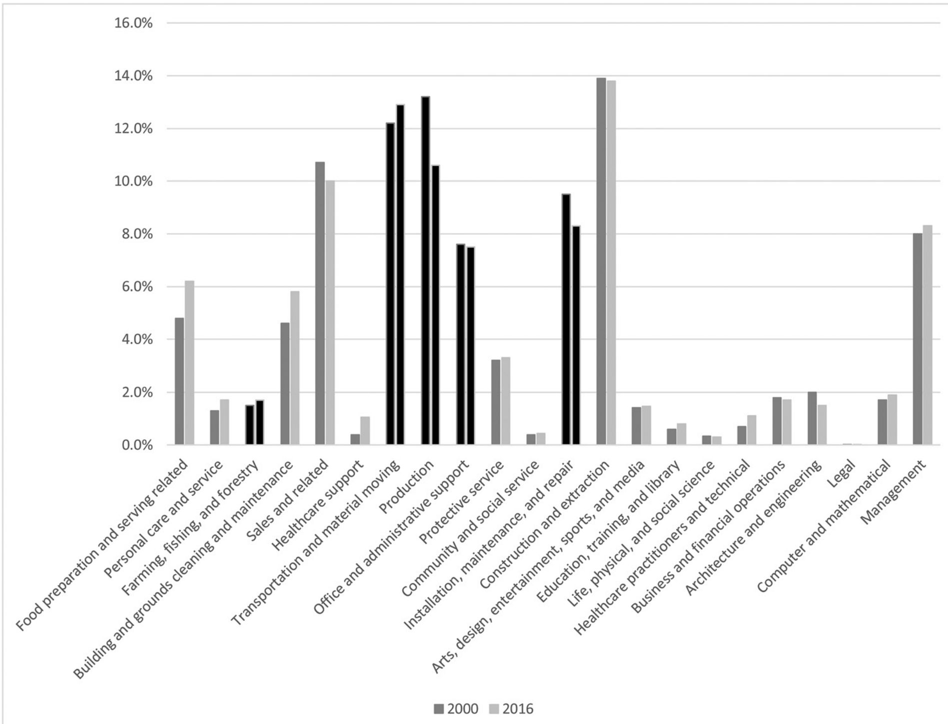

Note: Occupations listed in increasing order of average wages. Pairs of black bars denote occupations vulnerable to automation and offshoring.
Source: Author's tabulations of 2000 *Decennial Census* and 2016 *American Community Survey*.

a time. Autonomous truck/heavy vehicle technology is an example. Autonomous heavy vehicles now operate successfully in certain restricted (structured) areas, but the technology is not ready for highways. (There is little current discussion of autonomous trucks in city driving, a much more difficult problem.)

The possibility of self-driving trucks gained substantial public attention in October 2016 when an autonomous semi-trailer developed by Otto, a San Francisco start-up, drove 50,000 cans of Budweiser beer from Fort Collins Colorado to Colorado Springs (Davies, 2016). The BLS reports 1.87m persons working as 'heavy and tractor-trailer truck drivers'. The Otto demonstration and other tests led to articles suggesting a 'tsunami' of displacement (Sanchez, 2017; Ford, 2017).

Near-term truck driver job losses are likely to be small, with some acceleration starting in about 5 years. The Otto demonstration was equivalent to success in the laboratory: driving in good weather on a route that had been digitally mapped in advance[19] with few unanticipated obstacles. While AI can progress in unexpected leaps, it will

[19] Current autonomous vehicle technologies require a digital map of the road as well as information from onboard sensors.

likely take some time before these and other limits can be relaxed (Markoff, 2017). In addition, autonomous trucks underline the limits of the 'science push' model since the technology will diffuse under both federal and state regulations and insurers' restrictions. Taken together, these factors suggest the near-term diffusion of autonomous truck technology will emphasize two simpler problems: operation in controlled environments and long-distance truck platooning.

Autonomous operations in controlled (structured) environments are already in wide use. Heavy-duty carriers pick up and distribute cargo containers in major ports, including Rotterdam and Sydney (Saulwick, 2015). Autonomous trucks, earth loaders, and trains operate in Rio Tinto's Australian Pilbara mines and mines in Sweden (Ellem, 2016). In these environments, all routes can be digitally mapped in advance and administrators can limit unanticipated contingencies. Because these applications involve large capital expenditure, ramp-up will be slow, but by the end of 7 years, there should be a number of fully automated US port facilities, mines, and other industrial facilities. These applications will likely involve a modest loss of driver jobs (Cosbey *et al.*, 2016)

Truck platooning is a highway driving strategy in which one semi-trailer truck closely follows another, a strategy similar to stock car drafting. The strategy reduces overall air drag and allows both trucks to reduce gasoline consumption. Several firms, including Peloton Technology, have developed digital controls in which the acceleration and braking on the follower truck are controlled by a combination of Wi-Fi, radar, and other sensors. The controls allow a shorter (and safer) gap between trucks resulting in greater gasoline savings. Technology-enabled platooning faces few regulatory barriers and the technology is sufficiently inexpensive that it should rapidly come into use.

Platooning technology currently assumes a human driver in the follower truck to steer and to handle emergencies.[20] Over a 5-to-7-year horizon, the technology will likely expand to control steering and the follower truck will become autonomous over highways (Janssen *et al.*, 2015). This automation will result in some loss of long-distance truck driver jobs.

The follower truck in a platoon has only limited autonomy. Full autonomy—no lead truck—introduces additional engineering problems. All autonomous truck routes require digital maps. More important, navigating without a lead truck, as the Otto truck did, requires long-distance sensors including Lidar. Lidar's current performance, unlike radar's performance, degrades badly in rain or snow.

Both past and present BLS occupational projections for heavy and tractor-trailer truck driver jobs assume no jobs eliminated by technology. The BLS projections for 2014–24 projected the jobs to grow by 5 per cent in line with total employment growth in the economy. The explanatory notes mentioned platooning but assumed it would not emerge by 2024 (personal communication, September 2017). Autonomous trucks were not mentioned. The newly released 2016–26 projections assume a 6 percent growth in employment. Explanatory notes now mention the platooning but assume platooning has no employment effects and autonomous trucking again is not mentioned.[21]

[20] As part of the technology, the driver of the follower truck has a screen that allows him to see the road as if he were driving the lead truck.

[21] Discussion of employment projection at: https://www.bls.gov/ooh/transportation-and-material-moving/heavy-and-tractor-trailer-truck-drivers.htm#tab-6

My own assumption is that technological development will permit demonstrations of fully autonomous trucks within 4 or 5 years. If these demonstrations prove successful and the technology begins to spread, operating in dedicated highway lanes, it is reasonable to project that in 2024, heavy and tractor-trailer truck driver jobs will have increased by 2 per cent in 2024, a reduction 76,000 jobs below the original 2014–24 BLS projection.

(ii) Automated customer service responses

Multiple start-ups now work on improving automated customer service responses[22] but simplified technology has been in the workplace since at least 1992 when AT&T announced it would replace up to one-third of its 18,000 long-distance operators with speech recognition software (Gellene, 1992). This early speech automation was feasible because it was replacing highly structured scripts. For example: 'This is the long-distance operator. I have a collect call from [Caller Says Name]. Will you accept charges? Please say yes or no.'[23]

Despite continued technical development and significant call centre offshoring (Lee, 2015), the BLS estimates that 2.8m customer service representatives currently work in the United States. The number reflects the fact that while some customer questions (called or texted) are highly structured and can receive automated responses, many others are not. One goal of software development is to distinguish between these cases in ways that satisfy the customer.

Current customer service technology involves three software components:

- automated speech recognition (ASR) to properly represent the words in the customer's call. ASR is not required in text communication;
- natural language processing (NLP) to extract the meaning from the call or text: what does the customer wants to accomplish?
- information retrieval (IR) to retrieve the 'best' answer from the call centre's database.

Advances in machine learning have led to substantial progress in both ASR and NLP. The issues surrounding IR are familiar to anyone who has used Google search. If the search is for a factoid—e.g. a 'who' or a 'when' or a 'how much'—the software immediately produces a highly accurate answer that could be an appropriate response to a customer. If the search requires a more complicated answer—a 'how' or a 'why'—the software produces a list of links to text that are possible answers. A typical caller wants a single answer—not a list of links.

Given the technology's current limits, automated customer service responses operate as elaborated versions of Apple's Siri or Amazon's Alexa. The software starts by estimating the probability that it recognizes the incoming question (voice or text) as part of a question/answer pair in its database. If the probability is above a threshold established by the call centre, the customer receives the database answer. If the probability is below the threshold, the caller is routed to a human operator. The fraction of caller responses

[22] For example, https://www.digitalgenius.com

[23] The first version of the automated script did not include the phrase 'Please say yes or no'. As a result, the software received many unanticipated responses to the question that it could not process—'That bum—I'll never speak to him again as long as I live.'

that are successfully automated varies substantially from industry to industry: a large fraction of calls to an electric utility involve a small set of billing questions, while calls to a retailer such as L. L. Bean are much more varied.

Over time, developers have expanded this framework in several dimensions, as follows.

- A frequently asked question may have a complicated, but unchanging answer: 'How do I locate the registration number on my television set?' In this case, it may be possible to form an answer that is retrieved as if it were a factoid. The language processing problem is reduced to determining whether the customers' words are, in fact, the frequently asked question.
- After receiving an automated response, a customer is given the option of requesting to speak to an operator. The request is equivalent to a rejection of the automated response and so can be used to further train the software.
- When software cannot produce a single correct answer, it can produce a list of possible answers or a template for a correct answer that can be displayed on the screen of a customer service representative, thereby saving her time.
- In a call centre where different representatives have different expertise, software can route the call or text to the correct representative.
- Historically, a customer service representative would fill out a digital 'ticket' for each call—information entered into a searchable database to identify problems, improve responses, etc. These tickets are increasingly filled in automatically by software, saving the representative's time.

Near-term employment projections for customer service representatives are shaped by two opposing forces. Because the technology resides in software that is increasingly migrating to the cloud, technology improvements can diffuse through the economy much more quickly than technology embodied in expensive capital equipment like autonomous semi-trailer trucks. Conversely, the growing substitution of online sales for brick and mortar stores means that questions that had been posed to store personnel are now directed to customer service representatives. In its 2014–24 occupational projections BLS had predicted the number of customer service representatives would grow by 10 per cent between 2014 and 2024. The projection assumed no employment impact, positive or negative, of software innovations (personal communication, May 2017). The updated 2016–26 projection takes note of technology stating:

> some companies are increasingly using Internet self-service or interactive voice-response systems that enable customers to perform simple tasks, such as changing addresses or reviewing account billing, without speaking to a representative. Improvements in technology will gradually allow these automated systems to perform more advanced tasks.[24]

Due in part to this trend, the 2016–26 projection lowers estimated growth in customer service representatives from the previous 10 per cent to 5 per cent. My own assumption is that technology will advance more quickly than even the revised BLS projection and the number of customer service representatives will be no different in 2024 than it is today, a reduction from the original 2014 projection of 260,000 jobs.

[24] https://www.bls.gov/ooh/office-and-administrative-support/customer-service-representatives.htm#tab-6

(iii) Industrial robotics

Industrial robotics, like automated customer service responses, exists in all three phases of diffusion. Work in the laboratory currently focuses on increased 'human' capacities: easy reprogramming to perform multiple tasks; an ability to sense, grasp, and manipulate objects encountered for the first time; an ability to work safely side-by-side with humans. Current development also emphasizes cost reduction to make robotic arms and other small devices affordable for small manufacturers.

Estimating robots' employment impact is hampered by limited data. Several recent papers use data on industrial robot sales from the International Federation of Robotics (IFR, 2016) (Graetz and Michaels, 2015; Acemoglu and Restrepo, 2017; Presidente, 2017). These data extend back to 1993 and represent a significant increase in our knowledge, but they count only 'automatically controlled, reprogrammable, and multipurpose [machines]'. A robot dedicated to a single purpose—for example, machinery that automatically assembles circuit boards—is not counted even though it may substitute for human workers.

It follows that projected employment impacts based on IFR data are lower bound estimates of the number of production workers displaced by robotics. Presidente (2017) uses IFR data to estimate a stock of US robots (as defined above) of 234,000 in 2015. Over the year term, the IFR projects that the stock of US robots will grow by about 4 per cent per year. If that growth rate were sustained for through 2024, the stock of robots would grow by 100,000 by 2024. If we conservatively assume that each robot replaces two production workers (including workers on different shifts), the increased US robot stock would cause production worker employment to be at least 200,000 less than it otherwise would have been.

BLS occupational projections tell a similar story. In 2014, BLS reported 1.8m 'assemblers and fabricators', a number projected to be 1 per cent lower (–18,000) in 2024. This projection made no assumptions about robotics (personal communication, September 2017). The updated BLS projection for 2016–26 now assumes job losses from automation:

> In most manufacturing industries, improved processes, tools, and, in some cases, automation will reduce job growth. Increasingly, new advances in robotics have enabled machinery to perform more complex and delicate tasks previously performed by workers. In addition, assemblers and fabricators are increasingly working alongside robots, also known as 'collaborative robotics'. These new robots can help workers perform tasks and increase efficiency. However, this increased efficiency may reduce the demand for some assemblers and fabricators.[25]

As a result, the 2016–26 projection has been lowered from a 1 percent decline to a 13 per cent decline. My assumption is that the 13 per cent reduction is reasonable, leading to 216,000 fewer jobs in 2024 than the original BLS 2014–24 projection.

[25] https://www.bls.gov/ooh/production/assemblers-and-fabricators.htm#tab-6

(iv) Summary

Between 1999 and 2011, the US lost 4.5m manufacturing production jobs through a combination of Chinese imports, automation, and the Great Recession.[26] By that standard the estimated near-term job impacts of these three AI technologies are modest: roughly 552,000 fewer jobs than BLS had originally projected for 2024.

At the same time, the three technologies will raise awareness of AI-related job losses beyond their limited numbers. Autonomous trucks driving in dedicated lanes will appear on the roads and in multiple news stories. More stories will appear about people working side-by-side with robots (Harlan, 2017). Stories will describe other AI-related job losses—e.g. the (partial) role of online sales in retail sales layoffs (Townsend *et al.*, 2017). When stories describe AI-induced layoffs in a particular occupation, other persons in that occupation will assume they are at risk as well.

V. Automation and populism

There is good evidence that manufacturing job losses, driven in part by Chinese imports, played a role in the 2016 presidential elections (e.g. Autor *et al.*, 2017). Since automation anxiety is already high (Smith and Anderson, 2017), it is likely that AI-induced job losses will eventually create their own political reaction.

The question is whether that reaction will take a populist form. The beginning of an answer appeared in a series of columns by Jed Kolko, Chief Economist at Indeed.com (2016*a*,*b*,*c*). This section elaborates on Kolko's work.

Let us begin with definitions. The US population has grown ideologically polarized[27] over the last two decades (Pew Foundation, 2014), but ideological polarization can involve conflict between equally powerful groups. Populism describes a particular kind of group conflict by adding the ideas of victimization and illegitimacy. As Oliver and Rahn (2016) write:

> At its core, populism is a type of political rhetoric that pits a virtuous 'people' against nefarious, parasitic elites who seek to undermine the rightful sovereignty of the common folk. . . . Its tone is Manichean, casting politics as a bifurcated struggle between 'the people', on one hand, and a self-serving governing class undeserving of its advantaged position, on the other. Its goal is restorative, replacing the existing corruption with a political order that puts the people back in their proper place and that is more faithful to their longings and aspirations. Its worldview is apprehensive, suspicious of any claims to economic, political, or cultural privilege; for populists, the good is found in the common wisdom of the people rather than the pretensions of the expert.

[26] Data on manufacturing production jobs come from the St Louis Federal Reserve Bank—https://fred.stlouisfed.org/series/CES3000000006#0. Acemoglu *et al.* (2016) estimate that about 560,000 manufacturing jobs (both production and supervisory) were lost due to Chinese imports with the rest due to automation and recession. Jaimovich and Siu (2012) argue that layoffs in recessions create opportunities to accelerate adoption of automation technology.

[27] Occupational polarization refers to growing concentrations of employment at higher- and lower-wage jobs. In political polarization, members of a group can share an ideological position while having different incomes.

A populist politician defines 'the people' by the rhetoric he uses and the audiences he courts. Actual populist campaigns are infrequent[28] but in the 2016 election, President Trump ran such a campaign. His strong supporters were 'the people'.[29] He frequently reminded his supporters of how they were disrespected and lied to by 'the elite', something many supporters already believed (Cramer, 2016; Hochschild, 2016; French, 2017; Stern, 2017).

Compared to the general population, polling identified Trump's voters as more likely to live outside large urban areas, more likely not to live on the east or west coasts, less likely to have a Bachelor's degree, and more likely to be white.[30] There is no analogous polling that defines the elite, but a reasonable description would be highly educated professionals (not 'the rich'[31]) who live in coastal cities.

Many voters in the 2016 electorate had good reason to view the political establishment as incompetent and self-serving,[32] but the broader, populist indictment gained support in part through social media propaganda and 'fake news' (Holan, 2016). In 1951 Hannah Arendt in *The Origins of Totalitarianism* described propaganda's attraction in a time of economic and cultural dislocation:

> What convinces the masses are not facts, and not even invented facts, but only the consistency of the system of which they are presumably part. . . . They are predisposed to all ideologies because they explain facts as mere examples of laws and eliminate coincidences by inventing an all-embracing omnipotence which is supposed to be at the root of every accident. . . . Totalitarian propaganda can outrageously insult common sense only where common sense has lost its validity. (Arendt, 1976, revised edn, pp. 351–2)

Not every populist movement is authoritarian, but the 2016 versions of US and European populism included strong elements of authoritarianism and nativism (Ballard-Rosa *et al.*, 2017; Colantone and Stanig, 2017), reason enough to ask whether AI-induced job losses, like job losses from trade, will make populism more attractive.

There are several reasons why this might *not* happen.

- Populist movements appear to require multiple pre-conditions and so have been rare and short-lived (Oliver and Rahn, 2016). By the time computer-induced job displacement becomes significant, the political landscape may have shifted.
- Compared to the loss of trade-induced plant closures 1999–2011, most AI job losses are likely to be slower (section IV). They are also likely to be more

[28] Prior to Donald Trump, the last candidate to run a widely recognized populist campaign was George Wallace in 1968.

[29] Not all Trump voters were strong supporters. CNN's election 2016 exit polls indicate that 37 per cent of Trump voters strongly supported Trump, while 28 per cent of Trump voters were voting against Clinton: http://edition.cnn.com/election/results/exit-polls/national/president

[30] CNN *Election 2016 Exit Polls*. Note that income is not included in the list. Roughly speaking, Clinton attracted high- and low-income voters, while Trump attracted high- and middle-class voters.

[31] On working class contempt for professionals as distinct from the rich, see Lamont (2002) and Williams (2016).

[32] The origins and aftermath of the 2008 financial crisis were particularly important in discrediting the elite. Only one high-level banker went to jail after 2008 compared to 1,100 persons who were prosecuted in the 1989 Savings and Loan crisis (Eisinger, 2014). Similarly, the post-1999 loss of manufacturing jobs drew little attention from Washington compared to 1981when Ronald Reagan negotiated 'voluntary' restraints with Japan to slow imports of Japanese cars (Crandall, 1987)

geographically diffuse than manufacturing job losses and so harder to politicize (e.g. Arnold, 1990, pp. 132–5).

- At least some candidates for near-term automation—customer service representatives, heavy and tractor-trailer truck drivers—experience high rates of turnover (e.g. Phillips, 2015) such that job losses need not translate into layoffs.
- Historically, populist movements have focused more on job losses from trade ('cheap foreign labour') and immigration than on technology (Rodrik, 2017).

Suppose, however, that a politician wanted to frame AI-induced job losses in populist terms. 'The people' would be those who had lost jobs or were threatened with job loss due to automation. The 'elite' would be the people who were creating the software—software developers, venture capitalists, and so on. In many dimensions, the two groups would represent a replay of the 2016 presidential election.

With respect to geography, Figure 5 shows the geographic distribution of 'software application developers' and 'heavy and tractor-trailer truck drivers'. Darker shading in the figures identifies areas with higher occupational concentration ratios. If, for example, a particular area is home to 5 per cent of all US employed persons and 10 per cent of all US long-distance truck drivers, the area's concentration ratio area for truck drivers is 2.0. With a few exceptions—Salt Lake City, Austin Texas, Madison Wisconsin—software application developers are concentrated along the east and west coasts, typically in urban areas. Heavy and tractor-trailer truck drivers are concentrated in states between the two coasts, usually outside of metro areas (Kolko, 2016*a*).

Similarly, the five highest concentrations of 'assemblers and fabricators' are located in Kentucky, Ohio, Indiana, Illinois, and Oregon, and the five highest concentrations of 'customer service representatives' are located in Arizona, Utah, Florida, Idaho, and Georgia.

With respect to race and education, *CNN Exit Polls* permit direct comparison of Trump and Clinton voters with persons in the four occupations. As shown in the first two columns of Table 1, white men and women whose education stopped before a BA comprised 48 per cent of Trump's vote and slightly under 20 per cent of Clinton's vote. Columns 3–6 of the table use the combined American Community Surveys for 2011–14

Table 1: Demographic composition (%) of Clinton and Trump voters, long-distance lorry drivers, customer service representatives, assemblers and fabricators, and software developers and programmers

	Clinton voters	Trump voters	Long-distance lorry drivers	Customer service representatives	Assemblers and fabricators	Software developers and programmers
White women < BA	12.1	22.7	4.4	38.3	23.3	4.6
White men < BA	7.7	24.9	66.1	17.7	40.4	15.8
Non-whites < BA*	25.4	7.0	24.3	24.0	31.1	4.3
White women ≥ BA	21.3	19.3	0.3	8.4	0.9	10.3
White men ≥ BA	13.9	19.8	3.7	6.3	2.4	40.3
Non-whites ≥ BA*	19.6	6.3	1.2	5.3	1.9	24.7
Totals**	100	100	100	100	100	100

Notes: *CNN samples were too small to disaggregate non-white voters by gender and education; ** totals exclude responses of 'other/no answer'.
Source: *CNN Exit Polls*: http://edition.cnn.com/election/results/exit-polls

Figure 5: Geographic concentration of software application developers, and heavy and semi-trailer truck drivers

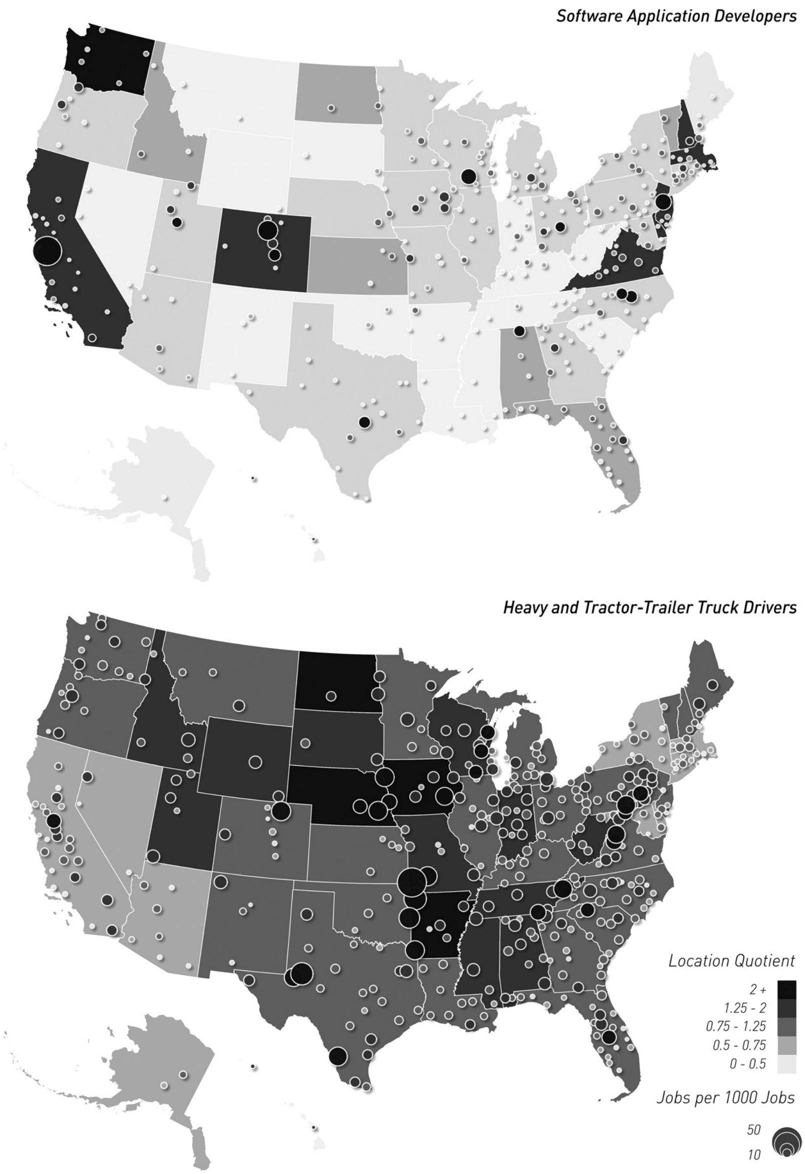

Source: US Bureau of Labor Statistics *Occupational Employment Statistics* for 2016.

to similarly show that white men and women whose education stopped before a BA account for 70 per cent of heavy and semi-trailer truck drivers, 56 per cent of customer service representatives, and 63 per cent of assemblers and fabricators, but 20 per cent of software application developers and programmers.

These data suggest that a populist politician who campaigned on AI-induced job loss would start with ready-made definitions of the ‘people’ and the ‘elite’ based on national fault lines that were sharpened in the 2016 presidential election. This politician could also point to an evident lack of respect: a set of highly educated coastal ‘elites’ who make a very good living developing robots to put ‘the people’ out of work.

As Arendt argued, the campaign’s focus on economic insecurity would create demands for fake news—demands that many domestic and foreign entrepreneurs would be willing to fill (Timberg, 2016; Gillin, 2017). Even today, raw material for fake news is provided by journalists who compete for readership by exaggerating the significance of laboratory demonstrations like Otto and by the dramatic claims of start-ups themselves as they fight for venture capital funding. In the campaign atmosphere, Otto could meet Hannah well before ‘the future of work’ had arrived.

VI. Conclusion

It is possible that artificial general intelligence—AI that out-performs humans in all ways—will arrive in several decades and the resulting employment disruptions will dwarf the disruptions described in this paper. But between today and 2040, the AI that already exists will disrupt the nation’s occupational structure. The question is whether these disruptions will seriously destabilize the country’s political and social structure.

Part of the answer involves the speed of technical disruptions. The evidence in this article is mixed, but on balance there appears to be some time to develop anticipatory policies to assist people who will lose jobs and other people who, under earlier circumstances, would have expected to take those jobs.

The shape of these policies is not yet clear (at least to me). What is clear is that over the near term, adjustments will begin in local labour markets that vary dramatically from technical centres to ‘meds and eds’ communities where hospitals and schools are the main employers. Sweeping national proposals—a guaranteed income, reforming the educational and training systems—are not near-term substitutes for the local knowledge these adjustments require: how many people will be affected by automation, what employers are not affected by automation, how the area is trying to attract new jobs, the particular competencies of community college training programmes, and so on. Learning these details for particular areas, like learning the details of particular technologies, is a step towards a more rational discussion.

A second part of an answer involves the presence of political will to facilitate local areas in making adjustments. We have been down this road before in failing to deal with the fallout from international trade and offshoring:

> In principle, US governments could have followed the European model. It could have complemented trade agreements—NAFTA, the WTO, and China’s WTO entry—with much more robust social insurance mechanisms and active labor-market programs and protections. Even though there was much talk about such supports during the Clinton administration, little was in fact done beyond

tinkering with the existing Trade Adjustment Assistance (TAA) mechanisms. (Rodrik, 2017, p. 11)

Until this past autumn, it appeared that, in the United States, disruptive policies including both expanded trade and technology could be implemented with no adverse consequences, at least for the people doing the implementation. That is no longer true and it is appropriate that discussions of technology take account of that fact.

References

Acemoglu, D., and Autor, D. (2011), 'Skills, Tasks and Technologies: Implications for Employment and Earnings', ch. 12 in O. Ashenfelter and D. E. Card (eds), *Handbook of Labor Economics*, vol. **4**, 1043–171.

— Restrepo, P. (2017), 'Robots and Jobs: Evidence from US Labor Markets', National Bureau of Economics Working Paper 21285, March.

— Autor, D., Dorn, D., Hanson, G. H., and Price, B. (2016), 'Import Competition and the Great US Employment Sag of the 2000s', *Journal of Labor Economics*, **34**(1), pt. 2, s141–98.

Agrawal, A., Gans, J., and Goldvarb, A. (2016), 'The Simple Economics of Machine Intelligence', *Harvard Business Review*, 27 November, https://hbr.org/2016/11/the-simple-economics-of-machine-intelligence xxxxx

Arendt, H. (1976), *The Origins of Totalitarianism*, revised edn, Houghton Mifflin Harcourt.

Arnold, R. D. (1990), *The Logic of Congressional Action*, New Haven CT, Yale University Press.

Autor, D. H., Levy, F., and Murnane, R. J. (2003), 'The Skill Content of Recent Technical Change: An Empirical Investigation', *Quarterly Journal of Economics*, **118**(4), 1279–334.

— Dorn, D., Hanson, G., and Majlesi, K. (2017), *'A Note on the Effect of Rising Trade Exposure on the 2016 Presidential Election'*, Working Paper, Department of Economics, MIT, January.

Ballard-Rosa, C., Malik, M., Rickard, S., and Scheve, K. (2017), *'The Economic Origins of Authoritarian Values: Evidence from Local Trade Shocks in the United Kingdom'*, Stanford University Working Paper.

Bessen, J. (2015), *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*, New Haven, CT, Yale University Press.

Buyst, E., Goos, M., and Salomons, A. (2018), 'Job Polarization: An Historical Perspective', *Oxford Review of Economic Policy*, **34**(3), 461–74.

Colantone, I., and Stanig, P. (2017), 'The Trade Origins of Economic Nationalism: Import Competition and Voting Behavior in Western Europe', BAFFI CAREFIN Centre Research Paper Series No. 2017–49, January.

Cosbey, A., Mann, H., Maennling, N., Toledano, P., Geipel, J., Brauch, M. D. (2016), *Mining a Mirage?*, International Institute for Sustainable Development and Columbia University Institute for Sustainable Development, available at http://www.iisd.org/sites/default/files/publications/mining-a-mirage.pdf

Cramer, K. J. (2016), *The Politics of Resentment: Rural Consciousness in Wisconsin and the Rise of Scott Walker*, Chicago, IL, University of Chicago Press.

Crandall, R. W. (1987), 'The Effects of US Trade Protection for Autos and Steel', *Brookings Papers on Economic Activity*, **1**, 271–88.

Davies, A. (2016), 'Uber's Self-Driving Truck Makes its First Delivery: 50,000 Beers', 25 October, https://www.wired.com/2016/10/ubers-self-driving-truck-makes-first-delivery-50000-beers/

Eisinger, J. (2014), 'Why Only One Top Banker Went to Jail for the Financial Crisis', *New York Times Sunday Magazine*, 30 April, https://www.nytimes.com/2014/05/04/magazine/only-one-top-banker-jail-financial-crisis.html

Ellem, B. (2016), 'Geographies of the Labour Process: Automation and the Spatiality of Mining', *Work, Employment and Society*, **30**(6), 932–48.

Ford, M. (2017), 'Driverless Trucks: Economic Tsunami May Swallow One of the Most Common US Jobs', 16 February, https://www.theguardian.com/technology/2017/feb/16/self-driving-trucks-automation-jobs-trucking-industry

French, D. (2017), 'Trump's Tweets are Damaging the Republican Character', *National Review*, 12 October, http://www.nationalreview.com/article/452618/donald-trump-twitter-damaging-republican-character

Frey, C. B., and Osborne, M. A. (2013), 'The Future of Employment: How Susceptible are Jobs to Computerization', working paper, Oxford Martin School, 17 September.

Gellene, D. (1992), 'AT&T Will Cut Up to a Third of Its Operators: Communications: The Company Will Replace as Many as 6,000 with a Voice-recognition System that Can Place Collect Calls', *Los Angeles Times*, 4 March, http://articles.latimes.com/1992-03-04/business/fi-3209_1_voice-recognition-systems

Gillin, J. (2017), 'Fake News Website Starts as Joke, Gains 1 Million Views within 2 Weeks', http://www.politifact.com/punditfact/article/2017/mar/09/fake-news-website-starts-joke-gains-1-million-view/, 9 March.

Graetz, G., and Michaels, G. (2015), 'Robots at Work', Centre for Economic Performance, Discussion Paper # 1335, March.

Harlan, C. (2017), 'Rise of the Machines', *The Washington Post*, 5 August, https://www.washingtonpost.com/national/rise-of-the-machines/2017/08/05/631e20ba-76df-11e7-8f39-eeb7d3a2d304_story.html?utm_term=.0aae2cd6c88e

Hochschild, A. R. (2016), *Strangers in Their Own Land: Anger and Mourning on the American Right*, The New Press.

Holan, A. D. (2016), '2016 Lie of the Year: Fake News', *Politifact*, 13 December, http://www.politifact.com/truth-o-meter/article/2016/dec/13/2016-lie-year-fake-news/

IEEE Spectrum (2015), 'A Compilation of Robots Falling Down at the Darpa Robotics Challenge', posted 6 June, https://www.youtube.com/watch?v=g0TaYhjpOfo

IFR (various years), *World Robotics: Industrial Robots*, International Federation of Robotics, Statistical Department, http://www.worldrobotics.org

Jaimovich, N., and Siu, H. E. (2012), 'The Trend is the Cycle: Job Polarization and Jobless Recoveries', NBER Working Paper w18334.

Janssen, R, Zwijnenberg, H., Blankers, I., de Kruijff, J. (2015), *Platooning: Driving the Future of Transportation*, TNO—Innovation for Life, https://www.tno.nl/en/about-tno/news/2015/3/truck-platooning-driving-the-future-of-transportation-tno-whitepaper/

Kolko, J. (2016*a*), 'Republican-leaning Cities are at Greater Risk of Job Automation', www.FiveThirtyEight.com, 17 February.

— (2016*b*), 'Will Your Job Disappear? Economic Anxiety, Demographics, and the Future of Work', www.JedKolko.com, 1 September.

— (2016*c*), 'Trump Was Stronger Where the Economy Is Weaker', www.FiveThirtyEight.com, 10 November.

Lamont, M. (2002), *The Dignity of Working Men*, Cambridge, MA, Harvard University Press.

Lee, D. (2015), 'The Philippines Has Become the Call-center Capital of the World', *Los Angeles Times*, 1 February, http://www.latimes.com/business/la-fi-philippines-economy-20150202-story.html

Levy, F., and Murnane, R. J. (2004), *The New Division of Labor: How Computers are Creating the Next Job Market*, Princteon, NJ, Princeton University Press and the Russell Sage Foundation.

Mandel, M. (2017), *The Creation of a New Middle Class? A Historical and Analytic Perspective on Job and Wage Growth in the Digital Sector, Part I*, Progressive Policy Institute, March.

Markoff, J. (2011), 'Armies of Expensive Lawyers, Replaced by Cheaper Software', *The New York Times*, 5 March, A1.

— (2017), 'A Guide to Challenges Facing Self-driving Car Technologists', *New York Times*, 7 June, https://www.nytimes.com/2017/06/07/technology/autonomous-car-technology-challenges.html?_r=0

Oliver, J. E., and Rahn, W. M. (2016), 'Rise of the *Trumpenvolk*: Populism in the 2016 Election', *ANNALS, AAPSS*, 667, September, 189–206.

Pew Foundation (2014), 'Political Polarization and the American Public', 12 June, available at http://www.people-press.org/2014/06/12/political-polarization-in-the-american-public/

Phillips, E. E. (2015), 'Truck Driver Turnover Rate Slides Down Sharply', *The Wall Street Journal*, 14 July, https://www.wsj.com/articles/truck-driver-turnover-rate-slides-downward-sharply-1436908863

Polanyi, M. (1967), *The Tacit Dimension*, New York, Anchor Books

Presidente, G. (2017), 'Labor Services at Will#8232;Regulation of Dismissal and Investment in Industrial Robots', paper presented at NBER Summer Institute, 15 July.

Remus, D., and Levy, F. (2017), 'Can Robots Be Lawyers? Computers, Lawyers and the Practice of Law', *Georgetown Journal of Legal Ethics*, **30**, 501–58.

Rodrik, D. (2017), 'Populism and the Economics of Globalization', working paper, J. F. Kennedy School of Government, Harvard University, June.

Sanchez, E. A. (2017), 'Volvo Testing Autonomous Garbage Trucks in Sweden', *Truck Trend*, 26 May, http://www.trucktrend.com/news/1705-volvo-testing-autonomous-garbage-trucks-in-sweden/

Saulwick, J. (2015), 'Sydney's Patrick Terminal Goes Automated, with Fewer Staff but Dancing Robots', *Sydney Morning Herald*, 18 June, http://www.smh.com.au/nsw/sydneys-patrick-terminal-goes-automated-with-fewer-staff-but-dancing-robots-20150617-ghqc24.html

Smith, A., and Anderson, M. (2017), *'Automation in Everyday Life'*, Pew Foundation, 4 October, http://www.pewinternet.org/2017/10/04/automation-in-everyday-life/

Stern, K. (2017), 'Inside how Trump Won the White Working Class', *Vanity Fair*, 5 January, http://www.vanityfair.com/news/2017/01/how-trump-won-the-white-working-class

Timberg, C. (2016), 'Russian Propaganda Effort Helped Spread "Fake News" During Election, Experts Say', *Washington Post*, 24 November, https://www.washingtonpost.com/business/economy/russian-propaganda-effort-helped-spread-fake-news-during-election-experts-say/2016/11/24/793903b6-8a40-4ca9-b712-716af66098fe_story.html?utm_term=.16b26ddc8eaa

Townsend, M., Surane, J., Orr, E., and Cannon, C. (2017), 'America's "Retail Apocalypse" is Just Beginning', Bloomberg.com, 8 November, https://www.bloomberg.com/graphics/2017-retail-debt/

Williams, J. C. (2016), 'What So Many People Don't Get About the US Working Class', *Harvard Business Review*, 10 November, https://hbr.org/2016/11/what-so-many-people-dont-get-about-the-u-s-working-class