December 20, 2016

Projections about Artificial Intelligence and Employment

Roxanne Heston
Tulane University

Abstract

Technology has a fortunate history of improving welfare. Despite the fears of generations past, new waves of technology have historically displaced some workers but on average maintained or even grown the employment rate, complementing rather than replacing human labor. There is, however, an active debate about whether we should expect this trend to continue, as an increasing number of economic and computer science scholars believe that future technological development, in the form of artificial intelligence, is likely to break the fortunate trend.

In this paper I briefly describe artificial intelligence and consider both sides of the debate. Then, taking the pessimist's view of job displacement, I try to lay the framework for quantitatively predicting the impact of increasing automation on the labor market. I borrow from prediction science literature, focusing on describing the scope, meaning, and timeline of the technology's development. In this context, this describes the kinds of skills of which artificial intelligences are likely to be capable, how these skills appear in jobs, and when such technologies are likely to exist and become adopted in the workplace. I then break each section into the five factors the McKinsey Global Institute deems necessary conditions for the adoption of artificial intelligences in the place of humans, and discuss them in turn. I conclude by applying the framework to the

example occupation of wait-staff, making predictions while voicing reservations about the predictions and areas for further investigation.

Introduction

What exactly will be the effects of artificial intelligence development on employment? At what scale and rate should we expect the displacement to occur, and in which sectors? While there is quite a lot to discuss about whether artificial intelligence will cause mass (or total) unemployment, the aim of this paper is not to determine this conclusively. Instead, taking technological mass unemployment to be the assumed future, the goal is to start to determine how to make concrete predictions about how it might unfold, laying the framework for evidence-based, critiqueable arguments in a domain largely dominated by anecdotes and broad-strokes claims. I hope that making predictions about large-scale unemployment will quickly prove the strong claims wrong, showing them to be unwarranted fear in the face of a beneficial technological development. However, until this is known to be the case, it is imperative to test if there is truth to these claims. This paper begins to help us prepare us for that potential undesirable outcome.

A Brief History of Physical Labor Mechanization

With every wave of technological innovation comes a new wave of excitement and fear about how it will help and hurt the job market. We know that technologies often create frictional and short-term structural unemployment, but tend not to affect, or even increase, unemployment rates in the long run (Autor, 2015). Technologies often allow workers to specialize, decreasing the amount of resources it takes to maintain production or "increasing the size of the pie" by augmenting their total production. ATMs may have taken bank tellers out of the cash withdrawal business, but, rather than eliminating the bank teller job, simply reallocated their time and efforts to financial planning and management functions, which are better performed in person-to-person interactions (Chace 52). Textile machines produce clothing far more quickly than any human could at a loom, increasing the quality and availability of clothing each worker can produce (Chace 11). While businessmen may shift capital between investments to better allocate what exists, entrepreneurs like Peter Thiel and Bryan Johnson believe technological innovation is the key to creating a post-scarcity economy and better world (Thiel; Dwoskin). We have witnessed astounding improvements in the human condition from previous feats of the Industrial Revolution, with the advent of a host of transportation and production innovations that have give even the poorest people in America better quality of life than royalty just a few centuries ago.

For instance, from 1870 to 1980, increased mechanization of the agricultural industry reduced the need for human labor to less than a tenth of its initial level, from 50 percent of the US labor force to just four percent, and increased the scale of production per laborer 14-fold (Daly, 1981). Much of this scaling is attributable to the introduction of machines like grain combines and large four-wheel-drive tractors, which perform functions once left to hired farm laborers. This has readily improved the availability and stability of food for a once hand-to-mouth population.

Reducing the need for many humans to perform physical labor also allowed them to divert their efforts to make use of their mental capacities, which often helped to further the scale and spread of production. In the mid-20th century millions of former farm laborers found themselves displaced as machines like the cotton harvester, tomato picker, and the like took over the jobs these laborers once performed (Daly, 1981). From the beginning to the end of the 20th century, the percent of the population working as farm labor decreased from 38 percent to just three percent (Fisk, 2001).

In circumstances like these, some displaced workers find new jobs in the same industry, while others move into different sectors, and other still have difficulty finding new jobs at all (Displaced Workers Summary). While some people remain in physical labor roles today, an ever-growing number are primarily putting to use their cognitive abilities as their means of employment, switching from production to technical and professional occupations (Fisk, 2001).

Fortunately, it seems that despite the rapid mechanization of jobs, the average unemployment rate has remained roughly stable over time, a fortunate side effect of the complementary nature of mechanical production and the so-far "uniquely human" skills we provide (Autor, 2015).

Prediction Regarding Intellectual Labor Mechanization

Today, less than one percent of the US labor force works in agriculture, thanks to the advent of technologies that perform intelligent functions. These so-called precision technologies — such as the GPS autosteer, automatic yield monitors, and variable rate technologies — help farm owners better perform some of the oversight and management functions that comprise much of the remaining work for those in the agriculture business, producing yet another reduction in costs and increase in yields (Schmidt and Moss, 2015). It is safe to view the improvements on the farm as emblematic of those in other sectors, as technological improvements have fundamentally changed how we do other production today.

With the allure of such potential welfare in mind, modern technologists chase the next big thing in technological development: artificial intelligence, or "AI." Deemed anything from "the fourth industrial revolution" (Schwab) or "second machine age" (Brynjolfsson and McAfee) to a technological revolution of a whole different kind (Chace 15), these technologies, much like the precision technologies mentioned above, seek to perform human functions at human-or-greater ability. In the way that the prior industrial revolutions outmoded the need for animal and human physical labor in many industries, AI shows promise to do the same for mental labor.

As machines start to become able to read facial expressions, diagnose illnesses, and even write music (Chace 37-45), it is becoming decreasingly clear which skills, if any, will remain uniquely human. Jobs we never imagined a machine could be capable of performing are becoming automated in droves, begging the question as to whether humans will have a use in the workplace in the long-run, and if so, in what capacities.

Debate About Artificial Intelligence Rendering Human Labor Moot

Some artificial intelligence researchers only see time as the limiting factor on almost complete automation (Ford; McAfee and Brynjolfsson; Susskind and Susskind), while others believe that, in due time, we will encounter another line of retreat, a set of skills that are fundamentally non-automatable and therefore invulnerable to the technological trend (Autor, 2015; Hanson 8; Cowen 5). AI researcher and science fiction author Calum Chace explains his skepticism of the latter claim in his book *The Economic Singularity*. Here he argues that humans are reaching "peak human" and are soon to be unnecessary for production, much in the way that horses became unnecessary on farms and in cities in the 1900s (Chace 17). As there is nothing that intrinsically guarantees that new technologies create as many or more jobs than they take, Chace and others claim that the complementaries between man and machine thus far are a temporary aberration, unreflective of what we should expect going forward.

This is, admittedly, a rather strong claim. There are many reasons to believe that we humans will continue to remain employable, particularly if one projects from historical precedent. The machines that exist today are far from able to take over the roles many humans serve. As Chace concedes, current artificial intelligences primarily perform narrow functions, require explicit instructions, make many more mistakes than humans in domains where accuracy is key (Autor, 2014), sometimes face pushback from unions and users alike, and often are most productive when integrated with humans rather than when performing on their own (Chace 22). However, all of these arguments demonstrate temporary resistance to complete automation in the short run, holding far less weight when applied to the next 50-100 years.

It is fair to observe that the potential for technology to perform human functions seems unlikely to follow from its former and current capabilities. However, market demand for cheaper, easier production is high, and history shows that market demand tends to overcome social resistance in the long-run. That trend, I believe, is reasonable to expect to continue.

The demand, it seems, is likely to be there. Artificial intelligences are expanding in functional capacity, both through man-made improvements in their range of applicability and accuracy of execution, and on their own with self-teaching “deep learning” algorithms. The simultaneous improvements make them able to perform more autonomously, relative both to prior technologies and even to their human counterparts. The last edge we are likely to have over machines is our humanness itself, and even that is coming to seem less and less crucial in both the goods and services industries, viewed as a luxury good for those who can afford it rather than a component of mass-market demand. While machines may be expensive to adopt at first, they are almost always cheaper in the long run, more consistent once built well, and open up the potential for a greater number of beneficiaries (Chace 15-19).

With time, we should expect the social barriers and hiccups to disappear, and, like it or not, for the overwhelming profit incentives to overpower concerns of the risks to personal autonomy and malicious hacking. From British Luddites smashing textile machines to taxi unions restricting ridesharing companies, history demonstrates that popular protest is little more than a temporary hindrance to the stronger forces of supply and demand. Human pushback, it seems, is likely to do little more than slow the technology's rate of adoption and showcase the ever-growing tension between human- and machine-powered production (Chace 20-22).

Structure of the Discussion

While there is certainly more to discuss on both sides of this debate, the aim of this paper is not to determine conclusively whether artificial intelligence will cause mass (or total) unemployment. Instead, taking technological mass unemployment to be the assumed future, the goal is to start to determine how to make concrete predictions about how it might unfold, laying the framework for evidence-based, critiqueable arguments in a domain largely dominated by anecdotes and broad-strokes claims. I hope that making predictions about large-scale unemployment will quickly prove the strong claims wrong, showing them to be unwarranted fear in the face of a beneficial technological development. However, until this is known to be the case, it is imperative to test if there is truth to these claims. This paper begins to help us prepare us for that undesirable outcome.

Given this aim, I will start along a path of setting the framework for figuring out the magnitude, if any, of AI-induced technological unemployment (AIITU). Should these fears be warranted, I think this is important for making tractable efforts to mitigate the harms of mass unemployment. Policymakers are unlikely to know how to move forward without knowing what to expect. If technologists fail to provide economic models substantiating their fears, those with the power to prepare for a very different future will be at a loss for what policies they could or should pursue. Predictions should also help to improve the caliber of the conversations in this domain. Where the prevailing arguments in this field are vague and grandiose, precise predictions with articulated reasoning make claims refutable. In this paper I will look to the works of economics researcher Robin Hanson, FiveThirtyEight blogger Nate Silver, and *Superforecasting* author Philip Tetlock, amongst others, who advocate for broader adoption of the practice of making predictions that one can critique and then prove wrong or right as time progresses (Hanson 5; Silver; Tetlock and Gardner).

If predictions are so valuable to researchers and policymakers, it begs the question of why people have not made quantitative predictions about AIITU already. Despite it being a new field, AIITU has received a fair amount of interest, including much speculation about its likely effects. Chace dedicates a whole section of his book to laying out a possible scenario for the coming decades, but intentionally shies from concrete projections about which types of jobs will be outmoded and at what scale. In fact, Chace prefaces his timeline by noting that it is simply a rhetorical device that "may in some small way help us to construct a valuable body of scenarios." I believe he speaks for many future-thinking researchers when he notes that he chose not to write in the equivocal manner of a formal predictor because, despite their superior forecasting power, they "are often not the people who get listened to in discussions about the future" (Chace 99).

Similarly, Hanson makes predictions in his book *The Age of Em* about a future dominated by technological sophistication, yet remains anecdotal and focused on a narrow and more engaging part of that future. He makes it clear that his aims are to spur interest in prediction-making, and to create emotional salience around a future vastly different from the present (Hanson 1-3). In most other contexts experts and pundits skirt discussing concrete predictions entirely, undoubtedly worried about putting their reputations on the line when those who do find themselves subject to hefty criticism for inaccurate predictions.

So despite the incentives otherwise, I will lay out a framework for making critiqueable predictions here, then attempt to make a prediction myself using the outlined model. I do so knowing that they will be far from perfect, but that it is the only way to figure out in which ways our projections are systematically right or wrong, and make plans based on a more accurate model.

Components of a Good Prediction

According to Philip Tetlock, “the meaning, scope and timeline must be clear for meaningful forecasts capable of being tested for accuracy” (Tetlock and Gardner, 2016). The meaning of predictions is often less clear than one might assume at first glance. It is easy to see the need for conceptual clarification in the intelligence sector, for instance, as Tetlock found in one study that there are up to 60 percentage-point spreads in the probabilities intelligence agency personnel assign to words like “very likely” and “probable” (Barnes, 2015).

Rigorous researchers may define terms, but with loose interpretability matching their rightfully high level of uncertainty. For instance, Frey and Osborne define a job as being at “high risk” of unemployment if “associated occupations are potentially automatable over some unspecified number of years, perhaps a decade or two” (Frey and Osborne, 2013). Similarly, Google’s CEO Larry Page stated in an interview that “computers will be better-suited to take on most jobs” (Eadicicco). It is unclear what he means by “better-suited” (better-suited in skill alone, or also cheap enough to be used?) and by “most jobs” (just over half, or almost all jobs?). Famous technology predictor Ray Kurzweil gestures at “major social unrest” from AIITU and “major danger to humanity generally,” although what exactly he means by that is left to the reader to guess (Kurzweil). With this sort of rhetoric, it is entirely plausible that those who claim that AIITU will be huge and disruptive, and those who believe it will be much like what we have seen in the past, may actually be in greater agreement than they might seem.

Lack of concrete predictions also has repercussions on what timelines we should anticipate for certain developments. The game of Go used to be often referenced as the paradigmatic example of a game that is impossible to computerize, given that there are more possible moves than atoms in the universe. Researchers around the globe were shocked when, in 2014, Google Brain’s AlphaGo beat human world champion Lee Sedol, just hours after first encountering the game (“A Game-Changing Result”). While superior prediction techniques would not make genies of humans, it would certainly facilitate the filtration process of super-predictors from rhetorical pundits. Disparity of timelines is at worst detrimental for business competitors at this stage; in later stages, we should be worried that such mis-anticipations can bring about massive market failure.

Prediction Framework

Now that I have contextualized the field and importance of concrete prediction-making, I will begin to make a prediction of the kind someone can refute, given little fear and much anticipation of being wrong, both in my methodology and predictions themselves. My goal is to spur others to rise to the challenge with me. I hope others will join me in establishing a concrete basis, albeit back-of-the-envelope and insufficient, for calibrating our expectations to the coming shifts in the labor market, and thereby to the population at large.

As I go through this example I will explain the framework I am using. I am intermingling two models to create one, comprehensive structure for making AIITU-related predictions. The first is Tetlock's three components of a good predicting, each of which I talk about in turn:

1. The kinds of intelligences that exist, or "meaning";
2. The jobs have those intelligences and how many people have those occupations, or "scope"; and
3. The dates by which we should expect to be able to automate each of them, or "timeline."

I have then subdivided each of those sections into accordance to the five factors the McKinsey Global Institute (MGI) articulates as being necessary conditions for the adoption of artificial intelligences in place of humans.

Meaning

The first factor I will talk about is the kinds of intelligences used in different forms of employment. As has been found with very simple "intelligent" technologies, a machine does not need to have full human capabilities to replace a human, just the ones that the human uses on the job. *Frames of Mind* identifies nine types of human intelligence: linguistic, logic-mathematical, musical, spatial, bodily, interpersonal, intrapersonal, existential and naturalistic (Gardner). These seem like a reasonable way in which to break up job automation.

In fact, some of these skills, particularly mathematical and spatial ones, are already automated to a reasonable extent, while others, particularly those humans find to be the simplest or most intuitive, are seen to be the most difficult to automate (Chace 52-54). Research has shown that currently we are witnessing technological unemployment of the middle-skill jobs, once with repetitive tasks like secretarial work, some forms of accounting, scanning through vast troves of medical records or legal cases, or accounting (Chace 56-66; Autor, 2015). Low-skill jobs have physical intuitions, while high-skill jobs require abstract abilities, both of which Autor believes will be harder to impart into machines (Autor, 2015).

However, to ease the process of evaluation, I will use the more granularly segmented model developed by the McKinsey Global Institute, which I will describe later.

Scope

The relevance of each of these skills only goes so far as each job uses them. Chace notes that what constitutes an "essential function" of an occupation is at minimum complex, if not impossible, to disentangle, since jobs may explicitly require one skill but in fact require many (Chace 52). For instance, you are unlikely to find "spatial and interpersonal ability" listed in the qualifications required of doctors, but both are obviously important for the job. Another caveat that makes this difficult is that just because a job requires some skill now does not mean that the person who has that job is going to lose their jobs as soon as a machine acquires that skill. Using the example from earlier, a narrow AI — that is, one that just performs one function, rather than a suite of complementary functions — might have banking skills that make it a good ATM machine, but not the interpersonal skills necessary for financial management, so bank tellers simply pivoted their functions, rather than being outmoded (Chace 74). The McKinsey Global Institute does not use Gardner's intelligence breakdown but does go into more usable and greater detail, breaking up jobs into "2,000-plus work activities for more than 800 occupations" (Chui, Manyika, and Miremadi).

Since historically there are many examples of jobs with multiple required skills continuing to exist in an augmented form, I will make my predictions more modest by assuming that in any occupation where multiple skills are seen as essential to the job will not go out of existence until all of the skills have become automated, and that employees will be retained in the same ratio as the percent of the job left unautomated. For example, if 75% of an occupation becomes automatable and the rest of the factors necessary for AI adoption are in place, I would assume that 25% of the original employees would be retained and that all of their time would go to the remaining 25% of their former occupational activities. This, of course, is unlikely to perfectly match how employment works in practice, where employee roles change and take on new responsibilities, or employers decide to keep employees from profit-unrelated reasons.

Timeline

We should be skeptical about predictions of this sort for many reasons. Forecasters who are not held to their predictions tend to do pretty poorly in making predictions. Machine-intelligence research Katja Grace talks about some of the systematic biases she has encountered in looking at AI development predictions, whereby people predict things that look eminent to occur about ten years from now, and things for which there is no foreseeable path of accomplishment but is likely to happen at some time in the unforeseeable future to be about 50 years from now. In fact, expert predictions are not that different from those of lay people, implying that the special information is not novelly informing their projections all that much. It turns out that for the past half century, experts in the field have thoughts that the same certain developments in artificial intelligence were 20-50 years away from whenever they made the prediction — which, of course, should not hold over time if the predictions are meaningful. Without that gap closing, we should feel quite uncomfortable taking literally the predictions we read, even if they are being made by experts in technological development (Grace, "Update on all the AI predictions"). There also seem to be strong correlations between the professions of those who make predictions and their timelines, with computer scientists being more skeptical about the rate of development than future-focused researchers (Grace, "Group Differences in AI Predictions").

Even when the technologies are developed, there are reasons to believe that there will be delays in adoption. In *The Economic Singularity* Chace discusses the variety of lengths of business planning cycles, wherein companies in rapidly moving sectors like technology may make large changes to their plans and infrastructure on an annual or semi-annual basis, where "lumbering giants” in e.g. law, manufacturing, and finance may only change course every five years or so (Chace 74-83). The lag that the latter industries create are clear from McKinsey’s findings; according to their research, 45% of existing jobs could already be automated away (Chui, Manyika, and Miremadi). Similarly, Frey and Osborne deem 47% of US jobs to be at risk of automation based on their assessment of current or upcoming technological capabilities and the functions that current workers serve (Frey and Osborne).

Of course, business planning timelines are far from the only source of delays in technological adoption. As I hinted at earlier, the MGI articulated five things that they believe to be necessary for AIITU. Having clarified the definitions of “meaning,” “scope,” and “timeline,” I will finally progress into describing what exactly these criteria are, contextualizing each one as applied to a prediction about the automation of a specific profession: that of a waiter or waitress.

Framework

In accordance with the McKinsey Global Institute’s brief description, I will consider AIITU of a specific job to be likely to happen when the following are true:

1. It is technically feasible. This means that a technology exists that has all of the major skills relevant to the job. You are more likely to find this for industries that operate on short timelines and re-evaluate their processes often, as AI companies have less incentive to create the capabilities certain industries could theoretically use if those industries will take years to adopt them.

2. The costs of automation make doing so both possible (not cost-prohibitive) and worthwhile. Drawing from classical labor economic theory, you are more likely to find labors substituted by capital in industries where the technology costs less than the workers who perform those functions, and in industries where the costs of labor are high relative to other inputs.
3. The workers who might otherwise do the activity are scarce, skills are hard to find or difficult to train, or costly. Again, this is likely to be the case in companies in which employees make up a large fraction of the company's costs. While this might not have always been true, the recent polarization of jobs into primarily low- and high-skill jobs with few in between (Autor and Dorn, 2013) means that companies are likely to want to automate jobs that require higher-level education or change too rapidly for a human to keep up.
4. The benefits of automation are beyond the labor-cost substitution. While this has some overlap with (2) and (3), this specifically points at the things that make technology not just sufficient but actually preferable to humans in certain contexts. This is going to be common in fields where precision, creativity, or "the human touch" are less important than consistency, speed, and endurance.
5. Automation will be accepted socially and permitted regulatorily. This is likely to depend a lot on current expectations for the job and the dispositions of the people in the area in which the technology is debuted. This, perhaps more than any of the other factors, suggests that a job may be entirely automatable and even become entirely automated in some regions but not in others. For instance, cities lacking strong unions associated with

the endangered profession are more likely to see automation, or cities where there is a clear line of retreat for those employees.

There are a few more caveats I would like to add. It seems important to distinguish between an activity going out of human execution and a job going out of existence. As I mentioned in the "Scope" section, I will consider it necessary for all activities of a job to disappear before the job itself disappears. Should some of the tasks remain un-automatable, or the other factors fail to align with automating some of the tasks, I will assume that the job will just take another form comprised of the remaining activities.

I will also distinguish between a job going out of existence and a human becoming unemployed. This is a harder thing to anticipate, as there are too many unknowns to anticipate whether everything will be in place for an employee in Industry X to transition to Industry Y. However, I will at least maintain the concept that a person will be unemployed "forever" if the other skills they are likely to have have also been automated away and/or there is insufficient demand for those jobs.

Application

Let us now apply this to a specific job, by looking at the meaning, scope, and timeline of AIITU in the context of the waiter/waitress job.

Meaning

The meaning of "technical feasibility" here is already rather complex. This is a rather manual labor-heavy task, including functions like arranging and assisting tables or dining areas, and collecting dirty dishes or other tableware (McKinsey Global Institute). A machine would have to be able to do all of these functions in order for this job to go away completely. "Ability" here encompasses a few things: not just accomplishing the task, but also as well as or better than a human. We already have evidence that this is necessary; causal restaurants like Numero Uno Pizzeria, Chili's, and Applebee's already employ hybrid-model wait-staff, with table-top electronic "waiters" called Ziosks at every table taking orders, but retain live waiters running over to delivery orders to tables and doing post-visitor cleanup (Vanek Smith and Smith).

The costs for automating some of their functions are low — just a Ziosk per table and its electricity, as well as additional revenue from faster service — but so are the wages of their employees (see Figure 1). The costs of automating the rest of their functions would have to be lower than the extra revenue generated in order for the jobs to disappear.

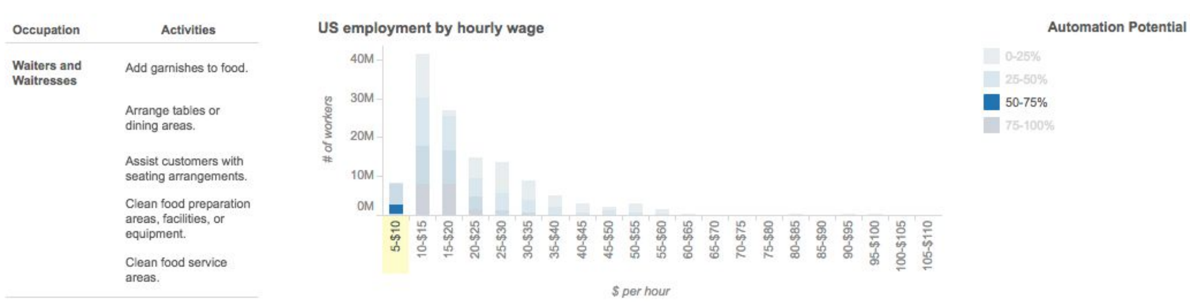

Figure 1. The activities, hourly wage, and automation potential of waiters and waitresses. Source: McKinsey Global Institute's "Automation and US Jobs" Interactive Infographic

The benefits beyond the direct substitution of existing labor is their superior sales abilities; some customers report preferring ordering on a machine, since it “doesn’t judge them for ordering dessert after a large meal,” leading to a 30% increase in dessert sales. Additionally, machines process orders more quickly than human servers, getting customers in and out fast enough to reduce the time each customer spends at the restaurant by about 10 minutes (Vanek Smith and Smith).

Regulations in this space are low as there is little unionization for high turnover jobs like waiting tables, and social acceptance of technologies in restaurants is high, at least in the cosmopolitan cities in which the technology has been partially introduced.

Scope

Next, I shall estimate the scope of AIITU for waiters and waitresses.

The scope of technical feasibility for this job is already quite large. According to McKinsey, 64% of the activities a waiter performs are automatable with existing technologies (McKinsey Global Institute). The things that AIs cannot yet do but humans can easily do are those that are unpredictable and environmentally sensitive. Machines with human-level communication abilities and the ability to navigate complex, sensitive environment — e.g. clearing up a variety of dirty dishes — would make the job nearly moot. Here it seems safe to assume that, given the flexible hours, requirements, and number of waiters, waiters will continue to be hired for the fractional time of the remaining, un-automatable activities, serving in just the remaining functions. Since 64% of their jobs are automatable now, I will estimate that the industry would continue to retain 36% of their waiter hours were there to be no further technological developments. With 2,505,630 people in the industry as of May 2015 (“Waiters and Waitresses”), we should have already seen that number drop to 1,603,603.2 people-equivalents (64% of total wait-staff jobs) were all of the other necessary factors in place.

Since fewer than 64% of waiting jobs have been automate, it seems that the rate-limiting factors are beyond technical feasibility — in particular, the cost of automation. According to the Bureau of Labor Statistics, waiters were paid a median wage of $9.25 as of May 2015. The range of costs correlates with what is expected of them; waiters in the bottom 10% bracket, making $8.08 per hour, are paid primarily for performing their job functions, where those in the 90th percentile, $17.14-per-hour bracket have both a broader range of expectations and higher pay for their “personal touch” (“Waiters and Waitresses”). However, some of this pay comes in the form of tips, costing the restaurant between $2.13 per hour under the Fair Labor Standards Act up to $10 in states like Oregon and California (“Minimum Wages for Tipped Employees”).

Training costs are low, since waiters and waitresses usually require little training, but recently benefits have increased in prevalence as it has become more difficult to find the employees they need. A restaurant that offers health insurance benefits ("Why Is It Getting Harder to Hire in the Restaurant Industry?"), just 14.4% of restaurants, can expect to pay $536 per month for single workers and $1,512 per month for families ("2016 Employer Health Benefits Survey”). While the cost of the Ziosk varies greatly depending on the scale of adoption, I estimate the price as being that of a tablet plus the subscription fee for using similar software services. I have combined a simplified model of the relevant costs in Figure 2 and tentatively determined that restaurants who adopt the technology already stand to almost double their monthly revenue by adopting this technology, even without the cost benefits to employers of replacing workers (See Figure 2). But since the technologies can only serve the 64% of non-automatable jobs, we should only expect this to affect 64% of waiter/waitress jobs.

REVENUE	Dessert	Drinks
Average rate of purchase[1]	13.84%	70.00%
Increase in sales using Ziosk (gender- and BMI-ambivalent)[2]	30.00%	30.00%
Projected percent increase in average rate of purchase	4.15%	21.00%
Average percent of meal revenue[3]	24.39%	14.63%
Projected percent increase in meal revenue	1.01%	3.07%
Total projected percent increase in meal revenue	4.09%	
Revenue from average meal[4]	$20.50	
Total projected dollar increase in meal revenue per meal	$0.84	
Total projected dollar revenue from average meal	$21.34	
Patrons served per table per day[5]	5	
Projected percent increased average rate of patrons served[6]	120.00%	
Total projected increased average rate of patrons served per table per day	6	

[1] Computed using figures from Döring and Wansick, 2015 on the effects of BMI and gender on dessert and drinks purchases, and the Bureau of Labor Statistics' gender breakdown in the industry.
[2] As discussed by the restaurant owner in "I, Waiter," with Vanek Smith and Smith.
[3] A fraction of "Revenue from average meal," using $5 and $3 as the average dessert and drinks prices, respectively, as listed in Pavesic, 2005.
[4] The average of the $16-25 range in which roughly one-third of consumers said they spend when they eat at a casual restaurant, according to Deloitte.
[5] Using Robson's estimated five table turns per day in small-to-medium restaurants.
[6] As discussed by the restaurant owner in "I, Waiter," with Vanek Smith and Smith.

Average days per month open[7]	30	
Total projected increase in patrons served per table per month	180	
Total projected dollar increase in meal revenue per meal	$0.84	
Total projected dollar increase in meal revenue per table per month	$150.77	
EXPENSES		
Tablet upfront cost[8]	$499.00	
Tablet lifetime (in months)[9]	24	
Tablet cost per month (average)	$20.79	
Restaurant software cost per month[10]	$59.00	
Total tablet and software cost per table per month	$79.79	
NET INCOME (assuming no worker displacement)		
Projected revenue increase with tablet	$150.77	
Less projected cost increase with tablet	$79.79	
Projected net income per table per month with tablet	$70.98	

Figure 2. Simplified model of projected additional profits per table for restaurants that adopt table tablets, assuming employment remains constant.

Now I will look at the relative scarcity, skills, and cost of workers who might otherwise do their activities. As I already noted, workers in this industry are rather replaceable. They are not paid a lot because they do not require lots of training or skills that are difficult for humans to learn and perform. Save for a few, high-end locations, restaurants often employ students and part-time employees to perform most of the functions articulated in the waiter/waitress job description. I therefore do not expect this to be a large limiting factor in technological displacement, and so

[7] As averaged from the Numero Uno Pizzeria, Applebee's, and Chili's websites, three large casual dining chains that have already adopted table tablet technologies.
[8] Using the price of a 128GB iPad Air 2 purchased directly from Apple.
[9] Using two years, the low end of the two-to-three-year projected average lifespan of an iPad according to Gartner's Market Trends report, since they state that iPads become slow after two years and these are likely to receive heavier-than-average use.
[10] Using the price of a monthly subscription to CAKE Corporation's guest manager software that serves functions like those articulated in the paper.

will only assume that this hinders adoption for jobs with high volumes of wait-staff employment and higher skill in preparation: those at drinking places and doing special food services, who comprise 17.43% and 9.99% of waiting jobs, respectively ("Waiters and Waitresses"). Bartenders may attend bartending school but usually must do so on their own dime, learning some on the job in informal training. I expect that most of the difficulty for employers is in finding those with the necessary skills, rather than the cost of training them themselves.

The benefits, as I listed above, are increased efficiency of taking customers, and increases in food purchases. If the 10-minute increase in efficiency experienced by Numero Uno Pizzeria scales to other restaurants who adopt the technology without compromising pay, we should expect a 20% increase in revenue, reducing the average table turns from 50 minutes (Tobin and Huffman, 2006) to 40. Additionally, if, as the pizzeria employees reported, a restaurant has a 30% increase in dessert and drink sales (Vanek Smith and Smith), comprising 5-10% (Pavesic, 2005) of the counterfactual $16-25 average order amount (Deloitte), I will expect that same increase in additional revenue.

Finally, the scale of regulatory and social acceptance will be, I believe, primarily location- and caliber-dependent. I expect it to be socially accepted in all lower-caliber restaurants where personal service expectations are lower or speed of service is priority. This is substantiated with a finding of the Toast, Inc. who, in their Restaurant Technology in 2016 Industry Report, reveal that "79% of diners agree that restaurant technology improves their guest experience" (Toast). Needless to say this is a rather broad statement, encompassing things like Yelp reviews and credit card machines, but it is some indication of openness to its broader adoption. Regulation, then, seems to pose far more of a challenge. According to the 2016 Restaurant Operations Report, restaurant operators list “government” as their primary challenge, but with regards to regulations not directly affecting their ability to adopt technologies like the ones in question here (Deloitte and National Restaurant Association).

Timeline

Finally, I will determine at what rate I expect these hurdles to automation will be overcome.

For the 36% of activities not already technically feasible, expectations are reasonably difficult to anticipate. Pieter Abbeel, the creator of the “Berkeley Robot for the Elimination of Tedious Tasks” believes that machines with fine motor movement are 20-30 years away (Henn). Since his towel-folding machine still takes four minutes to fold a towel, this would not count as “technically feasible” for replacing a laundry job. His machine attempts to do many of the things that are as-of-yet un-automatable in the waiter/waitress job, so I will apply his predicted year range to the 36% of the remaining tasks.

Now I will explore when I expect the costs of automation to be sufficiently inexpensive for it to become widely adopted, given that the other factors are in place. Given the calculation above for the costs of a machine relative to a human on currently automatable tasks, I will take the rates of the real price declines in technology and compare them to our expectations for real wage increases for waiters and waitresses. With the exception of cable, satellite radio, and cable television, the real prices of technology have been decreasing fairly consistently since 1997, a trend we can expect to continue (“Long-term Price Trends for Computers, TVs, and Related Items”). Since the aforementioned table tablets have many of the components of typical consumer technology — as it is much like a personal tablet and attached credit card reader — I will assume their prices, too, will decrease. Furthering the incentive for technological adoption, while real wages have been declining relative to the marginal increases in labor productivity (Imbery), other costs like benefits and recruitment, as well as a reduction in consumers (Braverman), negate those effects, making labor more expensive on net (Green and Williams, 2010).

As for the availability of the necessary labor, restaurant owners are finding workers increasingly scarce ("Why Is It Getting Harder to Hire in the Restaurant Industry?"). Should this decline continue and wages continue to grow to compensate, I will assume the two effects will combine to lead to near-total technological adoption and labor replacement with existing technologies between 2020 and 2025.

As for when the benefits of automation will be beyond the labor-cost substitution, this, as I have explored, is already the case in some circumstances. The important thing to predict here is when their detriments will no longer be prohibitive. Calum Chace notes that adoption of technological substitutes is likely to be slower here than in some other industries where faulty technologies are more forgiving. Since high rates of accuracy are very important in the food services sector, I will refer to the estimate I made earlier, when the technology will exist at sufficient speed and accuracy, as my estimate here as well.

Lastly, I will consider the technology already sufficiently socially accepted to be adopted as soon as the other factors are in place. It may take a lot long, if not indefinitely, for the vast majority of the positions to receive social acceptance, as restaurant patrons may in fact be paying for "the human touch" at higher-end locations. I will assume otherwise, though, as one might have argued the same of coach-and-buggy services, hotels, etc. I can look at the rate at which people have accepted technologies doing social functions in the past. For instance, despite ATMs existing for commercial use since 1969 (History.com Staff), they only have the first ATM installed in the city center of Yangon, Myanmar, not for lack of need for secure monetary transactions and storage but because of fear unreliability, that the money will not be returned consistently (Thuy Vo and Smith). This could downgrade my estimate of when the technology will be adopted, but, given that I am making a prediction about US-based employment, does little to sway my prediction.

As I have said before, all of the factors above need to be in place for the technology to be adopted. However, adoption can happen variably across the sector. For instance, if all of the pieces are in place but the fixed costs of adoption are high, big companies might adopt the technology — as, in fact, they already have — while small companies may continue to hire humans for longer. Alternatively, as we have already witnessed with car sharing services, some cities, both the people and their regulatory bodies, may be more amenable to automation than others, and therefore adopt the technology accordingly.

Concise Prediction Statement

To stitch a rather large set of considerations together, let us now consolidate the meaning, scope, and timeline of AIITU for waiters and waitresses to get a prediction one can critique and refute.

Given that I believe economic benefits and business cycles to be the limiting factor on adoption of the extant automation technologies, and technical feasibility for the remaining factors, I expect 64% of waiter/waitress jobs in the US to be gone in the next 5-10 years, and most of the remaining 36% to be gone in the next 25-35 years. Given that people are willing to pay a premium for human service in some contexts, I expect 5-10% of the jobs otherwise projected to exist along current growth trajectories to remain, principally in high-end dining facilities.

Conclusion

This paper tackles just a sliver of the challenge we have at hand. There are hundreds of occupations in the United States and many more around the globe, many or all of which stand to witness dramatic effects to their viability in an era of ever-growing automation. I hope that others will take this model and approach as a basis for investigating and predicting what each industry will experience, either running with the model I have articulated or crafting their own in light of my mistakes. I believe that the three-factor model of clear prediction-making, combined with the five necessary components for technological displacement, create a solid initial basis for testing our best guesses as to how the next few decades will unfold.

Specifically, I would like to challenge the US Department of Labor to accept and grapple with the implications of artificial intelligence before it has a chance to have dire economic consequences. I also believe it to be in the best interest of sector-specific associations and guilds to apply this model to their individual fields, to curb the repercussions of rapid unemployment should that be the default effect of AI development if we fail to act. I hope to see these sorts of finding published publicly and transparently, as once-independent parts of our economy become increasingly interdependent and governing bodies increasingly powerful.

I hope, as I am sure everyone does, that artificial intelligence only ends up augmenting society's production and quality of life. Unfortunately, ignoring the potential risks or leaving their magnitude and timing to guesswork does not make them go away. The more quickly powerful stakeholders accept that the risks this technology might pose, rather than simply looking forward to the benefits it offers, the better equipped we will be to mitigate its downsides should they come to pass.

Works Cited

Autor, David H., and David Dorn. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103.5 (2013): 1553-597. *MIT Economics*. Web. 20 Dec. 2016.

Autor, David H. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29.3 (2015): 3-30. *MIT Economics*. Massachusetts Institute of Technology, Department of Economics. Web.

Autor, David. *Polanyi's Paradox and the Shape of Employment Growth*. Working paper no. 20485. The National Bureau of Economic Research, Sept. 2014. Web. 12 Sept. 2016.

Barnes, Alan. "Making Intelligence Analysis More Intelligent: Using Numeric Probabilities." *Intelligence and National Security* 31.3 (2015): 327-44. *Taylor & Francis Online*. Web. 20 Dec. 2016.

Bostrom, Nick. "Ethical Issues in Advanced Artificial Intelligence." *Nick Bostrom*. N.p., 2003. Web. 12 Sept. 2016.

Braverman, Beth. "Here's Why Restaurant Food Prices Are Rising So Fast." *The Fiscal Times*. N.p., 17 Dec. 2015. Web. 20 Dec. 2016.

Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W.W. Norton, 2014. Print.

Chace, Calum. *The Economic Singularity: Artificial Intelligence and the Death of Capitalism*. N.p.: Three Cs, 2016. Print.

Chui, Michael, James Manyika, and Mehdi Miremadi. "Where Machines Could Replace Humans--and Where They Can't (yet)." *McKinsey & Company*. McKinsey Global Institute, July 2016. Web. 20 Dec. 2016.

Daly, Patricia A. "Engineering Enrollments: Has The Long Decline Ended?" *Chemical & Engineering News* 52.16 (1974): 4. *Bureau of Labor Statistics*. Bureau of Labor Statistics, Nov. 1981. Web. 20 Dec. 2016.

Deloitte, and National Restaurant Association. *Restaurant Operations Report*. Rep. National Restaurant Association, 2016. Web. 20 Dec. 2016.

Deloitte. "On an Average (per Person), How Much Do You Spend per Visit at the following Restaurant Formats?" *Statista*. N.p., n.d. Web. 20 Dec. 2016.

"Displaced Workers Summary." *Bureau of Labor Statistics*. Bureau of Labor Statistics, 25 Aug. 2016. Web. 20 Dec. 2016.

Dwoskin, Elizabeth. "Putting a Computer in Your Brain Is No Longer Science Fiction." *The Washington Post*. WP Company, 15 Aug. 2016. Web. 20 Dec. 2016.

Eadicicco, Lisa. "Google CEO: Computers Are Going To Take Our Jobs, And 'There's No Way Around That'" *Business Insider*. Business Insider, 31 Oct. 2014. Web. 20 Dec. 2016.

"Find a Chili's Near Me." *Chili's Grill & Bar Restaurant*. N.p., n.d. Web. 20 Dec. 2016.

Fisk, Donald M. "American Labor in the 20th Century." *Compensation and Working Conditions* (2001): n. pag. *Bureau of Labor Statistics*. Bureau of Labor Statistics, 30 Jan. 2003. Web. 20 Dec. 2016.

Ford, Martin. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic , a Member of the Perseus Group, 2015. Print.

Frey, Carl Benedikt, and Michael A. Osborne. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114 (2013): 254-80. *Oxford Martin School*. Web. 20 Dec. 2016.

"A Game-Changing Result." *The Economist*. The Economist Newspaper, 19 Mar. 2016. Web. 20 Dec. 2016.

Gardner, Howard. *Frames of Mind: The Theory of Multiple Intelligences*. New York: Basic, 1983. Print.

Grace, Katja. "Group Differences in AI Predictions." *AI Impacts*. Future of Life Institute, 24 May 2015. Web. 20 Dec. 2016.

Grace, Katja. "Update on All the AI Predictions." *AI Impacts*. Future of Life Institute, 5 June 2015. Web. 20 Dec. 2016.

Green, Yvette, and John Williams. "An Examination of Operating Costs within a State's Restaurant Industry." *Caesars Hospitality Research Summit* 11th ser. (2010): n. pag. *University Libraries UNLV*. Web. 20 Dec. 2016.

Hanson, Robin. Introduction. *The Age of Em: Work, Love, and Life When Robots Rule the Earth*. Oxford: Oxford UP, 2016. N. pag. Print.

Hendry, Erica R. "What Happens When Artificial Intelligence Turns On Us?" *Smithsonian Magazine*. Smithsonian, 21 Jan. 2014. Web. 12 Sept. 2016.

Henn, Steve. "Robots Are Really Bad At Folding Towels." Audio blog post. *National Public Radio*. Planet Money, 19 May 2015. Web. 20 Dec. 2016.

History.com Staff. "Automated Teller Machines." *History.com*. A&E Television Networks, 2010. Web. 20 Dec. 2016.

Imbery, Lecia. "Low-Wage Workers See Largest Drop in Real Wages." *Coalition on Human Needs*. Coalition on Human Needs, 04 Sept. 2015. Web. 20 Dec. 2016.

"IPad Air 2." *Apple*. N.p., n.d. Web. 20 Dec. 2016.

Kurzweil, Ray. "Don't Fear Artificial Intelligence." *Kurzweil*. KurzweilAINetwork, 30 Dec. 2014. Web. 12 Sept. 2016.

"Locations & Delivery." *Numero Uno Pizza Corporate Website*. N.p., n.d. Web. 20 Dec. 2016.

"Locations." *Applebee's*. N.p., n.d. Web. 20 Dec. 2016.

McKinsey Global Institute, comp. "Automation Potential and Wages for US Jobs." *Tableau Public*. N.p., n.d. Web. 20 Dec. 2016.

Nguyen, Tuong Huy, and Ranjit Atwal. "Market Trends: Slowing Innovation Lengthens Consumers' Device Replacement Cycles." *Technology Research*. Gartner, Inc., 20 Feb. 2013. Web. 20 Dec. 2016.

Pavesic, Dave. "A Sweet Deal: How to Market and Merchandize Desserts to Boost Startup Sales and Profits." *Hospitality Faculty Publications* 1 (2005): 30-35. *Food and Beverage Management Commons*. Web. 20 Dec. 2016.

"Pricing." *CAKE Corporation*. N.p., n.d. Web. 20 Dec. 2016.

Robson, Stephani. "Small Wonder: The Case for Smaller Restaurants and How to Maximize Them." *School of Hotel Administration Collection* (2013): 42-45. *Cornell University School of Hotel Administration*. Web. 20 Dec. 2016.

Schmidt, A., and C. B. Moss. "Mechanized Agriculture: Machine Adoption, Farm Size, and Labor Displacement." *The Journal of Agrobiotechnology Management & Economics* 6th ser. 18.3 (2015): 278-96. *AgroBioForum*. Web. 20 Dec. 2016.

Schwab, Klaus. *The Fourth Industrial Revolution*. Geneva: World Economic Forum, 2016. Print.

Silver, Nate. *The Signal and the Noise: Why so Many Predictions Fail--but Some Don't*. New York: Penguin, 2012. Print.

Susskind, Richard, and Daniel Susskind. *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford: Oxford UP, 2015. Print.

Tetlock, Philip E., and Dan Gardner. Introduction. *Superforecasting: The Art and Science of Prediction*. New York: Crown, 2015. N. pag. Print.

Tetlock, Philip E., and Dan Gardner. "Superforecasting: The Art and Science of Prediction." *Futurecasts Journal*. N.p., Apr. 2016. Web. 20 Dec. 2016.

Thiel, Peter. "Peter Thiel's Keynote." Effective Altruism Summit. Leverage Research, Oakland, CA. 5 July 2013. Web. 20 Dec. 2016.

Thuy Vo, Lam, and Robert Smith. Audio blog post. *National Public Radio*. National Public Radio, 4 June 2013. Web. 20 Dec. 2016.

Toast. *Restaurant Technology in 2016: Dining Edition*. Rep. Toast, Inc., n.d. Web. 20 Dec. 2016.

Tobin, Eric R., and Lynn M. Huffman. "Examining the Impact of Service Times on Overall Guest Satisfaction Perception in the Casual Dining Environment." *Hospitality Review* 5th ser. 24.1 (2006): 42-48. *FIU Digital Commons*. Florida International University, 1 Jan. 2006. Web. 20 Dec. 2016.

"2016 Employer Health Benefits Survey." *The Henry J. Kaiser Family Foundation*. Kaiser Family Foundation, 14 Sept. 2016. Web. 20 Dec. 2016.

United States of America. Department of Labor. Bureau of Labor Statistics. *Long-term Price Trends for Computers, TVs, and Related Items*. TED: The Economics Daily, 13 Oct. 2015. Web. 20 Dec. 2016.

United States of America. Department of Labor. Wage and Hour Division. *Minimum Wages for Tipped Employees*. U.S. Department of Labor, 1 Aug. 2016. Web. 20 Dec. 2016.

U.S. Bureau of Labor Statistics. Labor Force Statistics from the Current Population Survey. 10 Feb. 2016. Raw data. N.p.

Vanek Smith, Stacy, and Robert Smith. "I, Waiter." Audio blog post. *Planet Money*. National Public Radio, 16 Mar. 2015. Web. 20 Dec. 2016.

"Waiters and Waitresses." *Bureau of Labor Statistics*. Bureau of Labor Statistics, May 2015. Web. 20 Dec. 2016.

Wansink, B., and T. Döring. "The Waiters Weight: Does a Servers BMI Relate to How Much Food Diners Order?" *Environment and Behavior* (2015): n. pag. *SAGE Publications*. Web. 20 Dec. 2016.

"Why Is It Getting Harder to Hire in the Restaurant Industry?" *Indeed Blog*. Indeed, 27 Oct. 2015. Web. 20 Dec. 2016.