MELBOURNE INSTITUTE
Applied Economic & Social Research

Working Paper Series

Behind the headline number: Why not to rely on Frey and Osborne's predictions of potential job loss from automation

Michael Coelli

Jeff Borland

Working Paper No. 10/19
October 2019

THE UNIVERSITY OF MELBOURNE

MELBOURNE INSTITUTE

Behind the headline number: Why not to rely on Frey and Osborne's predictions of potential job loss from automation*

Michael Coelli
Department of Economics, The University of Melbourne

Jeff Borland
Department of Economics, The University of Melbourne

Melbourne Institute Working Paper No. 10/19
October 2019

* We are grateful for helpful comments from David Autor, Josh Healy and Mark Wooden. Research in this article has been funded by ARC Discovery Grant DP160102269.

Melbourne Institute:
Applied Economic & Social Research
The University of Melbourne
Victoria 3010 Australia
T +61 3 8344 2100
F +61 3 8344 2111
E melb-inst@unimelb.edu.au
W melbourneinstitute.unimelb.edu.au

Abstract

We review a highly influential study that estimated potential job loss from advances in Artificial Intelligence and robotics: Frey and Osborne (FO) (2013, 2017) concluded that 47 per cent of jobs in the United States were at 'high risk' of automation in the next 10 to 20 years. First, we investigate FO's methodology for estimating job loss. Several major problems and limitations are revealed; especially associated with the subjective designation of occupations as fully automatable. Second, we examine whether FO's predictions can explain occupation-level changes in employment in the United States from 2013 to 2018. Compared to standard approaches which classify jobs based on their intensity in routine tasks, FO's predictions do not 'add value' for forecasting the impact of technology on employment.

JEL classification: J21, O33

Keywords: employment; technology; prediction; job loss; AI and robotics.

1. Introduction

Currently there is an explosion of interest in the new technologies of Artificial Intelligence (AI) and robotics - and nothing is exciting more interest than how the technologies will affect the future of work. Research on this topic is taking a variety of directions.[1] Some studies have adopted a 'backward-looking' approach, examining the effects of automation and computerisation on historical labour market outcomes, such as impacts on the occupational composition of employment and on labour's share of income. Other studies have applied a 'forward-looking' approach, seeking to forecast the future impact of the new technologies on job destruction and the aggregate demand for labour.

The seminal study making predictions of potential job loss due to automation and computerisation is by Frey and Osborne (2013, 2017) (hereafter, FO). They concluded that 47 per cent of jobs in the United States were at 'high risk' of automation in the next 10 to 20 years. The FO study has been and continues to be widely cited (over 4,100 citations on Google Scholar – 29 August 2019) and noted in the popular press and in government.[2] FO's estimates of the probabilities of specific occupations being automated have been applied to construct predictions of job losses in many other countries – including Australia (Durrant-Whyte et al., 2015; Edmonds and Bradley, 2015); Finland and Norway (Pajarinen et al., 2015); Singapore (Lee, 2016); United Kingdom (Deloitte, 2013; Lawrence et al., 2017); Germany (Brzeski and Burk, 2015); Japan (David, 2017); South Africa (le Roux, 2018); and 40 developing countries covered by the World Development Report (World Bank, 2016).

In this study we go 'behind the headline number' of the FO study. Our objective is simple – to establish how much FO's predictions of job loss due to automation and computerisation should be believed. We do this in two ways. First, we provide an in-depth analysis of the method FO use to construct their predictions of potential job loss. Second, we compare FO's predictions against actual changes in occupation-level employment in the United States for the five years since they were constructed (2013 to 2018).

It is important to say at the outset that our analysis of FO's study is undertaken with a full appreciation of the exact task FO set themselves. Their stated objective is to forecast the job

[1] For recent reviews on technology and the future of work, see Autor (2015); Brynjolfsson et al. (2017); and Acemoglu and Restrepo (2019).
[2] An extensive list of press coverage is provided by Carl Frey on his website: http://carlbenediktfrey.com/in-the-media/.

loss by occupation that would occur if the capacity for new technologies to replace labour through automation and computerisation was the only influence on employment. It follows that their forecasts are not intended to predict actual changes in employment, which will also depend on other factors apart from technological change.[3]

Understanding FO's purpose has two important implications for the questions we ask in this paper. First, our analysis of FO's method asks specifically whether it is a useful and appropriate way to predict *potential* job loss due to new technologies. Second, in our comparison of FO's predictions of potential job loss with actual changes in employment in the US, we deliberately do not ask whether their predictions match exactly with actual changes in employment. Rather, we ask only whether their predictions have any explanatory power for actual changes in employment; that is, whether their predictions are significantly related to relative changes in employment by occupation.

We conclude there are major problems with FO's method and predictions. First, the method FO use is problematic and opaque. The method is built on subjective assessments of the potential for individual occupations to be fully automated. Those assessments appear to have been based on limited information about the job content of the occupations. The outcome is a set of predicted probabilities of occupations being automated which are upward-biased and inconsistent with FO's own model of the determinants of technology-induced job loss. Second, FO's predictions of the probability of automation and job loss by occupation do not provide additional information for forecasting actual changes in occupation-level employment that occurred in the United States between 2013 and 2018, once account is taken of the now standard approach that economists use to think about the relation between technology and labour demand: classifying occupations as routine/non-routine and manual/cognitive. This contradicts FO's claim (2017, p.255) that existing methods for understanding the impact of technological change in the labour market are inadequate.

Our analysis of FO's method and predictions demonstrates the importance of looking behind headline-grabbing predictions before relying on them; and ultimately the need to calibrate the weight placed on forecasts used in decision-making to the quality of the underlying

[3] FO (2017, p.268) state: '…we focus on the share of employment that can potentially be substituted by computer capital, from a technological capabilities point of view, over some unspecified number of years. We make no attempt to estimate how many jobs will actually be automated….'; and (p.265) '…we make no attempt to forecast future changes in the occupational composition of employment'. At some points, however, they do become a little stronger in their claims (p.265): '…high probability occupations are likely to be substituted by computer capital relatively soon'.

methodology. It also suggests the value in continuing to apply 'tried and tested' methods for understanding labour market outcomes – in this case, using the concept of routinisation to investigate the impact of technological change.

The outline of the paper is as follows. Section 2 presents an overview of the FO method. Section 3 describes the main limitations of the FO method as we perceive them. Section 4 evaluates the predictions of job loss made by FO against recent actual employment outcomes in the United States. Concluding comments are in section 5.

2. The Frey and Osborne method

FO construct their estimates of potential job loss from automation in the United States using a four-stage procedure. The essence of their method is to subjectively code a subset of US 6-digit occupations as fully automatable or not fully automatable, and to then use that coding of automatability together with nine specific characteristics of occupations to predict the likelihood of full automation for all 6-digit occupations in the United States.

Stage 1

70 6-digit US Department of Labor Standard Occupation Classification (SOC) occupations were chosen by FO to be subjectively hand-labelled as to whether they could be fully automated.[4] The 70 occupations were taken from the 702 6-digit occupations for which FO were able to obtain O*NET information and measures (in 2010, using version 15 of O*NET) (FO, 2017, p.263).[5] FO (2017, p.264) state that the method for choosing the 70 occupations was to select '… those occupations whose computerisation label we are highly confident

---

[4] Appendix Table A1 provides a list of the 70 6-digit occupations, organised according to the Autor et al (2003) classification of occupations as routine/non-routine and cognitive/manual. Throughout this study we classify occupations into four categories of routine/non-routine by cognitive/manual in two steps. First, we classify each occupation into one of the ten occupation categories in Acemoglu and Autor (2011). Second, we classify those ten occupation categories as follows: non-routine cognitive (managers; professionals; technicians); routine manual (production; operators); routine cognitive (office/administration; sales); and non-routine manual (protective; food/cleaning; personal services).

[5] The Occupational Information Network or O*NET is developed under the sponsorship of the United States Department of Labor/Employment and Training Administration (USDOL/ETA) and is a primary source of occupational information. The O*NET database contain many standardized and occupation-specific descriptors on almost 1,000 occupations covering the entire US economy, It is continually updated from input by a broad range of workers in each occupation. More details on O*NET can be found at: https://www.onetcenter.org/overview.html.

about.' The hand-labelling of the automatability of each occupation was done by answering the question:

"Can the tasks of this job be sufficiently specified, conditional on the availability of big data, to be performed by state-of-the-art computer-controlled equipment?" (FO, 2017, p.263).

FO undertook the hand-labelling in collaboration with a group of machine learning (ML) experts gathered at a workshop on automation of tasks done by labour held at Oxford University's Engineering Sciences Department in early 2013. Each of the 70 chosen occupations was assigned a 'one' if all tasks undertaken as part of the job were considered automatable (after allowing for 'task simplification' in some cases) and 'zero' otherwise (FO, 2017, p.263). No information is provided on how the opinions of the ML experts were aggregated to assign each occupation as fully automatable or not fully automatable. Of the 70 occupations hand-labelled, 37 (53 per cent) were assigned as being fully automatable. The hand-labelling by FO is described as having been based on '…eyeballing the O*NET tasks and job description of each occupation' (FO, 2017, p.263). These text descriptions of tasks and job were specific to each occupation. Significantly, the classification of job characteristics for individual occupations available in the O*NET database was not used in this stage.

Stage 2

FO identify what they regard as the three main 'bottlenecks' to automation of tasks performed by labour. One proposed bottleneck, 'perception and manipulation', is argued to derive from robots not being as able as humans in the manual manipulation of odd-shaped objects. The second and third bottlenecks are respectively 'creative intelligence' and 'social intelligence'. These bottlenecks are suggested to exist due to AI being still unable to deal with many creative and social tasks.

FO choose nine specific O*NET job characteristics to capture these three bottlenecks. The job characteristics selected by FO are taken from the complete O*NET database where characteristics are grouped into categories as follows (with the number of characteristics in each category provided in parentheses): abilities (52), interests (6), knowledge (33), skills (35), work activities (41), work context (57), work styles (16) and work values (6). The three proposed bottlenecks and FO's nine specific O*NET job characteristics (and the O*NET categories they are in) are listed in Table 1.

Table 1**: Frey and Osborne (2017) bottlenecks to computerisation, O*NET job characteristics**

Bottleneck	**O*NET job characteristic**	**O*NET category**
Perception and manipulation	Finger dexterity	Abilities
	Manual dexterity	Abilities
	Cramped work space, awkward positions	Work context
Creative intelligence	Originality	Abilities
	Fine arts	Knowledge
Social intelligence	Social perceptiveness	Skills
	Negotiation	Skills
	Persuasion	Skills
	Assisting and caring for others	Work activities

Source: Table 1 of Frey and Osborne (2017).

It is worth noting that the job characteristics chosen by FO to represent the proposed main ‘bottlenecks’ on automation are drawn from 246 job characteristics available in the O*NET.[6] No detailed justification is provided for the nine job characteristics chosen – at least not in comparison with other job characteristics that might have been selected from O*NET to represent the bottlenecks. As an example, alternative abilities and skills that might have been applied to represent the bottleneck of ‘creative intelligence’ are visualisation; writing; complex problem-solving; design; or flexibility of closure.[7]

FO estimate models to establish the association between their hand-labelling of whether an occupation can be fully automated (from stage 1) and the nine O*NET job characteristics for the 70 6-digit occupations.[8] Three alternative models are estimated, all of which are essentially non-linear regression (Logit) models of the indicator of whether an occupation can be fully automated on the nine O*NET variables. The base model is the standard Logit,

[6] The O*NET also has information on the specific tasks undertaken within each occupation, as well as on the education, training, experience and licensing requirements of each occupation.
[7] Autor (2013) has noted that ‘…researchers who wish to use these databases [O*NET or DOT] as sources for task measures are essentially *required* to pick and choose among the plethora of scales available…While I have found that task measures distilled from DOT and O*NET can serve as powerful proxies for occupational tasks, I am at best only moderately comfortable with these tools because their complexity and opacity places little discipline on how they are applied and interpreted.’
[8] For some job characteristics in O*NET there are multiple measures available. For eight of the nine job characteristics used by FO (the exception being ‘cramped work space, awkward positions’) O*NET provides both a level measure and an importance measure. For those variables FO use the level measure. The ‘cramped work space, awkward positions’ characteristic is in the work context category and has a single measure of frequency which is used by FO.

where a linear function of the nine variables is used as the index function underlying the model. FO then allow in a non-parametric way for non-linear and potentially interacted relationships between the nine O*NET variables. They use the Gaussian Process Machine Learning (GPML) toolbox of Rasmussen and Nickisch (2010) during estimation and try two different ways of characterising the covariance matrix when defining the Gaussian process: the exponentiated quadratic and rational quadratic. The exponentiated quadratic is found to perform the best of the three models using the criteria of log likelihood and area under the receiving operator curve (AUC), narrowly outperforming the rational quadratic. Both non-parametric methods considerably outperform the base Logit.

The non-linear (Logit) estimation method results in the predicted probabilities of automation for occupations clustering near zero and one. This is evident from Figure 1, which shows the unweighted distribution of predicted probabilities for the 70 hand-labelled occupations.

Figure 1**: Frey and Osborne (2017) predicted probabilities for 70 hand-labelled occupations**

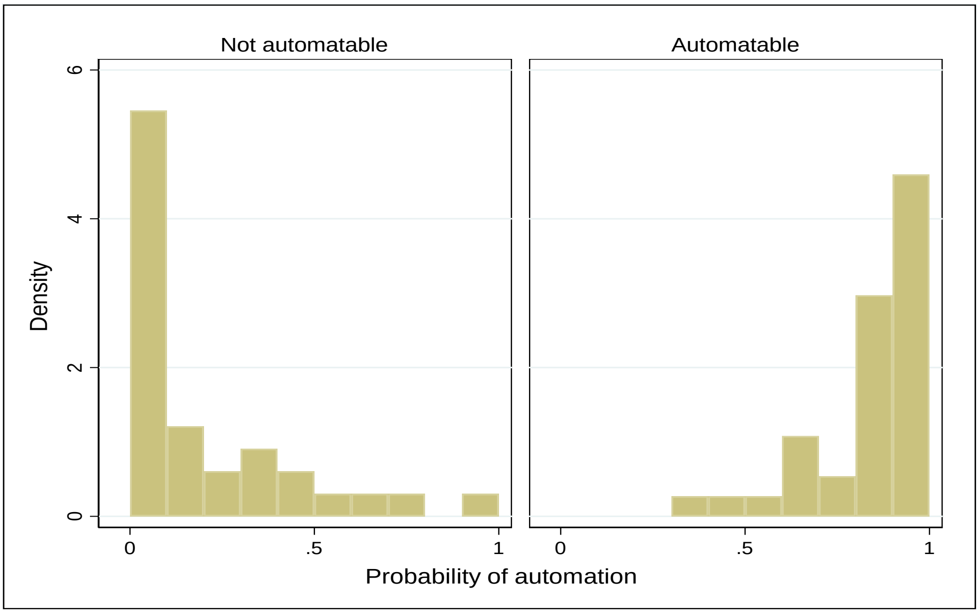

Source: Appendix A Table of Frey and Osborne (2017).

Stage 3

Automation risk probabilities for all 702 6-digit SOC occupations are constructed using estimates from the non-linear regression model from stage 2 together with the nine O*NET job characteristics listed in Table 1. The unweighted distribution of the predicted probabilities for all 702 occupations is shown in Figure 2. Not surprisingly, these predicted probabilities are also clustered around zero and one, with 57 per cent of occupations having probabilities greater than one-half of being fully automated.

Figure 2**: Frey and Osborne (2017) predicted probabilities for all 702 occupations**

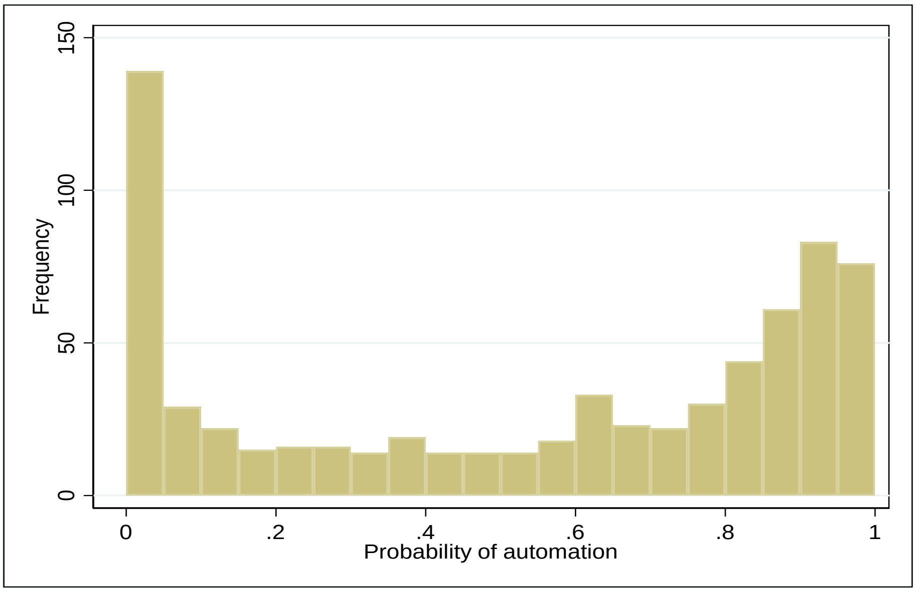

Source: Appendix A Table of Frey and Osborne (2017).

Stage 4

Any occupation with a predicted probability of being fully automated of 0.7 or above is classified to be at 'high risk' of automation, and it is assumed that all workers in that occupation could be replaced by automation within the next decade or two. Using this definition of 'high risk' and data on US employment in 6-digit occupations for 2010, FO estimate that 47 per cent of US workers are at 'high risk' of their jobs being entirely automated.

3. Investigating the Frey and Osborne method

In this section we evaluate FO's methodology. We work sequentially through the stages in the method.

3.1 Stage 1: Hand-labelling

Precisely how the group of ML experts gathered at the Oxford workshop determined whether all tasks in each of the 70 occupations were automatable is not described in much detail. FO provide the question that the group was asked to answer, but little other background. From their description of the process of hand-labelling it appears that the information provided to the group about each occupation was limited. That seems the likely explanation for several problematic features of the assignment of occupations as fully automatable.

First, the classification of some of the 70 occupations suggests that the ML experts were not well-versed in the set of tasks undertaken in those occupations. As one example, FO labelled the occupation ‘Accountants and Auditors’ (code 13-2011) as fully automatable. The O*NET job description for this occupation is:

> “Examine, analyze, and interpret accounting records for the purpose of giving advice or preparing statements. Install or advise on systems of recording costs or other financial and budgetary data.”

The list of tasks undertaken by ‘Accountants and Auditors’ is provided in Appendix Table A2. Given this job description and list of tasks, it is difficult to see how the group of ML experts managed to agree and label this occupation as automatable. Many of the specified tasks require interpretation of information about organisational performance, with interpretation usually being regarded as outside the scope of what will be substituted for by AI (Agrawal et al., 2017, 2019). Other tasks require interviewing of staff at organisations whose performance is being reviewed. There are also tasks which suggest that a role of the occupation is to design and implement systems that can automate analysing and reporting on the performance of organisations.

Second, the hand-labelling of occupations by the ML experts was conducted “conditional on the availability of big data”. Unfortunately, by being instructed to assume away the need to gather information, the ML experts were being told to ignore what is in fact an important task undertaken in many occupations. As an illustration, of the thirteen non-routine cognitive occupations labelled as automatable, twelve of them have average or above measures for the importance and level of the “getting information” work activity in the O*NET. The question is, where is the “big data” coming from, if not from the collection efforts of such high-skill workers?

As well as being limited in the information available to inform their predictions, the ML experts may also have exhibited a natural bias among technology experts towards thinking automation will have substantial effects. Evidence for a predisposition for experts to believe their area of practise has an impact has been found in other contexts – such as judgments by surgeons about the impact of their medical treatments (see for example the research on this topic by Archie Cochrane summarised in Goldacre, 2008, pp.43-44). As well, some behavioural scientists have suggested that training in deterministic sciences (such as chemistry and physics) where causal connections are 'tidier' can limit the capacity to forecast compared to a training in probabilistic sciences (such as psychology and economics) which deal with phenomena that are not perfectly predictable and where causal relations are rarely necessary and sufficient (Gilovich, 1991, pp.190-91).

3.2 Stages 2 and 3: Predicting automation probabilities for the 702 occupations

In stages 2 and 3, FO construct predicted probabilities of the automatability of all 702 US occupations. This is done by extrapolating from the estimated relation between the hand-labelled predictions of whether occupations are fully automatable and the nine O*NET job characteristics for the 70 chosen occupations.

FO's method of prediction is motivated by their argument that the nine O*NET job characteristics capture bottlenecks that are a key influence on whether occupations will be automatable. But there is a major problem with this argument. When the relation between FO's predictions of the automatability of occupations and the O*NET variables they choose to represent bottlenecks is estimated, it suggests that several of these job characteristics do not constitute bottlenecks on automation.[9]

Table 2 shows the results of simple regressions of the relation between FO's predictions of the automatability of occupations and each of the nine O*NET variables. The regression estimates in the left-hand column use the sample of 70 hand-labelled occupations with an indicator for whether the ML experts assigned an occupation as fully automatable as the explanatory variable. The regression estimates in the right-hand column use the sample of 698 6-digit US occupations with FO's predicted automation probabilities as the explanatory variable.[10] Each element in the table shows the result from a regression with one of the nine O*NET variables as the dependent variable.

---

[9] FO (2017, p.266) note that their model predicts that occupations with high levels of the three 'perception and manipulation' measures of finger dexterity, manual dexterity and cramped work space are more likely to be in the medium and high ranges of automation probability risk. If these occupation abilities/work contexts are truly bottlenecks to computerisation, such predictions are difficult to understand. Having acknowledged this issue, FO choose not to pursue it further.

[10] FO provide automation probabilities for 702 occupations. We find, however, that four of these occupations ("all other" aggregated occupations) do not have O*NET data available in the version 15 release, and hence we are unable to assign them values for the nine O*NET variables.

Table 2: **Relation between Frey and Osborne (2017) predictions of automation and the nine O*NET variables**

O*NET measure	Labelled	Predicted
Perception and manipulation		
Finger dexterity	-0.067	0.270
	(0.163)	(0.071)
Manual dexterity	0.255	1.074
	(0.295)	(0.117)
Cramped work space, awkward positions	-0.019	0.421
	(0.162)	(0.070)
Creative intelligence		
Originality	-0.997	-1.599
	(0.162)	(0.045)
Fine arts	-0.494	-0.833
	(0.155)	(0.111)
Social intelligence		
Social perceptiveness	-0.909	-1.238
	(0.150)	(0.050)
Negotiation	-0.682	-1.079
	(0.160)	(0.055)
Persuasion	-0.815	-1.210
	(0.157)	(0.051)
Assisting and caring for others	-1.101	-1.327
	(0.296)	(0.116)
Occupations	70	698

Notes: Each estimate is taken from a separate regression of each O*NET measure on an indicator for an occupation being labelled as subject to automation (column 1), and on the prediction of automation risk of an occupation (column 2). Heteroskedasticity-robust standard errors are provided in parentheses.

We begin by looking at the three variables relating to 'perception and manipulation' tasks. In the left-hand column these variables are unrelated to occupations being hand-labelled as fully automatable (large p-values). In the right-hand column the variables are positively related to the predicted automation probabilities (p-values below 0.001).[11] Hence, FO's procedure is

[11] This is likely due to the negative correlations between the three 'perception and manipulation' variables and five of the other six O*NET variables considered. Those correlations are embedded in the predicted probabilities of occupations being automated through the model that is used to make those predictions which was estimated with the set of nine O*NET variables.

predicting higher automation probabilities among occupations with characteristics they have argued are bottlenecks to automation. We now turn to the six O*NET variables intended to capture the bottlenecks of 'creative intelligence' and 'social intelligence'. Consistent with the argument that these are bottlenecks to automation, they are all significantly lower in occupations that are hand-labelled as being fully automatable.

A possible explanation for the contrary findings on 'perception and manipulation' as bottlenecks is the method used by FO to construct their predictions, a non-parametric estimator for the latent index underlying the Logit model (a Gaussian Processor). To investigate this issue, we construct predicted probabilities of automation using the standard Logit model, where the nine O*NET variables enter the estimation linearly in the underlying latent index function. The coefficients and marginal effects from the standard Logit estimation using the 70 hand-labelled occupations are reported in Table 3.[12] In this simple model, several of the nine O*NET variables intended by FO to capture bottlenecks to automation are not negatively related to occupations being hand-labelled as fully automatable. Finger dexterity and negotiation are positively related to labelling of an occupation as fully automatable, conditional on the other included measures. As well, several other O*NET variables are not significantly related to the predictions of automatability for the 70 chosen occupations. It must be kept in mind, however, that the O*NET variables within the three categories of bottlenecks are correlated, which is likely to affect these estimates.

Analysis of the relation between predictions of the probability of automation and the nine O*NET variables can also be done using our simpler Logit model to make the predictions for the full set of 698 6-digit occupation. The findings are reported in Table 4. The left-hand column shows results from simple regressions of each of the nine O*NET variables against the automation predictions from our linear Logit model of Table 3. The right-hand column shows results from the same set of regressions where predictions are from an extended non-linear Logit model.[13] Positive or near zero relationships are again found to exist between the

[12] The marginal effects from a Probit estimation of the model were similar.

[13] This extended Logit model for generating predicted probabilities of automation includes all nine O*NET variables linearly, the squares of all nine variables plus the interactions of the leading principal components (one for each of the three bottlenecks). Models with more interactions were not able to be estimated with the limited number of observations (70).

'perception and manipulation' O*NET variables and the predicted automatability of occupations.[14]

Table 3: **Relation between FO's hand-labelling of automation risk and the nine O*NET variables using a Logit model**

O*NET measure	Logit coefficients	Marginal effects	Mean (standard devn.)
Finger dexterity	2.356	0.189	2.470
	(1.178)	(0.0937)	(0.670)
Manual dexterity	-1.780	-0.142	1.787
	(0.589)	(0.0466)	(1.221)
Cramped work space, awkward positions	-2.183	-0.175	1.784
	(0.815)	(0.0632)	(0.667)
Originality	-2.047	-0.164	2.837
	(1.348)	(0.108)	(0.829)
Fine arts	-0.987	-0.0790	0.418
	(1.753)	(0.139)	(0.661)
Social perceptiveness	-7.922	-0.634	3.144
	(2.735)	(0.181)	(0.754)
Negotiation	4.643	0.372	2.707
	(2.154)	(0.165)	(0.737)
Persuasion	-2.487	-0.199	2.850
	(1.892)	(0.150)	(0.759)
Assisting and caring for others	-0.137	-0.0110	2.942
	(0.407)	(0.0322)	(1.307)
Constant	26.49		
	(6.536)		
Observations	70	70	70

Notes: Heteroskedasticity-robust standard errors provided in parentheses.

Using alternative methods to construct predictions of automation (the linear Logit and extended Logit models) therefore does not resolve the puzzle of why the relation between the predictions and the O*NET variables representing 'perception and manipulation' tasks is inconsistent with them constituting a bottleneck to automation. Hence, it seems likely that the finding on the perception and manipulation variables is more to do with FO's assignment of occupations as fully automatable in the stage of hand-labelling, rather than being caused

[14] Although the relationships are not nearly as strong as those with predictions derived using the FO estimation method. Without more information regarding how FO's predictions were constructed, it is not possible to reveal why these differences are arising.

by their empirical method for constructing predictions of automation for all occupations.[15] This is an issue we take up again later in the paper.

Table 4: **Relation between predictions of automation from our Logit models and the nine O*NET variables**

O*NET measure	Linear	Non-linear
Perception and manipulation		
Finger dexterity	0.097	0.154
	(0.063)	(0.055)
Manual dexterity	0.408	0.363
	(0.101)	(0.091)
Cramped work space, awkward positions	-0.101	0.018
	(0.071)	(0.066)
Creative intelligence		
Originality	-1.260	-0.936
	(0.054)	(0.053)
Fine arts	-0.754	-0.095
	(0.090)	(0.079)
Social intelligence		
Social perceptiveness	-1.109	-0.819
	(0.043)	(0.043)
Negotiation	-0.798	-0.743
	(0.054)	(0.048)
Persuasion	-0.970	-0.857
	(0.050)	(0.046)
Assisting and caring for others	-1.136	-0.886
	(0.101)	(0.094)
Occupations	698	698

Notes: Each estimate is taken from a simple regression of each O*NET measure on the predicted probabilities of automation risk from the two Logit models estimated. Heteroskedasticity-robust standard errors provided in parentheses.

[15] Another fundamental problem with FO's method may exist, however. Taking seriously the idea of a bottleneck suggests that if even one of the aspects of tasks captured by the O*NET variables constitutes a barrier to automation then that would be sufficient to prevent automation. It is unclear whether FO's method is flexible enough to allow a specification where automation risk depends on any of the three main bottlenecks identified by FO being present or reaching a critical threshold level in the set of tasks in an occupation.

A further point on FO's predictions of the automatability of occupations relates to how those predictions are interpreted. FO (2017, p.266) suggest that: '…the probability axis [in Figure 3 which shows the distribution of the probabilities of individual occupations being automated] can be seen as a rough timeline, where high probability occupations are likely to be substituted by computer capital relatively soon'. Following this logic, FO (2017, p.266-67) then state that the U-shaped distribution of probabilities implies '…two waves of computerisation separated by a 'technological plateau''. In other words, the clustering of occupations at high probabilities can be interpreted as showing that some occupations can be automated relatively quickly; and the absence of occupations at mid-range probabilities and clustering at low probabilities shows that following an initial wave of automation there will be a '…subsequent slowdown in computers for labour substitution, due to persisting inhibiting engineering bottlenecks to computerisation'. In fact, the U-shaped distribution of probabilities of automation generated by FO is much more likely to be a simple artefact of them having used a Logit model to derive those predictions together with the hand-labelling having assigned an even split of occupations between being and not being capable of full automation.

3.3 Stage 4: The threshold

To move from predictions of the automatability of individual occupations to their forecast of potential aggregate job loss in the US economy, FO make two main assumptions: first, that occupations with a predicted probability of being fully automated above 0.7 will be automated; and second, that all jobs will be lost in an occupation which is fully automated.

The findings from sensitivity analysis which shows how the cut-off predicted probability for 'high risk' of automation affects the predictions of potential job loss are reported in Table 5. Obviously, the predicted proportion of jobs automated varies inversely with the cut-off level for 'high risk' of automation. For the FO method, the extent of variation with the cut-off is quite large. For example, a cut-off of 0.5 yields 64.2 per cent of jobs in the 'high risk' category, whereas with a cut-off of 0.9 the proportion of jobs is only 31.2 per cent. For our Logit models, the extent of variation is less for cut-offs below 0.7, but sizable above that level. Since the choice of cut-off for the 'high risk' category is necessarily arbitrary, finding that the predicted proportion of jobs which could be subject to automation varies so much with the cut-off becomes an important dimension of judging the validity of the predictions.

Table 5: **Sensitivity of job loss predictions – method and threshold**

Threshold	Frey-Osborne	Logit	Non-linear Logit
0.50	0.642	0.575	0.568
0.55	0.635	0.563	0.563
0.60	0.597	0.557	0.563
0.65	0.529	0.548	0.562
0.70	0.482	0.539	0.550
0.75	0.459	0.501	0.542
0.80	0.426	0.493	0.532
0.85	0.391	0.447	0.529
0.90	0.312	0.376	0.512
0.95	0.167	0.305	0.465

Notes: Figures are the proportion of workers in occupations with a predicted probability of automation above the relevant threshold, using 2010 US Bureau of Labor Statistics employment by occupation figures (employees only, no self-employed).

It is also possible to question FO's assumption that automation of an occupation will cause all jobs in that occupation to cease to exist. In most cases where new technologies do substitute for labour, this seems an unlikely scenario.[16] Arntz et al. (2016, 2017) propose that individual workers in an occupation may not undertake the same mix of tasks with precisely the same relative emphasis, and hence the likelihood of their jobs being automated will differ. To test this idea, they estimate the relationship between task measures reported by individual workers in the US version of the Program for the International Assessment of Adult Competencies (PIAAC) and FO's estimates of the probabilities of occupations being automated. That estimated relation is then applied to derive predicted probabilities for job loss due to automation for individual workers. Using this method Arntz et al. (2017) conclude that only nine percent of jobs in the United States are at risk due to computerisation and automation.[17]

[16] Technological change also affects the types of tasks workers undertake within their jobs. Learning new skills to use these new technologies has also been a feature of employment since the advent of the personal computer in the early 1980s. Changes in work tasks are likely to continue in the future as new technologies are adopted in workplaces.

[17] However, a major caveat is that the difference in predicted job loss between FO and Arntz et al. (2017) does not seem to be due to heterogeneity in the automatability of tasks undertaken by workers within the same occupation. Instead, the difference appears to mainly reflect the inability of analysis using the PIAAC data – which is available only for 39 aggregated occupation groups and a specified set of tasks – to adequately reflect differences in the probabilities of automation between occupations

4. How well have FO's predictions fared?

Evaluating the soundness of the method used to generate predictions provides an indirect way to assess their likely accuracy. Ultimately, however, it could be argued that the value attached to FO's predictions should depend on their realised accuracy. Since sufficient time has elapsed since FO first made their predictions, an initial assessment of their accuracy can now be made.

We do this by comparing FO's predictions of the automatability of occupations with actual changes in occupation-level employment in the United States. In making this comparison, our main test will be whether FO's predictions have explanatory power in a regression model for occupation-level employment changes. If technology is a key driver of labour market outcomes, and FO have accurately represented how automation will vary between occupations, then it should follow that their predictions have some explanatory power for occupation-level changes in employment. Note that this test is not asking that FO's forecasts should be the only, or even a strong predictor, of employment outcomes – simply that they have explanatory power for the evolution of the occupational composition of employment in the United States.

Table 6 reports the results of simple regressions of the relation between the percentage change in employment at the 6-digit occupation level using data from the Bureau of Labor Statistics (BLS) and FO's predictions of the likelihood of an occupation being fully automated. In all regressions the dependent variable is the percentage change in occupation-level employment from 2013 to 2018. Estimation is done using 2013 employment by occupation levels as weights.[18] This period from 2013 to 2018 captures the interval from publication of the working paper version of FO's predictions through to the most recent data on employment by occupation available for the US from the BLS. We also estimate models using employment growth over the longer period of 2010 to 2018 as the dependent variable, with results presented in Appendix Tables A3 and A4. Our main findings remain[19]

within those 39 groups and in tasks undertaken in routine jobs compared to non-routine manual jobs (see Coelli and Borland, 2019).

[18] The estimates cover 701 of the 702 occupations FO provide automation risk probabilities for. We are unable to include the occupation Hunters and Trappers (SOC code 45-3021) as the BLS does not report employment in that occupation.

[19] Our motivation for also analysing the longer period is that FO's predictions of automation were based on O*NET information from 2010 and on the mix of jobs that existed at that time and were seeking to forecast what would happen in '…perhaps a decade or two'(FO, 2017, p.265).

Table 6: **Employment changes by occupation, 2013 to 2018 (per cent)**

Regressor	(1)	(2)	(3)	(4)	(5)	(6)
FO high risk of automation	-3.49 (2.41)					
FO probability of automation		-6.59 (2.40)	-2.69 (2.85)		-1.59 (4.19)	
Routine cognitive			-9.03 (3.18)	-10.70 (2.38)		
Routine manual			0.394 (2.44)	-1.011 (2.06)		
Non-routine manual			1.25 (3.72)	0.09 (3.94)		
Routine task intensity (RTI) index					-0.108 (0.071)	-0.118 (0.052)
Constant	10.89 (1.05)	13.01 (1.35)	12.63 (1.46)	12.07 (1.31)	16.21 (3.08)	15.82 (3.72)
Occupations	701	701	701	701	666	666
R-squared	0.012	0.024	0.084	0.082	0.054	0.053

Notes: BLS employment by occupation in 2013 used as weights. Heteroscedasticity-robust standard errors provided in parentheses.

Column (1) shows results where the only explanatory variable is an indicator for whether an occupation is predicted by FO to be at 'high risk' of automation (that is, an automation risk probability of 0.7 or above). Occupation-level employment growth in the United States is only weakly related (p-value = 0.148) to whether an occupation was predicted by FO to be at 'high risk' of automation.[20] Column (2) replaces the indicator for an occupation being identified as having a 'high risk' of automation with FO's full set of predicted probabilities of each occupation being automated. For this representation of FO's predictions, there is a more significant negative relation (p-value = 0.006) with changes in occupation-level employment.

Column (3) show results when indicators for broad categories of occupations are added as explanatory variables. The broad categories consist of the now standard ALM classification of jobs according to whether they are routine/non-routine and cognitive/manual (Autor et al.,

[20] Note as well that, combining the point estimate for occupations identified as being at high risk of automation with the estimated change in employment in occupations not at risk (measured as the coefficient on the constant term), the overall change in employment in occupations deemed high risk of automation was growth of around 7.4 per cent over the five-year period. This is a very different outcome to the prediction of disappearing jobs, at least at this stage.

2003).[21] Once these broad occupation categories are controlled for, the point estimate on FO's predicted probability of automation is reduced considerably in size and significance (p-value = 0.345). By contrast, the ALM categories are jointly highly significant (p-value = 0.003). In column (4) the relation between the percentage change in employment and the ALM categories remains largely unchanged and increases in significance when FO's predicted probability of automation is omitted (p-value = 0.0001). These estimates reveal relatively weak employment growth of about 1 per cent in routine cognitive occupations and stronger growth of about 11 to 12 per cent in the other three occupation groups.

Column (5) shows results when the Routine Task Intensity (RTI) index (Autor and Dorn, 2013) is included as an explanatory variable together with FO's predicted probability of automation instead of the ALM categories. The relation between FO's predictions and employment growth by occupation is small and highly insignificant (p-value = 0.704). The relation with the RTI is negative, implying that employment growth has been weakest in occupations that include tasks which are most subject to routinisation. The association with the RTI has a much higher level of significance than FO's predictions (p-value = 0.125).[22] In addition, the RTI becomes more significant when included as the only explanatory variable in column (6) (p-value = 0.023).

It can be concluded from Table 6 that FO's predictions do not improve the capacity to forecast occupation-level changes in employment compared to using the concept of routinisation and distinguishing between manual and cognitive jobs.[23] This is at odds with the claim by FO (2017, p.255) that: 'Although there are indeed existing useful frameworks for examining the impact of computers on the occupational employment composition, they seem inadequate in explaining the impact of technological trends going beyond the computerisation of routine tasks.'

---

[21] The omitted base category is non-routine cognitive occupations.

[22] We are only able to estimate this model including the RTI index for 666 occupations as it was not possible to construct the RTI index for those occupations that have emerged since 1970, the timing of the collection of the data underlying the construction of the RTI index. This index is based on measures for occupations defined under the Dictionary of Occupation Titles (DOT), the precursor to the O*NET. The occupations without RTI data comprise 4% of employment in 2013 and are clustered in the IT and medical areas.

[23] One occupation that stands out among the set of non-routine manual occupations is personal care aides. Predicted by FO to be at high risk of automation, employment nearly doubled in this occupation, from 1.14 million in 2013 to over 2 million workers in 2018.

Part of the explanation for why FO's predictions do not have value added for forecasting employment growth is that their predictions are correlated with the ALM occupation categories. Table 7 shows the proportions of employment within each of the five occupation categories that are classified at 'low/medium' and 'high' risk of automation by FO. It confirms a high level of correlation. For example, only 12 per cent of non-routine cognitive employment is deemed at high risk of automation, while 81 per cent of routine cognitive employment is at high risk.

Table 7: **US employment by occupation group and automation risk, 2013**

	Risk of automation	
Occupation group	Low / Medium	High
Non-routine cognitive	88.4	11.6
Routine cognitive	18.8	81.2
Routine manual	39.3	60.7
Non-routine manual	51.2	48.8
Total	51.7	48.3

Notes: Per cent of 2013 US employment using BLS data.

The final issue we consider is whether the explanatory power of FO's predictions is limited by the inconsistency between those predictions and FO's conceptual framework for the determinants of automation. As has been noted, FO identify occupations with high scores for the three O*NET job characteristics relating to 'perception and manipulation' as being likely to have a low probability of automation, yet their predictions of the likelihood of automation are positively correlated with those variables. What then if 'perception and manipulation' really does constitute a bottleneck? If that is the case, it may be the absence of that bottleneck which explains why FO's predicted probabilities of automation are not useful in predicting employment outcomes. In other words, hand-labelling of the 70 occupations has caused FO to predict that employment growth will be negatively related to 'perception and manipulation' skills, whereas actual employment growth has been positively related to those skills. To investigate this further we estimate models for the change in occupation-level employment from 2013 to 2018 including indicators for occupations being in the low and high terciles of the 'perception and manipulation' bottleneck proposed by FO.[24]

Results are reported in Table 8. In the model in column (1) the only explanatory variable is FO's predictions. The model in column (2) also includes the tercile indicators for the 'perception and manipulation' bottleneck. The model in column (3) combines the explanatory variables from the preceding models. In the model in column (4) the RTI is added as an explanatory variable.

Table 8: **Employment changes by occupation and bottlenecks, 2013 to 2018 (per cent)**

Regressor	(1)	(2)	(3)	(4)
FO probability of automation	-6.60		-7.72	-2.56
	(2.41)		(2.07)	(3.72)
Low perception and manipulation		-6.30	-7.26	-6.49
		(3.20)	(3.11)	(3.27)
High perception and manipulation		-2.71	-3.00	-1.91
		(3.04)	(3.08)	(2.86)
Routine task intensity (RTI) index				-0.114
				(0.069)
Constant	13.01	12.29	17.18	20.04
	(1.35)	(2.75)	(2.78)	(4.55)
Occupations	697	697	697	662
R-squared	0.024	0.027	0.058	0.085

Notes: BLS employment by occupation in 2013 are used as weights. Heteroscedasticity-robust standard errors provided in parentheses.

Two main findings emerge from the results in Table 8. First, there is evidence that job characteristics relating to 'perception and manipulation' have been associated with changes in employment in the direction expected by FO. The percentage change in employment is lower for occupations in the bottom tercile of jobs based on the need for perception and manipulation than the middle tercile (p-value = 0.049). While jobs in the top tercile for perception and manipulation had lower rates of employment growth than the middle tercile,

[24] These terciles were constructed as follows. To begin, the leading principal component from the three O*NET job characteristics representing 'perception and manipulation' bottleneck was extracted. Terciles were then constructed using this leading principal component with BLS occupation employment in 2013 used as weights.

this association is only weakly significant (p-value = 0.374). Hence, while there is some support for this explanation for why FO's predictions cannot explain changes in employment, it is not conclusive. Second, the explanatory power of the RTI index remains similar when included with the O*NET controls (compared for example to the results in Table 6). This is a further demonstration of the value of that method of classifying occupations for understanding the evolution of employment by occupation.

5. Concluding remarks

In this paper we have argued that FO's predictions of job loss in the US are not built on a solid foundation. Rather, the foundation of their predictions, the hand-labelling of whether occupations are fully automatable, appears to lack rigour - specifically regarding its potential replicability, internal consistency and subjectivity of key elements. The hand-labelling results in predictions of the relative likelihood of occupations being automated which are inconsistent with FO's own beliefs about the determinants of automation. Furthermore, the hand-labelling gives rise to predictions of occupations being automated which are not informative for forecasting changes in occupation-level employment in the US. The analysis reported in this paper suggests that a much better foundation for understanding the consequences of new technologies on employment outcomes is by using the standard ALM approach – that is, distinguishing between jobs using the categories of routine/non-routine and cognitive/manual.

References

Acemoglu, Daron and David Autor (2011), ‘Skills, tasks and technologies: Implications for employment and earnings’, *Handbook of Labor Economics*, 4: 1043-1171.

Acemoglu, Daron and Pascual Restrepo (2019), ‘Automation and new tasks: How technology displaces and reinstates labor’, *Journal of Economic Perspectives*, 33(2), 3-30.

Agrawal, Ajay, Joshua Gans and Avi Goldfarb (2017), ‘What to expect from artificial intelligence’, *MIT Sloan Management Review*, 56(3): 23-26.

Agrawal, Ajay, Joshua Gans and Avi Goldfarb (2019), ‘Artificial intelligence: The ambiguous labor market impact of automating prediction’, *Journal of Economic Perspectives*, 33(2): 31-50.

Arntz, Melanie, Terry Gregory and Ulrich Zierahn (2016), ‘The risk of automation for jobs in OECD countries’, OECD Social Employment and Migration Working Papers, no. 189.

Arntz, Melanie, Terry Gregory and Ulrich Zierahn (2017), ‘Revisiting the risk of automation’, *Economics Letters*, 159, 157-160.

Autor, David (2013), ‘The “task approach” to labor markets: an overview’, *Journal of Labour Market Research*, 46(3): 185-99.

Autor, David (2015), ‘Why are there still so many jobs? The history and future of workplace automation’, *Journal of Economic Perspectives*, 29(3): 3-30.

Autor, David and David Dorn (2013), ‘The growth of low-skill service jobs and the polarization of the US labor market’, *American Economic Review*, 103(5): 1553-97.

Autor, David, Frank Levy and Richard Murnane (2003), ‘The skill content of recent technological change: An empirical examination’, *Quarterly Journal of Economics*, 118: 1279-1333.

Brynjolfsson, Erik, Daniel Rock and Chad Syverson (2017), ‘Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics’, National Bureau of Economic Research, Working paper no. 24001.

Brzeski, Carsten and Inga Burk (2015), ‘Die roboter kommen: folgen der automatisierung für den deutschen arbeitsmarkt’, ING DiBa Economic Research, accessed at: https://www.ing.de/binaries/content/assets/pdf/ueber-uns/presse/publikationen/ing-diba-economic-analysis-die-roboter-kommen.pdf on 3 September 2019.

Coelli, Michael and Jeff Borland (2019), ‘Estimating the automation risk of Australian jobs’, mimeo, University of Melbourne.

David, Benjamin (2017), ‘Computer technology and probable job destructions in Japan’, *Journal of the Japanese and International Economies*, 43: 77-87.

Deloitte LLP (2013), ‘Agiletown: The relentless march of technology and London’s response’, accessed at: https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/uk-futures/london-futures-agiletown.pdf on 2 February 2019.

Durrant-Whyte, High, Lachlan McCalman, Simon O’Callaghan, Alistair Reid and Daniel Steinberg (2015), ‘The impact of computerisation and automation on future employment’, chapter 1.4 in Committee for Economic Development of Australia, *Australia’s Future Workforce* (CEDA, Melbourne).

Edmonds, Daniel and Timothy Bradley (2015), ‘Mechanical boon: Will automation advance Australia?’, Office of the Chief Economist, Commonwealth Department of Industry, Innovation and Science, Research paper 7/2015.

Frey, Carl and Michael Osborne (2013), ‘The future of employment: How susceptible are jobs to computerisation?’ Working Paper, Oxford Martin School, Programme on the Impacts of Future Technology, University of Oxford.

Frey, Carl and Michael Osborne (2017), ‘The future of employment: How susceptible are jobs to computerisation?’, *Technological Forecasting and Social Change*, 114: 254-80.

Goldacre, Ben (2008), *Bad Science* (Fourth Estate, London).

Gilovich, Thomas (1991), *How We Know What Isn’t So* (Free Press, New York).

Lawrence, Matthew, Carys Roberts and Loren King (2017), ‘Managing automation: Employment, inequality and ethics in the digital age’, Institute for Public Policy Research Commission on Economic Justice, accessed at: https://www.ippr.org/publications/managing-automation on 22/2/2019.

Le Roux, Daniel (2018), ‘Automation and employment: The case of South Africa’, *African Journal of Science, Technology, Innovation and Development*, 10(4), 507-15.

Lee, King (2016), ‘Automation, computerisation and future employment in Singapore’, Munich Personal RePEc Archive, mimeo.

Pajarinen, Mika, Petri Rouvinen and Anders Ekeland (2015), ‘Computerization threatens one-third of Finnish and Norwegian employment’, *ETLA Brief*, 34:22 April.

Rasmussen, Carl Edward and Hannes Nickisch (2010), ‘Gaussian processes for machine learning (GPML) toolbox’, *Journal of Machine Learning Research*, 11, 3011–15.

World Bank (2016), *World Development Report 2016: Digital Dividends* (World Bank: Washington DC).

Table A1**: Frey and Osborne (2017) hand-labelling of automatable occupations**

Automatable	**Not automatable**
Non-routine cognitive	
Surveyors; Cost estimators; Market research analysts and marketing specialists; Civil engineering technicians; Electrical and electronics drafters; Tax examiners, collectors, and revenue agents; Accountants and auditors; Farm labour contractors; Claims adjusters, examiners, and investigators; Credit analysts; Loan officers; Insurance underwriters; Technical writers.	Chief executives; Physicians and surgeons; Dentists, general; Clergy; Social and community service managers; Preschool teachers; Registered nurses; Marriage and family therapists; Education administrators, preschool etc; Civil engineers; Fashion designers; Substance abuse / behavioural counsellors; Lawyers; Landscape architects; Meeting, convention, and event planners; Healthcare practitioners and technical workers; Compliance officers; Electrical engineers; Physicists; Zoologists and wildlife biologists; Judges and magistrates: Economists; Transportation, storage, distribution managers; Athletes and sports competitors.
Routine cognitive	
Judicial law clerks; Human resources assistants, exc. p/t; Paralegals and legal assistants; Switchboard operators, inc. a.s.; File clerks; Data entry keyers; Credit authorizers, checkers, and clerks; Cashiers; Meter readers, utilities.	
Routine manual	
Sheet metal workers; Computer-cont. machine tool ops. m/p; Sewing machine operators; Electrical/electronic equip. assemblers; Bus drivers, transit and intercity; Light truck or delivery services drivers; Industrial truck and tractor operators; Taxi drivers and chauffeurs; Couriers and messengers; Motorboat operators.	Plumbers, pipefitters, and steamfitters.
Non-routine manual	
Dishwashers; Cooks, fast food; Parking lot attendants; Gaming dealers; Medical transcriptionists.	Childcare workers; Chefs and head cooks; Hairdressers, hairstylists, and cosmetologists; Concierges; Flight attendants; Waiters and waitresses; Maids and housekeeping cleaners; Hunters and trappers.

Table A2**: O*NET tasks for occupation 13-2011 Accountants and Auditors**

1. Prepare, examine, or analyze accounting records, financial statements, or other financial reports to assess accuracy, completeness, and conformance to reporting and procedural standards.
2. Report to management regarding the finances of establishment.
3. Establish tables of accounts and assign entries to proper accounts.
4. Develop, implement, modify, and document recordkeeping and accounting systems, making use of current computer technology.
5. Collect and analyze data to detect deficient controls, duplicated effort, extravagance, fraud, or non-compliance with laws, regulations, and management policies.
6. Prepare detailed reports on audit findings.
7. Supervise auditing of establishments, and determine scope of investigation required.
8. Report to management about asset utilization and audit results, and recommend changes in operations and financial activities.
9. Inspect account books and accounting systems for efficiency, effectiveness, and use of accepted accounting procedures to record transactions.
10. Examine records and interview workers to ensure recording of transactions and compliance with laws and regulations.
11. Examine and evaluate financial and information systems, recommending controls to ensure system reliability and data integrity.
12. Review data about material assets, net worth, liabilities, capital stock, surplus, income, and expenditures.
13. Confer with company officials about financial and regulatory matters.
14. Examine whether the organization's objectives are reflected in its management activities, and whether employees understand the objectives.
15. Prepare, analyze, and verify annual reports, financial statements, and other records, using accepted accounting and statistical procedures to assess financial condition and facilitate financial planning.
16. Inspect cash on hand, notes receivable and payable, negotiable securities, and cancelled checks to confirm records are accurate.
17. Examine inventory to verify journal and ledger entries.

Table A3: **Employment changes by occupation, 2010 to 2018 (per cent)**

Regressor	(1)	(2)	(3)	(4)	(5)	(6)
FO high risk of automation	-3.53 (3.81)					
FO probability of automation		-7.05 (3.87)	1.90 (4.73)		-0.63 (6.86)	
Routine cognitive			-16.82 (5.45)	-15.63 (4.07)		
Routine manual			-3.82 (3.96)	-2.82 (3.31)		
Non-routine manual			-0.54 (6.13)	0.27 (6.46)		
Routine task intensity (RTI) index					-0.131 (0.118)	-0.135 (0.087)
Constant	15.52 (1.77)	17.91 (2.28)	18.00 (2.50)	18.39 (2.32)	21.38 (5.30)	21.23 (6.35)
Occupations	701	701	701	701	666	666
R-squared	0.004	0.010	0.063	0.062	0.025	0.025

Notes: BLS employment by occupation in 2010 used as weights. Heteroscedasticity-robust standard errors provided in parentheses.

Table A4: **Employment changes by occupation and bottlenecks, 2010 to 2018 (per cent)**

Regressor	(1)	(2)	(3)	(4)
FO probability of automation	-7.07 (3.88)		-8.81 (3.58)	-1.96 (6.29)
Low perception and manipulation		-10.08 (5.29)	-11.19 (5.25)	-9.82 (5.59)
High perception and manipulation		-4.81 (4.94)	-5.19 (4.94)	-3.88 (4.62)
Routine task intensity (RTI) index				-0.142 (0.116)
Constant	17.89 (2.28)	18.96 (4.55)	17.38 (5.39)	27.58 (7.72)
Occupations	697	697	697	662
R-squared	0.010	0.024	0.039	0.049

Notes: BLS employment by occupation in 2010 are used as weights. Heteroscedasticity-robust standard errors provided in parentheses.

THE UNIVERSITY OF MELBOURNE

MELBOURNE INSTITUTE