



Making Sense of Predictions About New Technologies and Jobs in Developing Countries

Background Paper

Nicolas Lippolis



Nicolas Lippolis

Background Paper 13
February 2019

The Pathways for Prosperity Commission on Technology and Inclusive Development is proud to work with a talented and diverse group of commissioners who are global leaders from government, the private sector and academia. Hosted and managed by Oxford University's Blavatnik School of Government, the Commission collaborates with international development partners, developing country governments, private sector leaders, emerging entrepreneurs and civil society. It aims to catalyse new conversations and to encourage the co-design of country-level solutions aimed at making frontier technologies work for the benefit of the world's poorest and most marginalised men and women.

This paper is part of a series of background papers on technological change and inclusive development, bringing together evidence, ideas and research to feed into the commission's thinking. The views and positions expressed in this paper are those of the author and do not represent the commission.

Citation:

Lippolis, N. *Making Sense of Predictions about New Technologies and Jobs in Developing Countries*. Pathways for Prosperity Commission Background Paper Series; no. 13. Oxford. United Kingdom

www.pathwayscommission.bsg.ox.ac.uk

@P4PCommission

#PathwaysCommission

Cover image © Shutterstock.com



Abstract

Recent years have witnessed the proliferation of studies attempting to estimate the number of jobs at risk of automation. Studies of this kind have recently also been extended to the developing world. Although there is merit in bringing attention to the risks posed by technological developments, these analyses share common methodological problems and should be interpreted with caution. Instead, reports looking at the issue through a broader economic lens offer greater insights for developing countries, even as their arguments would benefit from firmer backing from rigorous economic research. This paper argues that further attempts to quantify exposure to automation in labour markets hold little promise, and that future research should be conducted more holistically, taking into account the variety of effects exerted by new technologies on the economies, politics and societies of developing countries.

Table of Contents

Introduction	2
1. Quantifying job losses due to technology	5
Frey and Osborne's machine learning method	5
Other varieties of quantification	12
Beyond technical feasibility: estimating job automation	16
Discussion: technology and public policy in developing countries	23
2. The economics of technology in developing countries	27
The UNCTAD Trade and Development Report 2017	27
The future of manufacturing-led development strategies	30
3. Conclusions	34
Modelling the effects of new technologies	34
Thinking holistically about technological paradigms	35
References	38

"Reflection on the forms of human life, hence also scientific analysis of those forms, takes a course directly opposite to their real development. Reflection begins *post festum*, and therefore with the results of the process of development ready to hand."

Karl Marx, Capital Vol. I (1976 [1867], p. 168)

Introduction

The question of how new automation and digital technologies will affect jobs around the world is increasingly gaining prominence, as evidence of these technologies increasing inequality is amassed.¹ New publications trying to estimate the anticipated economic changes caused by these technologies are constantly emerging, with overtones ranging from moderately optimistic² to highly alarmist.³ Most existing research and policy publications deal with the effects of new technologies on developed countries, which are the first ones to be affected by these trends. However, there is growing attention to the effect of new technologies on developing countries, and a recent spate of policy publications have attempted to gauge the effects that new technologies could potentially have on jobs in the developing world.

A broad question such as the effects of new technologies on developing economies is made up of many moving parts, and inevitably lends itself to a variety of approaches. Despite the pioneering contributions of David Autor, Daron Acemoglu, and other economists, very little academic research has attempted to shed light directly on how new automation and digital technologies will affect developing country labour markets. At the same time, there is growing demand for insights on how new technologies will affect developing economies and on the policy measures that can reduce any negative impacts. Consequently, the urgent need for predictions and guidance has led to the proliferation of work on the subject, as well as a large number of numerical predictions that are often cited in rather uncritical fashion.

This paper takes a step back from the debate on new technologies to make a critical assessment of existing strategies for analysing the effects of new technologies on jobs in developing countries. Table 1 lists the studies discussed in this paper, the main theme addressed by each one, and the countries included in the analysis (where applicable).⁴ Of course, given the sheer number of reports on this same topic, there could be a degree of arbitrariness in the choice. However, the publications analysed in this paper are among those that have received the most attention from the media and the policy world, and are broadly representative of what is currently available to policymakers and other decision-makers 'out there'.⁵ While the work of Frey and Osborne (2017), Arntz, Gregory

¹ Some of the seminal papers in this literature include Autor, Levy and Murnane (2003), Acemoglu and Autor (2011), and Autor and Dorn (2013).

² Some of the most prominent members of this camp include Joel Mokyr (2014a; 2014b; Mokyr, Vickers and Ziebarth 2015), Erik Brynjolfsson, and Andrew McAfee (Brynjolfsson and McAfee 2014; 2016).

³ See, for example, Cowen (2010), Vjig (2011), and Gordon (2016).

⁴ For a more encompassing review of this literature, see Schlogl and Sumner (2018), although they approach the topic from a different perspective and put forward their own theory.

⁵ *The 2019 World Development Report* is another study that could have been included. I have opted to leave it out because of its greater focus on policy design than on predictive analysis.

and Zierahn (2016), and Nedelkoska and Quintini (2018) does not deal directly with developing countries, their methodologies are applied in the studies conducted by Chang and Huynh (2016) and the McKinsey Global Institute (MGI), which do. Therefore, a discussion of their methodologies is essential for understanding the specific consequences of automation and digitisation for developing countries.

Table 1: Studies discussed in this report

Study	Theme and countries included (where applicable)
Frey and Osborne (2017)	Estimating US jobs' susceptibility to automation
Arntz, Gregory and Zierahn (2016)	Estimating the risk of automation in 21 OECD countries
Nedelkoska and Quintini (2018)	Estimating the risk of automation in 32 OECD countries
Chang and Huynh (2016)	Risk of automation and policy implications in the five original Association of Southeast Asian Nations (ASEAN-5) countries (Indonesia, Malaysia, Philippines, Singapore and Thailand).
PwC (2018)	Whether robots will steal jobs in developed economies.
World Development Report (2016)	Spread of digital technologies in the developing world and what can be done to harness their potential.
MGI (2017a)	Automation potential of the global economy, the factors that will determine the pace and extent of workplace adoption, and the economic impact associated with its potential.
MGI (2017b)	Expected workforce transitions and their implications.
UNCTAD (2017)	How to orient efforts to build inclusive economies and societies in today's global environment.
Hallward-Driemeier and Nayyar (2017)	The future of manufacturing-led development strategies in view of current and future technological progress.

The reports discussed here are also representative of the variety of existing approaches. There are notable differences in the questions asked by each report, and in the strategies used to answer them. Most public attention has centred on studies that attempt to directly estimate the number of existing jobs susceptible to being automated. Some of the studies building directly on Frey and Osborne's (2017) pioneering work focus solely on the automatability estimates, while others, such as the World Development Report (2016) and the reports by the MGI, combine estimates of automatability with more complex forecasts of the labour market. Another approach consists of taking a more aggregate perspective and examining the matter through the lens of traditional economic categories, as in the study by the United Nations Conference on Trade and Development (UNCTAD). Alternatively, the World Bank study produced by Mary Hallward-Driemeier and Gaurav Nayyar elaborates on new ways of approaching the issue of automation by identifying key economic variables that will determine the overall effect of new technologies on developing country economies. As will be shown in this paper, the different considerations brought up by each report provide evidences of the complexity of the questions being asked, particularly when this takes place in a context of fast-changing developing economies.

Given that we are still at an early stage in the employment of new automation and digital technologies, this paper seeks to answer the following question: what kind of analysis is the most informative for predicting the effects of these technologies on developing countries, and for drawing out their implications for citizens, governments and the private sector?

I compare the different strategies for coming to terms with the effects of automation and digital technologies on developing countries, and assess their relative strengths and weaknesses. Unlike the studies reviewed, and papers that include surveys of the existing literature, such as Schlogl and Sumner (2018), a 'meta-question' is posed here: instead of asking what the effects of new technologies will be, I ask how assessments of these effects differ, what factors might account for the differences, and what type of analysis can be the most useful for improving our understanding of current and future trends. Through this exercise, I seek to cover relatively new ground that is yet to be adequately addressed. This is largely due to the peculiar position of the debate about technology and jobs in the field of research. Despite the concern with the general question of the comparative value of different strategies, I delve into a relatively detailed analysis of the texts reviewed, since the overarching question cannot be adequately addressed unless we have an idea of the strengths and limitations of each text. Moreover, I believe that this exercise provides a public good, as the critical scrutiny of these studies has so far been limited.

While acknowledging that every approach to assessing the effects of new technologies on developing countries has its limitations, partly due to a lack of suitable data, I contend that estimates of job automatability are of limited utility. In the best-case scenario, such studies require highly complex models, which yield very uncertain estimates. At their worst, they might involve empirical strategies of limited credibility. Unfortunately, the latter have been the most common in the literature, and I argue that it would be best to direct researchers' energies towards other research strategies, while trying to incorporate some of the insights provided by this early literature. The paper then discusses the studies by UNCTAD (2017) and Hallward-Driemeier and Nayyar (2017). These bring up a host of new considerations on the macro-level impacts of automation on developing countries. In the conclusion, I argue in favour of a more holistic perspective for thinking about the effects of new technologies on jobs in developing countries, taking into account the variety of (often contradictory) effects they will have on the economy, politics and society. Armed with this perspective, there will be much greater scope to build on the insights provided by these reports through rigorous economic research.

1. Quantifying job losses due to technology

Among all the pieces of analysis attempting to understand the impact of new digital and automation technologies on developing countries, reports that try to directly quantify the number of jobs susceptible to automation have received the most attention. This approach yields precise numerical estimates, offering the clear advantage of providing simple and unambiguous statements of the potential risks posed by new technologies, thus facilitating the communication of their results. Even if we accept the now conventional approach to understanding the effect of technology on the labour market pioneered by Autor et al. (2003), the “task model” cannot, on its own, reveal the magnitude of the risks to employment created by new automation technologies. Thus, methods for estimating jobs at risk of automation can potentially be very pertinent, as they can allow us to assess the relative magnitudes of the different channels through which new technologies may affect the economy.

However, economic forecasts are difficult to perform correctly, and even more so when there is very little historical data with which to generate a forecasting model. One solution to this problem is provided in a seminal paper by Frey and Osborne (2017).⁶ To overcome the problem of forecasting jobs' automatability in the absence of historical data showing the relationship between technological change, job characteristics and automation, they make use of an innovative strategy that combines expert opinion with machine learning methods. Given that their approach has been highly influential and was subsequently replicated in many other studies, in this section, I critically examine their methodology, before moving on to discussing similar papers.

Frey and Osborne's machine learning method

The stated aim of Frey and Osborne's study is to predict how susceptible existing jobs are to “computerisation” – a term that they employ interchangeably with the more common “automation” – including technologies related to machine learning, such as “Data Mining, Machine Vision, Computational Statistics and other sub-fields of Artificial Intelligence (AI), in which efforts are explicitly dedicated to the development of algorithms that allow cognitive tasks to be automated” (Frey and Osborne 2017, p. 258). The study also discusses Mobile Robotics. They note that, contrary to common applications of Autor et al.'s (2003) task model of the labour market, these new technologies open up the possibility that machines replace humans in non-routine tasks, in addition to the routine tasks that had hitherto been seen as machines' preserve. This stems from the fact that machine learning and related techniques do not require *a priori* codification of a procedure in order to work; instead, as long as they have sufficient data and are provided with a measurable definition of success, algorithms can themselves estimate a solution for a given task. In view of these technological developments, predictions based on a decomposition of jobs into routine and non-routine tasks are no longer sufficient for assessing a job's automatability.

⁶ Frey and Osborne's study was first published as a working paper in 2013, and most other papers cite the working paper version. Here, I cite the published version, although they are identical.

To overcome these perceived limitations in the task model, as well as the abovementioned lack of historical data with which to estimate a forecasting model, Frey and Osborne “automate” the task of making a prediction of automation by adopting these same machine learning methods.⁷ They restrict themselves to estimating the technical feasibility of automating different jobs over the next one or two decades. Training data for their algorithm was created by gathering experts at an Oxford University engineering workshop and asking them to give their best assessment of whether or not it is technically feasible to automate jobs in a list of occupations from the US Department of Labor’s O*NET dataset. This led to 70 of the jobs in the dataset being labelled as automatable or not. They posit that a job’s automatability is a function of the degree to which it requires perception and manipulation, and creative and social intelligence, which still constitute “bottlenecks to computerisation”⁸ and use some of the O*NET data on jobs’ task content and their knowledge and skills requirements to score jobs in each of these dimensions. Finally, they use the 70 labelled jobs and their respective embodiment of computerisation bottlenecks to train their machine learning algorithm to estimate a job’s automatability as a function of these bottlenecks. Once the algorithm has been trained, they are able to apply it to the entire dataset of 702 jobs, leading to their estimate that 47% of US jobs are at high risk of automation.

The underlying rationale for Frey and Osborne’s strategy is that since technological advances are hard to predict, and since there is no obvious way in which to model these advances as a function of computerisation bottlenecks (ie. to *codify* the drivers of advances in engineering), delegating the assessment of a job’s automatability to experts and extending their assessments to a broader universe of jobs can allow us to harness this tacit knowledge to make more general predictions.

Unfortunately, the robustness of Frey and Osborne’s methodology and the validity of their results can be questioned due to the limited explanation given for some of the paper’s core methodological procedures. The whole exercise is based on the experts’ subjective assessment of whether jobs are automatable in the near future. Very little justification is given for the validity of this approach, except for a reference to a paper by Blinder (2009), who similarly assigns a subjective index of “offshorability” to jobs in the O*NET dataset. As Frey and Osborne themselves admit, a subjective approach runs the risk of not being replicable, and “may result in the researcher subconsciously rigging the data to conform to a certain set of beliefs” (Frey and Osborne 2017, p. 263). One might argue that, for this reason, transparency becomes essential. However, as a comparison of the research shows, Frey and Osborne fall short of Blinder in scrutinising the judgments of assembled experts. Readers are given very little information on the process used to make these assessments. For instance, Blinder’s paper has a section on ‘Ground Rules’, where he clearly outlines the assumptions embedded in his subjective approach and explains why it was chosen over a more objective labelling of occupations. He makes it very clear that his projections are for the medium term; that they extrapolate “normal technological progress”; that he does not account for future changes in the occupational distribution of US employment; and that he uses an *ordinal* rather than a *cardinal* scale. The bulk of Blinder’s

⁷ For accessible introductions to the use of machine learning methods in economics, see Varian (2014) and Mullainathan and Spiess (2017).

⁸ Despite the overtones of radical innovation of Frey and Osborne’s description of machine learning and related technologies, one cannot fail to note that their list of “bottlenecks to computerisation” is very similar to the standard list of non-routine tasks, which includes “flexibility, creativity, generalized problem-solving, and complex communications” (Autor et al. 2003, p.6).

paper goes to great length to explain the process for subjectively labelling each occupation. He clearly points out the flaws in previous attempts at providing objective rankings of offshorability, but still cross-checks his subjective rankings with an objective one, as well as with subjective rankings by an experienced human resources professional and a large replication study by a team at the Harvard Business School. He finds rank correlations of 0.16, 0.38 and 0.6, with each of these other assessments, respectively.

Contrasting Blinder's paper with Frey and Osborne's approach clearly evidences the latter's comparative lack of transparency. To start with, little is said about how the group of experts was assembled, who these experts were, the process through which they reached a decision regarding each job's automatability, and how assessments varied between them. Had there been large discrepancies between experts' judgments, readers would be a little less confident of the accuracy of the predictions. This point is particularly important in view of research showing the limitations in experts' ability to perform accurate forecasts in a variety of different fields (Kahnemann 2011). Frey and Osborne themselves cite a paper by Armstrong and Sotala (2012) on experts' inability to make predictions on AI, but rationalise their choice by claiming that, unlike the experts in Armstrong and Sotala's paper, they make no attempt to predict the number of years that it will take to overcome various engineering bottlenecks to computerisation, thereby missing the more general point about the difficulty of making forecasts. Frey and Osborne's strategy for dealing with the fallibility of subjective forecasts consists of only labelling occupations that they are "highly confident about" (Frey and Osborne 2017, p. 264), claiming this attenuates the risk of subjective bias. This does little to assuage those fears, however, since confidence is by no means incompatible with bias.

The presence of subjective bias in the experts' predictions could compromise Frey and Osborne's results. As argued by Mullainathan and Obermeyer (2017), in the presence of systematic mismeasurement of the outcome variable in an algorithm's training sample, the use of the algorithm to make predictions for the broader population may generalise the biases in the algorithm's predictions. In Frey and Osborne's study, this could be the case if the experts were more likely to misjudge the automatability of jobs with certain characteristics. For example, if they were biased towards judging jobs that they have first-hand experience of as non-automatable, then the algorithm would reflect this bias when predicting automatability. It would do this by giving a larger weight to the computerisation bottlenecks that are correlated with first-hand experience of the occupation. The final estimation would extend this source of bias to the entire sample, as occupations with higher values along these correlated dimensions would have their automatability underestimated. Given the highly speculative nature of the experts' assessment, and in the absence of more detail on the process through which these assessments were made, we cannot rule out the possibility that such forms of systematic bias are present in Frey and Osborne's results.

Another issue with Frey and Osborne's approach to labelling occupations' automatability is that, unlike Blinder's ordinal scale of offshorability, jobs in their study are labelled according to whether they are *fully* automatable, given current technological capabilities. The problem here is that jobs that fall short of being fully automatable are not differentiated according to their degree of automatability, and are all labelled simply as a "0". It is straightforward to see that, if there is some underlying relationship between the bottlenecks to computerisation and occupations' susceptibility to automation, placing occupations with different degrees of automatability in the same class and training a machine learning algorithm with these labels might negatively affect the performance of the algorithm when applied to data outside the training sample.

A more benevolent interpretation of this statement could be that they only label occupations at the extremes of the distribution of automatability, and that occupations assigned a "0" are those which would be assigned very low levels of automatability in a continuous scale. Indeed, viewing the list of occupations in the paper's appendix, we can see that most of the jobs assigned a "0" are ones that most people would consider eminently not automatable, such as dentists, nurses, fashion designers, athletes, and childcare workers. Thus, it would not create much harm to class them all in the same category, since they are all clustered at the lower end of the automatability scale.

Unfortunately, this more benevolent interpretation can be seen to create further complications. One of the key advantages of machine learning techniques, as compared to traditional econometrics, is that they require very few assumptions about the underlying data. One of the few assumptions that *is* needed is that the variables in the training set have the same joint distribution as the variables in the test set (Athey 2018).⁹ By extension, we can also reason that the algorithm is only likely to make relatively accurate predictions if the variables in the independent dataset in which we are trying to make predictions – in this case the O*NET dataset containing all 702 occupations – also have the same joint distribution. One of Frey and Osborne's justifications for making use of machine learning methods is that the relationship between engineering bottlenecks is likely to be non-linear. But if that is the case, then the relationship between X (the degree to which occupations involve tasks that still constitute computerisation bottlenecks) and Y (their degree of automatability) at the extremes of the distribution of Ys might be different than in the middle ranges of the distribution. It would then follow that the joint distribution of X and Y differs between the training set, made up of observations at the extremes of the distribution of Y, and the sample to which the algorithm is applied (that is, the 632 jobs not included in the training set), violating one of the key assumptions required for machine learning methods to work. Therefore, the algorithm estimated using the training set would not provide very accurate estimates of the automatability of the rest of the occupations in the O*NET dataset. To use machine learning jargon, the differences in the joint distribution of variables in the training set and in the population would be a source of variance. This means that the function linking computerisation bottlenecks and probability of automation would be very different than if it were estimated with a different set of occupations, further undermining our confidence in their results.

Limitations can also be observed in Frey and Osborne's estimation procedure. The first step is selecting the variables that they will use to predict an occupation's likelihood of being automated. They choose the degree to which occupations require social intelligence, creativity, and perception and manipulation tasks. They then select nine variables from the O*NET dataset that describe these attributes. Each occupation is ranked on a three-point scale (low, medium or high) for each of the nine variables. These rankings then form covariates used for estimating their algorithm. Here, the discussion becomes more ambiguous. While Figure 1 in the paper suggests that each occupation is assigned a value from 0 to 100 for social intelligence, creativity, and perception and manipulation, the discussion in section 4.2 suggests that the nine variables are used for estimation. However, there is no discussion of how these variables are scaled: does each job receive a value of 1, 2 or

⁹ Athey (2018, p.4) states that "The observations are assumed to be independent, and the joint distribution of X and Y in the training set is the same as that in the test set. These assumptions are the only substantive assumptions required for most machine learning methods to work."

3, corresponding to low, medium and high levels on the O*NET scale? Or is it 0, 50 and 100? Or are they scaled in some other way, perhaps with more granular distinctions? This is a potentially important omission since, unlike most traditional econometric estimators, in machine learning, the scale of variables matters for the predictions one obtains (Mullainathan and Spiess 2017). Once more, the study's opacity compromises the confidence one can have in the results.

After scaling, the next step in the estimation consists of choosing the probability classification algorithm to be employed. Here, unlike most academic papers that use machine learning methods, there is no in-depth discussion of the model chosen for the discriminant function.¹⁰ Frey and Osborne consider three types of discriminant function: a simple logistic function and two variants of 'Gaussian process classifiers', including the 'exponentiated quadratic' and the 'rational quadratic'. They claim that the Gaussian process classifiers are used to account for non-linearity, but do not discuss how the choice of non-linearity for the discriminant function varies according to the nature of the decision boundary (that is, the boundary between an observation being classified as '0' or '1' in the covariate space).¹¹ They make no mention of some of the algorithms most commonly employed in machine learning applications – such as tree-based methods – and do not include discussion on any other criteria for model choice. While they compare their three chosen models in terms of how accurately they predict occupations' probability of automation, they do not account for standard considerations that are usually made with respect to the level of complexity of the discriminant function or the bias-variance trade-off. This is puzzling, given that together they constitute some of the most important choices to be made when using this method.¹²

Another puzzling aspect of Frey and Osborne's paper is the superficial way they deal with empirical tuning and validation. At this stage, most applications of machine learning use methods such as 'cross-validation' and 'bootstrapping' (see below) to increase the confidence that the model fit is not dependent on the particular sample used.¹³

To validate their classifiers, Frey and Osborne divide their sample of 70 occupations into two sets of equal size, where one half is used as a training set and the other half as a test set. For each of 100 random selections of training and test sets, they test the accuracy of their classification algorithm in predicting occupations' labels, as measured by the area under the curve (AUC) – a standard measure for predictive accuracy – and the log likelihood. They report the average score of each classifier for each criterion, but do not report their standard deviation.

This would be important in assessing the classifiers' variance (that is, the degree to which the estimate parameters vary depending on the training sample). Moreover, it is well known that splitting the sample in half is not a particularly efficient form of cross-validation, since it does not make use of the

¹⁰ Compare this, for example, with the information reported in the Appendix to Blumenstock, Cadamuro and On (2015) and in Bjorkegren and Grissell (2017).

¹¹ See James et al. (2013, pp. 37-42 and 184-186) for a simulation demonstrating the importance of the nature of the decision boundary for the choice of the form of non-linearity in the discriminant function.

¹² At least if judged by the space dedicated to it in most publications using machine learning methods or providing advice on how to use them. See for example Domingos (2012) or Table 2 in Mullainathan and Spiess (2017).

¹³ Once more, the comparison with Blumenstock et al. (2015) and Bjorkegren and Grissell (2017) is instructive.

available data in the most efficient way.¹⁴ Instead, the standard forms of cross-validation in machine learning studies are techniques such as leave-one-out cross-validation (LOOCV) and k-fold cross-validation (where k usually equals 5 or 10), which are better able to harness the information available in the data. The efficiency of data use should be of particular concern for a study such as Frey and Osborne's, in view of their very small training dataset ($N=70$) and the overwhelming importance of dataset size for the performance of machine learning methods (Halevy, Norvig and Pereira 2009).¹⁵

Bootstrapping is another commonly used technique that would have helped increase the confidence in the robustness of Frey and Osborne's results. It consists of randomly drawing a set of observations of size n (where n is the size of the entire testing sample) from the testing sample and calculating a test statistic for it – in this case, the AUC or the log likelihood. This is repeated a large number of times to obtain a sampling distribution of the test statistic, allowing us to construct confidence intervals for the sampling statistic and thus “quantify the uncertainty associated with a given estimator or statistical learning model” (James et al. 2013, p. 187). Other methods that can be used for similar purposes include 'bagging', 'boosting', or a combination of these methods in 'random forests' (Varian 2014). Although it is unreasonable to expect studies to apply every existing method to increase confidence in the predictive power of a chosen algorithm, the fact that Frey and Osborne present only a small part of their results, make no mention of even the most standard practices, and provide so little explanation adds to the general feeling of opacity that permeates their study.

The unorthodox methodologies in Frey and Osborne's study would arguably have benefitted from increasing the clarity of the study's objectives and the interpretation of its results. In Section 4 of the paper (Frey and Osborne 2017, p. 263), it is stated that:

The hand-labelling of the occupations was made 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”. Thus, we only assigned a 1 to fully automatable occupations, where we considered all tasks to be automatable.

The way this passage is phrased suggests that rather than making a *prediction* of automatability, they are simply assessing whether existing capabilities allow a given occupation to be *currently* automated. Elsewhere in the paper, however, the phrasing of the text is more ambiguous, and could also be read as suggesting that the results are forecasts of *future* technological feasibility. In Section 1, for example, it is claimed that their methodology allows them to “examine the future direction of technological change in terms of its impact on the occupational composition of the labour market, but also the number of jobs at risk *should these technologies materialise*” [my emphasis] (Frey and Osborne 2017, p. 255). This suggests that they are trying to predict technological developments, and this would be at odds with the question asked to the experts.

¹⁴ James et al. (2013, pp. 176-183) perform simulations comparing the accuracy of estimates of the mean squared error for the “validation set approach” (that is, splitting the sample in half), LOOCV and k-fold cross-validation. They find that the validation set approach invariably yields the least accurate estimates of the mean squared error.

¹⁵ The appendix to Blumenstock et al. (2015) graphically illustrates how the accuracy of their estimates increases as the size of the training set increases.

It could be argued that the forward-looking aspect of their method does not have a bearing on the jobs judged as automatable by the experts, which presumably are already automatable, but on jobs at intermediate levels of automation. One could reason that, if automatability is a function of an occupation's embodiment of the computerisation bottlenecks, then a given occupation's "susceptibility to automation" (the expression most commonly employed by Frey and Osborne) will depend on the extent to which its requirements for each of the attributes that constitute computerisation bottlenecks resemble those of the jobs that are already known to be automatable. It is precisely this type of similarity between the occupations labelled by the experts at the engineering workshop and the rest of the occupations in the O*NET sample that the machine learning algorithm is meant to measure. By measuring this degree of resemblance, the automation probabilities yielded by the algorithm would allow us to establish a rough timeline – at least in terms of an ordering between jobs – for how close we are to the point when it will be possible to automate a given job.

An issue here is that the term "susceptibility to automation", which presumably is what the algorithm measures, is itself quite ambiguous. A job can be said to be "susceptible to automation" because: (i) the technology required to automate it already exists, but is yet to be deployed to substitute for workers; or (ii) the necessary technology does not currently exist, but given the nature of the occupation, we are able to make a prediction as to how likely and/or how soon it will develop. Interpretation (ii) contradicts Frey and Osborne's claim that they "avoid making any predictions about the number of years it may take to overcome various engineering bottlenecks to computerisation." (Frey and Osborne 2017, p. 268). If their estimation does not include any prediction of technological development, and does not otherwise incorporate any time component, then one cannot interpret "susceptibility to automation" in terms of how long it will take for the required technology to develop. This leaves us with interpretation (i), which contradicts Frey and Osborne's claim that their probability estimates "describe the likelihood of an occupation being *fully* automated" [my emphasis] (Frey and Osborne 2017, p. 268): at any point in time, given existing technology, an occupation is either fully automatable or it is not.

Admittedly, the line between current and future technological capabilities is more tenuous in machine learning than in other fields, since technological feasibility depends on the existence of an algorithm that is capable of learning a given task, as opposed to achieving full codification of a task. One could also conceive that the technology to replace workers in a task already exists but is yet to be deployed in a marketable way. However, these are points that could have been addressed through greater clarity when posing the question to the experts and when describing their results. Moreover, even if we think some existing machine will eventually be able to perform a given task, if this performance has not been achieved yet, then we are making a forecast of technological development, which Frey and Osborne themselves acknowledge is extremely difficult and not the aim of their paper.¹⁶ We can thus conclude that, even if we accept the validity of their methodology, the way Frey and Osborne frame their study renders their probability estimates less convincing. Unfortunately, this point often gets lost due to the ambiguity of the language used to frame and interpret their study.

¹⁶ See Makridakis, Spiliotis and Assimakopoulos (2018) for evidence on the (weaker) performance of machine learning when compared to other forecasting methods.

Through the present discussion, I have sought to highlight some of the limitations in Frey and Osborne's study on jobs' susceptibility to automation. Despite the value of pointing out how recent innovations in machine learning and related fields can alter received notions on the relationship between occupations' task content and their susceptibility to automation, their methodology fails to convincingly estimate jobs' exposure to automation. The description of the labelling exercise that forms the basis of their study lacks transparency; the risks posed by subjective bias are not adequately addressed; the algorithm is estimated using unorthodox methodological choices, offering little explanation; the usual procedures used to build confidence in the reliability of results are lacking; and an ambiguous use of language ultimately reveals contradictions in how the results are interpreted.

I have chosen to dedicate some space to the critique of Frey and Osborne's study due to the centrality it has acquired in the literature on the effects of new technologies on labour markets. It is important to emphasise the novelty of Frey and Osborne's approach, as they are among the first analysts to go beyond the task model and realise that the relationship between job characteristics and automatability in the machine learning age may be more complex than hitherto imagined. Their methodology for assessing jobs' likelihood of being automated is highly innovative, as they are among the first analysts to make use of the rich O*NET data on job characteristics to make labour market predictions for the next wave of technological progress. Accordingly, they inspired a series of subsequent efforts to grapple with the question. However, despite being a landmark study in the field, I have argued that it falls short of providing reliable evidence for some of its claims, which may also benefit from clearer articulation. These points are often missed amid the clamour for insights on the effects of new technologies on developing countries, leading to frequently acritical reporting of their results. In addition, the influence achieved by the study has set the standard for subsequent empirical work on the topic, thereby reproducing some of the same shortcomings. I discuss the effect on other literature in the next section.

Other varieties of quantification

Having discussed the seminal work of Frey and Osborne, I now turn to an analysis of other papers attempting to quantify job losses due to new technologies strictly from a technological capabilities point of view. Table 2 compares the studies that make use of this sort of estimation, detailing the countries included in the study, the data used, and the probabilities of automation found. Despite the similarities in the data sources and the way each study is set up, there are large discrepancies in the results, stemming from methodological differences. However, as I argue here, the fact that they all build on Frey and Osborne's study – and thus inherit some of its limitations – combined with their own methodological shortcomings, lend little confidence about the reliability of their results.

Table 2: Comparison of studies estimating jobs at risk of automation

Study	Countries included	Data used	Automation probabilities
Frey and Osborne (2017)	US	O*NET database	47% of US jobs at high risk of automation
Arntz et al. (2016)	21 OECD countries	Frey and Osborne automatability estimates plus the Programme for the International Assessment of Adult Competencies (PIAAC) database	On average, across OECD countries, 9% of jobs are automatable
Nedelkoska and Quintini (2018)	32 OECD countries	Frey and Osborne automatability estimates plus PIAAC database plus job-task surveys for the UK and Germany	About 14% of jobs in OECD countries included in the study are highly automatable
PwC (2018)	27 OECD countries plus Singapore and Russia	Frey and Osborne automatability estimates plus PIAAC database	Risk of automation varies from 20–25% in highly educated East Asian and Nordic economies to over 40% in some Eastern European countries
Chang and Huynh (2016)	ASEAN-5 countries	A variety of labour force surveys, harmonised with Frey and Osborne's data	56% of all employment at high risk of automation in the next couple of decades
World Development Report 2016	A large number of developing and developed countries	Skills Towards Employment and Productivity surveys for 10 countries and one Chinese province, and the International Labour Organization (ILO) Labor Statistics (Laborsta) database and various World Bank sources	Two-thirds of jobs in the developing world are susceptible to automation
MGI (2017a)	46 countries representing more than 80% of the global economy (developing and developed)	Core of the analysis uses World Bank, US Bureau of Labor Statistics and O*NET data	Around 60% of occupations have at least 30% of their activities that are automatable
MGI (2017b)	46 countries representing more than 80% of the global economy (developing and developed)	A large variety of sources, including data from the UN, ILO and OECD	Between 3–14% of the global workforce might need to switch occupational categories by 2030

The reports generally have a similar preamble, where they identify a shortcoming in Frey and Osborne's approach and discuss how their method corrects it. In Arntz et al. (2016), the perceived shortcoming is Frey and Osborne's inattention to the fact that jobs involve a variety of tasks, and that jobs with the same occupational title can often differ in the mix of tasks they involve. Therefore, they make use of the newly assembled (at the time) Programme for the International Assessment of Adult Competencies (PIAAC) dataset, which contains detailed information on occupations' task content at the individual level and is comparable across countries. They match Frey and Osborne's automatability estimates (it is unclear whether they use the labelling of the 70 occupations or the final probability estimates) to individual US jobs in the PIAAC dataset through a 'multiple-imputation' approach, as only two-digit International Standard Classification of Occupations (ISCO) codes are available in the PIAAC. Then, using the matched jobs data, they estimate susceptibility to automation as a function of workplace tasks, demographic and economic variables. Using the coefficients estimated for the US, they calculate automation probabilities for jobs in other countries.

Arntz, Gregory and Zierahn's inclusion of demographic and economic variables to predict automatability is puzzling, given that the dependent variable, Frey and Osborne's experts' predictions, was explicitly formulated holding only technical feasibility to account; the actual substitution of workers by machines was deemed too complex a question to be tackled. Even more puzzling is the fact that, although they criticise the subjective labelling of occupations' automatability,¹⁷ they make use of these labels, or at least of the probability estimates derived from these labels. Frey and Osborne's experts' labels are made for occupations in general, as described in the O*NET database. It naturally follows that the only variables that can be scrutinised to derive a model predicting automatability (or at least the label of 'automatable') are those used to decide on the label in the first place, as done by Frey and Osborne. These labels are not predictions on the automatability of every single job with the same title, and therefore there is a mismatch between the dependent variable (the labels of job automatability) and the independent variables, even if we ignore the inclusion of economic and demographic variables. To put it more plainly, we cannot try to derive a relationship between workplace tasks and automatability by combining Frey and Osborne's labels with the PIAAC data, since we are not able to tell if every specific job in the dataset bearing the same title would have been labelled in the same way by Frey and Osborne's experts. In view of this, the coefficients derived by Arntz et al. can at best be treated as indicative of the workplace, economic and demographic characteristics of US workers currently in the occupations considered more susceptible to automation. At worst, these coefficients are simply meaningless, since there is a mismatch between the PIAAC data and Frey and Osborne's labels, and since we may consider the underlying data used to estimate automatability unreliable.

The report by Nedelkoska and Quintini (2018) resembles Arntz et al. (2016) in that, despite being a critique, it is heavily based on Frey and Osborne. The paper also takes inspiration from Arntz et al., and claims that it builds on their study by broadening the country coverage and by identifying more closely Frey and Osborne's engineering bottlenecks. Nedelkoska and Quintini's methodology essentially consists of a logistic regression with Frey and Osborne's labels as dependent variable and the variables from the PIAAC dataset that most resemble the engineering bottlenecks as independent variables. Strangely, in spite of claiming that "the current study closely aligns in

¹⁷ "...the approach still reflects technological capabilities based on experts' assessments rather than the actual utilisation of such technologies which might lead [to] an overestimation of job automatability [sic]" (Arntz et al. 2016, p. 21).

spirit to the assessment of the potential automation deriving from the development of Machine Learning" [sic] (Nedelkoska and Quintini 2018, p. 22) they entirely ignore the possibility of a non-linear functional form, which is the main reason why Frey and Osborne decide to employ a machine learning method in the first place. More generally, little thought is given to questions of functional form or robustness of the results. They note that, by estimating their coefficients using Canadian data, they may be biasing the results due to the specific characteristics of the Canadian economy, but brush these concerns aside by simply stating that their coefficients had the same sign and significance as those of Arntz et al., except for the frequency of solving simple problems. Here, they make no reference to the sizes of these coefficients, which would have been an essential comparison for judging how consequential the use of Canadian data is for their results. In any case, this again is a highly problematic estimation procedure, due to its dependence on the labels assigned by Frey and Osborne's machine learning experts, as well as the mismatch between these occupational-level labels and individual level data, which is compounded by the use of only one country to perform the estimation.

The two OECD studies discussed in this section so far build heavily on Frey and Osborne, but their results vary for methodological reasons. The two other studies discussed in this section are also based on Frey and Osborne's work, and closely follow the methodology of the OECD studies. PwC (2018) builds on previous work done in PwC (2017) and recreates Arntz et al's (2016) dataset, but uses a different set of predictive features for their Expectation–Maximization algorithm. They claim that their predictive features improve the performance of their classifier. Also, simulating the results using different sets of predictive features, finds that "the more predictive the model, the higher the estimation of high automatability jobs." (PwC 2017, p. 47) The model is calibrated for the US and then extended to the UK, Germany and Japan in PwC (2017) and to 29 countries PwC (2018). Also, in the later study, "the initial set of labels, seeded from the study by FO,¹⁸ were simulated across a range of scenarios that varied the automation-rate estimates associated with both tasks and occupations" (PwC 2018, p. 40). This simulation is combined with output from a computable general equilibrium model of the economic impact of AI to predict the outcomes of various 'waves' of automation.

Chang and Huynh (2016) consists of a variation of Nedelkoska and Quintini's OECD study (2018).¹⁹ Chang and Huynh perform a logistic regression with data from Frey and Osborne as the dependent variable, but they differ in using the probability estimates for the 702 occupations, as opposed to the 70 labels from the machine learning experts. Like Arntz et al. (2016), they include variables beyond unrelated to technical feasibility, such as socio-demographic variables, status of employment, economic sector, and even job experience. Using the coefficients thus derived, the authors calculate a fitted value for each individual job's probability of automation across the five countries studied.

In spite of their methodological differences, the PwC and Chang and Huynh studies fail to convince for much the same reasons as the other OECD studies. They are heavily based on Frey and Osborne's labels and do not make proper use of them, since there is a mismatch between the logic behind these labels and the data with which it is combined in the estimation. Further consideration could

¹⁸ The 'FO' acronym refers to the Frey and Osborne paper.

¹⁹ The fact that the International Labour Organization (ILO) study precedes the work of Nedelkoska and Quintini implies that it cannot have taken inspiration from the latter, and I describe it as "a variation of Nedelkoska and Quintini's OECD study" for ease of exposition, and simply to highlight their similarities.

have been given to the inclusion of variables in the estimation and the choice of functional form. Finally, there is the seamless projection of the results obtained in one country onto other countries with no real consideration on how this might depend on the economic environment. With these limitations in mind, the next section discusses the different strategies used by the WDR and MGI reports to deal with the difficulties involved in estimating potential job losses due to automation.

Beyond technical feasibility: estimating job automation

The reports discussed in the previous section all focus, at least nominally, on the technical feasibility of job automation. However, we often find a number of other observations and qualifications on the labour market effects of technology scattered in their reports. For instance, Arntz et al. (2016) make the following comments on Frey and Osborne's work:

- (i) Experts tend to overestimate the potential of new technologies
- (ii) The classification of tasks into 'routine' and 'non-routine' is often problematic
- (iii) Even in a highly competitive economy such as Germany, firms' familiarity with new technologies is actually quite low, further slowing down the rate of technology adoption
- (iv) There may be a lack of workers qualified to handle these new technologies
- (v) Actual automation, as opposed to mere technological feasibility, depends on a host of economic, ethical, legal and societal factors
- (vi) Even when automation technology is introduced in the workplace, it may substitute for some jobs but complement others, and the end results might be a change in the mix of tasks performed at work
- (vii) The creation of labour-saving technologies might itself generate new jobs
- (viii) Income effects due to the introduction of new technologies might dominate their substitution effects, which can lead to increased employment or wages.

These are all valid observations, although they may not count as valid critiques of Frey and Osborne, who frame their work exclusively in terms of technological feasibility. Moreover, point (ii) and the references made in point (i) miss the mark by failing to acknowledge the loss of centrality of the distinction between routine and non-routine tasks, which Frey and Osborne see as the main innovation brought forth by machine learning and related technologies (though see Footnote 8, above).

Nonetheless, Arntz et al.'s list has value in encapsulating the main intervening variables that determine how technological change is reflected in the economy. Similar comments are present, for example, in Nedelkoska and Quintini (2018, p. 8), who note that the actual rate of technology adoption "could be influenced by several factors, including regulations on workers dismissal [sic], unit labour costs or social preferences with regard to automation". They also mention that some types of technology require an enabling environment for their operation. One example is electric cars, which require complementary infrastructure such as charging stations. PwC (2018) echoes Nedelkoska and Quintini on the importance of complementary infrastructure, and also highlight technological constraints such as the need to integrate and adapt technology into solutions before deploying it, economic constraints due to the cost of deploying new technologies, and the legal and regulatory mechanisms that need to be developed before firms can adopt new technologies. Finally, Chang and Huynh (2016) highlight some issues specific to developing countries, such as the prevalence of informal work arrangements, which "are less likely to adapt in response to technology innovations because of capital investment constraints" (Chang and Huynh 2016, p. 8). It also notes that the high prevalence of agricultural work may amplify the risk of automation relative to more industrialised countries, given the "extremely routine and manual nature of agriculture work" (Chang and Huynh 2016, p. 8). Automation will be held further back by the lower wages and skill levels of workers in ASEAN countries, as well as by social preferences. On the latter, they mention the example of the hotel industry where, despite a propensity to automation, surveys of industry leaders indicate that the demand for human services is increasing.

The studies discussed so far make no attempt to include these considerations in their estimation procedures. However, if our ultimate interest is in the effects of new automation technologies on the economy, rather than just their technical feasibility, we might consider it important to somehow integrate these factors into our analysis. Yet, the challenge that still behoves the research community is thinking about how to combine this multitude of factors to form a sensible prediction of the effect of new technologies on jobs worldwide. One possible strategy for dealing with this challenge is to attempt to quantify the variety of factors that may matter for job automation. But as we will see in this section, this is no easy feat, and the outcomes could be so uncertain as to be essentially meaningless.

The World Bank's World Development Report 2016 (WDR 2016) offers one strategy for directly addressing the labour market impacts of automation. Drawing from Santos, Monroy and Moreno's (2015) background paper, the WDR calculates automation probabilities at the country and sector level, adjusted for adoption time lags based on Comin and Hobijn's (2010) estimates. For each occupation, they take the automation probability from Frey and Osborne, after the occupational categories have been harmonised, and adjust it to account for the expected adoption lags. The latter are calculated by assigning countries to a percentile in the distribution of adopters according to their income levels. High-income countries are predicted to take 30 years to fully automate jobs at risk of automation; upper-middle-income countries are assumed to be at the 10th percentile of the distribution of adopters, based on the average distribution of adoption lags among 20th-century technologies from Comin and Hobijn (2010), and therefore take an additional 11.71 years; the corresponding figure is 18.57 years for lower-middle-income countries; and finally, lower-income countries are predicted to fully automate jobs at risk 28 years after high-income countries have done so.

With these numbers in hand, they are able to make the adjustment for each country. They offer the example of the adjustment made for Bolivia, where 61% of occupations are at risk of being automated.²⁰ As it is a lower-middle-income country, the expected time lag for full automation is 18.57 years. Based on this, they infer that the adjusted number of jobs at risk of automation is 38%. The logic behind this adjustment could be clearer, however, as the paper does not explain the link made between adoption time lags and the share of jobs at risk of automation.

One way of making sense of these numbers is by assuming that they are meant to indicate the share of jobs at risk of automation at a two- to three-decade horizon, based on statements made throughout the paper. In developed countries, the share of jobs at risk of automation that will be automated in this time horizon is 100%, given their assumption that full automation of an occupation at risk takes 30 years. The adjustment lag for developing countries would then imply a slower process of automation. Thus, for the case of Bolivia, the expected time it would take for all jobs to be automated is 48.57 years. If we further assume that the figure for 'adjusted risk of automation' denotes the number of jobs that will be automated after 30 years, and if the relationship between technology adoption and time is linear, then, out of the 61% of jobs in urban Bolivia at risk of automation, 61.8% of them (which is equal to $30/48.57$) will have been automated 30 years from now. This calculation yields the figure of 38% reported by the paper.

Besides the fact that the reasoning above is just an imputation, and that none of this is explained in Santos et al. (2015), the validity of the underlying logic hinges on a number of questionable assumptions. We are not told where the figure for 30 years for full automation of an occupation comes from. The authors "assume that the developed world will take 30 years to adopt the necessary technologies to realize the risk of computerization in their labour market." (Santos et al. 2015, p. 13). The meaning of "realizing the risk of computerization in their labour market" is unclear, since it is hard to tell what it means to "realize a risk". Here we have interpreted the phrase as meaning that all jobs at risk will have been computerized, which makes most sense in the context of the report, but it is important to emphasize its lack of clarity. The figure of 30 years conveniently coincides with "The time frame in Frey and Osborne in these changes to happen" [sic] of "around 2 to 3 decades". Unfortunately, nowhere in Frey and Osborne (2017) is a time frame of two to three decades proposed; instead, at multiple points in the paper, they make reference to a time frame of one or two decades.²¹ Lacking a valid connection to Frey and Osborne's time frame, the figure of 30 years lacks any kind of backing.

One possible defence of the choice of this number may be linked to the authors' stated intention of providing "lower bound" estimates of automation risks. The choice of 30 years for full automation, by going beyond Frey and Osborne's time frame, enables their estimates to be relatively conservative. However, it strongly influences the size of the adjustment for the risk of automation in late adopters. For instance, a figure of 20 years until full automation, keeping the same adoption lags, yields an

²⁰ We are not told if this is equivalent to Frey and Osborne's high-risk category. Elsewhere in the paper, this figure for Bolivia is not corroborated.

²¹ While at different points in the text, Frey and Osborne (2017) vaguely refer to automation taking place over the "next decades" or "over an unspecified number of years", they explicitly refer to a time frame of one or two decades on pages 262, 265, and 268.

estimate of 32% of Bolivian jobs at risk of automation. Of course, the interpretation of the numbers is different, since the figure of 38% refers to a three-decade horizon, while the 32% would refer to a two-decade horizon. But none of this is mentioned in the WDR itself, where there is no reference to the time horizons over which jobs will be automated. Therefore, the arbitrariness in the numerical assumptions have a significant effect on the numbers obtained.

It could be argued that the exact number of jobs at risk does not matter as much as the ordering of countries in terms of risk of automation. The same argument could be extended to justify the use of Comin and Hobijn's (2010) numbers for adoption lags in an incompatible context: their numbers refer to the time between the invention of a technology and its first adoption in a given country, which is a completely different concept to the way they are used in the paper, where it is interpreted as the additional time required – relative to the early adopters – until a given occupation is fully automated.²² In fact, Santos et al. (2015) explicitly mention that the fact that they do not have measures of technology diffusion “should not matter because (they) are more interested in the *relative position* between countries” [my emphasis] (Santos et al. 2015, p. 12). Still, one cannot take for granted that the adoption lag (the “extensive margin” in Comin and Hobijn's terminology) correlates with technology diffusion (the “intensive margin”) in view of evidence suggesting the declining significance of the former and the rising significance of the latter over time (Comin and Mestieri 2018). Moreover, the persistence of large productivity differences across sectors in the same country (Caselli 2005; Herrendorf, Rogerson and Valentinyi 2014; Gollin, Lagakos and Waugh 2014), and among firms in the same sector (Hsieh and Klenow 2009; Restuccia and Rogerson 2017), suggest that, even if a technology is adopted, a different set of drivers might determine its diffusion rate. But irrespective of the accuracy of their assumptions on technology dynamics, the biggest limitation of Santos et al.'s (2015) automation forecasts is their reliance on Frey and Osborne's results, with all the attendant weaknesses.

The WDR (2016) also contains a second numerical indicator meant to quantify the impact of new automation technologies on the labour market. Using data from a variety of labour force surveys, they construct an “ICT intensity index” for a variety of occupations in developed and developing countries. They use this index, together with other data from the labour force surveys, to illustrate trends in world labour markets, focusing on jobs' changing skills requirements. These trends are very much in line with the predictions of the task model of the labour market and highlight the growing importance of nonroutine cognitive, socioemotional, and ICT skills. The discussion also touches on evidence suggesting that these changes in labour markets are driven by technological change. In Figure 2.25 (WDR 2016, p. 131), they use a scatter plot to show the relationship between occupations' automation probabilities and the ICT intensity index. This is meant to show how these two factors will interact to shape the labour market: presumably, tasks with a higher automation probability are likely to be automated, irrespective of their ICT intensity. Meanwhile, among the jobs with a lower probability of automation, those with higher ICT intensity, such as managers, will

²² See Comin and Hobijn (2010, p. 2042). Moreover, as Comin and Hobijn's results show, the adoption lag has diminished over time, and for the latest technologies analysed in their paper, the internet, the 10th, 50th and 90th percentile of adoption have reached five, eight and 11 years, respectively. The use of these numbers would make Santos et al.'s (2015) adoption adjustments insubstantial.

undergo changes in skills requirements, while other jobs low in ICT intensity, such as hairdressers, will go largely unaltered. It is also mentioned that the effects of new technologies on earnings will depend on jobs' complementarity with technology, the price and demand elasticities of the goods and services produced by these workers, and the relative supply and demand for skills.

Their analysis culminates in the construction of a "labour market disruption index" for each country. The purpose of this index is to capture labour markets' vulnerability to technological change due to changing skills requirements, which presumably changes the *nature* of existing jobs, and automation, which changes the *set* of existing jobs. The index is constructed through an equally weighted combination of the (adjusted) automation probabilities and the ICT intensity index, though the latter is taken from countries "at the next level of development, to be more forward-looking" (WDR 2016, p. 132). The indicator for the "level of development" is unclear, but it probably is some measure of per capita income. Among the countries in their sample, very few of which are low-income, the expected labour market disruption turns out to be highest for countries like Luxembourg, Italy and Austria, and lowest for Georgia, Albania and, above all, Ethiopia. Figure 2.26 in the report then charts the relationship between this index and quality-adjusted years of schooling, which indicates countries' capacity to deal with these disruptions.

Compared to the reports discussed in the previous sections, the WDR's (2016) main strength is that the factors mediating the effect of technological change on labour markets are more directly addressed in the estimation. The key elements here are the adjustment for adoption time lags and the use of the ICT intensity and expected labour market disruption indices to emphasise that much of the impact of new technologies will occur through changes in jobs' skills requirements. However, as discussed above, besides the limitations inherent to the uncritical adoption of Frey and Osborne's results, the assumptions made to map countries' relative exposure to technological change are somewhat coarse. In addition, some of the factors listed earlier in this section are not addressed – for example, there is no attempt to predict how ethical, legal and societal factors might affect technology adoption, or what occupations will witness an increase in employment due to income effects from technological change (ie. the growth in income brought forth by increased productivity due to technological progress). Although these considerations are brought up, they are not incorporated into the technical analysis. On reading the report, one gets the impression that this is due to the peripheral role of this analysis, which serves merely to illustrate the themes discussed in the main body of the text.

The two MGI reports (MGI 2017a; 2017b) resemble the WDR (2016) in going beyond pure estimates of automatability and attempting a more detailed analysis of the labour market implications of new technologies. Yet, they differ dramatically in the scope of their quantitative analysis and in its centrality to the discussion in the main body of the text. MGI (2017a) deals with "the automation potential of the global economy, the factors that will determine the pace and extent of workplace adoption, and the economic impact associated with its potential." (MGI 2017a, p. vi). Meanwhile, the MGI (2017b) centres on "potential labour market disruptions from automation and some potential sources of new labour demand that will create jobs." (MGI 2017b, p. ii). By dealing with technology adoption and labour market implications separately, the reports are able to treat each question in greater depth, and they jointly offer the most detailed quantitative analysis of the labour market impacts of automation.

Like other studies discussed here, MGI (2017a) builds on Frey and Osborne to develop a methodology for estimating jobs' susceptibility to automation, using World Bank data for 45 developed and developing countries, and US Bureau of Labor Statistics data. The novelty in their approach consists of breaking down 800 occupations into more than 2,000 activities such as "greet customers", "clean and maintain work areas" or "process sales and transactions." These activities are in turn broken down into 18 capabilities, grouped under "sensory perception", "cognitive capabilities", "natural language processing", "social and emotional capabilities", and "physical capabilities". Each activity can be rated on a four-point scale with respect to each capability. Activities are scored in relation to capabilities through a machine learning algorithm by "matching keywords from the capability to the activity title" (MGI 2017a, p. 122), as well as some manual adjustments. Few details are given as to how exactly this matching is performed.

With the detailed breakdown of occupation characteristics in hand, the following step in the estimation consists of modelling automation adoption timelines for each capability. This is modelled as depending on the timelines for technical feasibility, solution development (that is, the creation of productive inputs that embody a given technology), economic feasibility (that is, relative costs), and adoption and deployment. The assumptions made to project the evolution of each of these areas are too numerous to be outlined in detail; it suffices to mention that they combine survey findings, extrapolation of current trends, other publications, historical examples, and a variety of forecasting models. The final inputs into their model are a series of economic projections on labour force evolution and GDP per capita, as well as data on time spent on each activity and hourly wages per activity. This allows them to calculate how many of the activities in currently existing jobs can be automated under two scenarios: one where GDP per capita is constant at today's level, and one where it grows according to their GDP per capita projections. In both scenarios, they triangulate GDP per capita projections with productivity data, as of 2014, and the projections for population growth to calculate the number of full-time equivalent (FTE) jobs consistent with their numbers. All these moving parts are combined to estimate the main figures reported in the study, including the impact of automation by job title, by activity, and by industry. Unlike in other studies, automatability is reported in terms of the share of activities, time, and wages in each occupation, thus providing a more granular view of new technologies' labour market impacts.

MGI (2017b) builds on the earlier study, focusing more closely on the labour market impacts of technological change. The evolution of labour markets is analysed with more nuance. Besides the substitutability of jobs with machines, they condition future labour market conditions on the choices of governments, business leaders, and individuals with respect to investment in infrastructure and buildings; investment in renewable energy, energy efficiency, and climate adaptation; and the "marketization" of previously unpaid domestic work. Moreover, they also consider three drivers of new job creation, including: rising incomes and consumption, especially in emerging economies; ageing populations; and the creation of jobs through the development and deployment of technology. These considerations are formalised in their quantitative estimates, where they assume that the number of jobs that can be automated in a given sector or occupation is proportional to the fraction of constituent activities that are automatable. A number of other assumptions are made with regards to the labour demand drivers and macroeconomic projections. Perhaps the most interesting assumptions concern the mapping of new labour demand and automation onto the future demand for occupations. On the one hand, labour substitution by machines is accounted

for by subtracting automated jobs from the projected occupational composition in 2030. On the other hand, the effects of the drivers of new labour demand are modelled on a driver-by-driver basis. Since most drivers create jobs in a specific sector, their projections include more jobs for the sectors affected by these labour demand drivers. These are complemented by projected increases of "indirect jobs", created due to linkages with other sectors, and calculated using input-output tables based on the World Input-Output Database.

The above description of the methodology employed in the MGI reports shows the number of moving parts that make up their projections of the effects of automation on labour markets. While other studies, building on Frey and Osborne, simply calculate the proportion of current jobs susceptible to being automated and make some qualifications for the factors that may slow down or speed up the process of automation, the MGI reports seek to integrate every possible consideration into their analyses and project different scenarios for the future. Of the studies reviewed here, these are the only ones to go beyond assessing the automatability of existing jobs and to make a serious effort at tackling the counterfactual question of how many of the jobs that would otherwise have been created will be replaced by machines. This is a particularly pertinent consideration for developing countries, where the pace of structural change is likely to be higher than in the developed countries which most studies focus on.

While it presents an ambitious and comprehensive method for quantifying the labour market implications of new technologies, MGI's approach evidences the complexity of performing such kinds of estimation and the need to often settle for relatively simplistic or arbitrary assumptions. For instance, it does not make much sense to assume that, in the absence of technological change, the occupational composition of the economy will remain the same, particularly in dynamic emerging economies. Similarly, it is unrealistic to disregard the future impact of new industries and occupations, changing business models in existing industries, the rise of the gig economy, or the changing task content of jobs. The authors of the report also make additional caveats on factors that may alter their predictions (see Box E2, MGI 2017b p. 21). However, figuring out exactly how to numerically forecast these undoubtedly important developments is probably beyond the capacity of even the most gifted analyst, and the reports do try to work within the bounds of realistic projections.

But despite constituting laudable efforts towards integrating the range of considerations that might influence the economic effects of new technologies, the MGI reports retain some of the opacity problems afflicting other studies. For instance, there is no discussion of how they assign capabilities to activities, or any detailed disclosure of the models and methods used to operationalise their assumptions, including their machine learning algorithm. Admittedly, the report is not meant for an academic audience, and perhaps should not be held to the same transparency standards as other studies discussed here. Still, as a result, the degree of confidence we can have in their predictions is, to a large extent, compromised. Given that, despite representing the most comprehensive, ambitious, and resource-intensive effort at quantification among all studies on automation, the MGI reports still rest on some relatively arbitrary assumptions, whose operationalisation is left largely unexplained, one may wonder about the prospects of this strategy for evaluating the labour market impacts of automation: it might be the best we can do, but is even this informative enough to influence policy decisions? The answer to this question can have profound implications for considerations on the best avenues to expend future research efforts.

Discussion: technology and public policy in developing countries

So far, I have focused on reviewing the empirical methodologies of studies attempting to quantify the labour market impacts of new automation technologies. But before proceeding to compare the merits of the different approaches, it is useful to consider the policies advocated by each study. Table 3 makes this comparison. Despite the methodological differences between these studies, and the level of detail with which they discuss policy issues, we can see that there is a striking uniformity in policy prescriptions, and that these fall into three areas: skills development, including training/retraining and education; social protection; and measures related to economic growth. Of course, we must examine the details, and even if there is a broad agreement on the policy priorities for the automation age, few of the studies address the specifics of policy design.²³

Table 3: Policy implications

Report	Policy Implications
Frey and Osborne (2017)	· Workers will need to acquire creative and social intelligence, which are non-susceptible to automation
Arntz et al. (2016)	· Providing retraining to low-skilled workers
Nedelkoska and Quintini (2018)	· Providing internship schemes for young people · Providing training and social protection to those most vulnerable to automation
PwC (2018)	· Boosting education and skills levels · Supporting job creation through government investment · Enhancing social safety nets
Chang and Huynh (2016)	· Cultivating skills and preparing the workforce for new ways of working through corporate education and government initiatives · Governments need to extensively revamp their economic structures, carefully think about automation trajectories, and take appropriate steps to diversify the country's economy
WDR (2016) ²⁴	· Supporting entrepreneurship and innovation to expand businesses and job opportunities · Ensure that education and training systems, labour regulations, and social protection institutions support workers in seizing opportunities generated by the internet
MGI (2017a)	· Policies to encourage investment · Market incentives to encourage innovation · Rethinking education and training, income support, safety nets, and transition support for those affected by automation
MGI (2017b)	· Ensuring robust demand growth and economic dynamism · Scaling mid-career job training to make lifelong learning a reality · Enhancing labour market dynamism and enabling worker redeployment · Modernising educational systems · Rethinking and strengthening transitions and income support for workers caught in the cross-currents of automation · Modernising data collection on the labour market · All these areas, taken together, "may require a Marshall Plan-like initiative".

²³ The exceptions are the MGI reports and the WDR 2016, though the latter focuses mostly on policies related to the digital economy.

²⁴ The WDR (2016) also discusses policy implications in other areas, but here I focus on those relevant to the analysis of automation conducted in Chapter 2.

This uniformity also extends to the underlying models and sources used to frame the studies. Frey and Osborne's work is the most proximate source of inspiration for most of this literature, and one cannot emphasise enough the influence that their methodological choices have had on subsequent publications.²⁵ But they, in turn, draw inspiration from the task model of the labour market pioneered by David Autor and co-authors, as they think of the drivers of automation in terms of occupations' task contents and the comparative advantages of workers versus machines. One can easily notice the resemblance between Frey and Osborne's list of "bottlenecks to computerisation" and Autor et al's (2003) list of non-routine tasks (see Footnote 8, above), for example.

It could be argued that the early contributions to this line of research created an ambivalent legacy for subsequent work on the topic. By putting the risks posed by machine learning technologies in the spotlight and providing a common analytical basis to a field of inquiry where it was previously lacking, these papers have helped direct research efforts towards matters of crucial importance. At the same time, since they provided – at least until recently – the *only* framework used for these types of analysis, they reproduced its weaknesses and may have retarded the emergence of alternative ways of thinking about the effects of new technologies on jobs in developing countries. As discussed above, most methodologies used to estimate exposure to the effects of new automation and digital technologies suffer from a combination of reliance on subjective indicators, lack of transparency, non-standard empirical strategies, and interpretive incoherence. Meanwhile, due to the complexity of the phenomenon being modelled, and despite the greater resources at their disposal, the reports by the MGI rely on so many assumptions that their forecasts are more than a little uncertain. At the moment, it does not seem like there remains much scope for further research of this kind.

Faced with this scenario, one could give up hope of forming quantitative forecasts and argue that any attempts are useless. However, we cannot ignore two purposes that may be served by this type of quantitative analysis. In the first place, it allows us to compare countries' relative exposures to automation trends. Even if we do not believe that exactly 26% of jobs in Japan will be automated by 2030, we do know that, because of its relatively large manufacturing sector, the high wages paid, and advanced technological development, its economy is more likely to undergo automation than an unindustrialised, low-wage and technologically lagging economy (though with many pockets of high-tech enterprise) like India. We might also like to have a rough idea of how different sectors will be affected, and who might be the winners and losers under different scenarios, even if we do not take the numbers too literally. For this purpose, even if the assumptions are unrealistic – which, given the sheer difficulty of making some economic forecasts, they are bound to be – as long as we agree that the countries are more or less ranked correctly in the different dimensions of the model, then the model can be informative when making comparisons.

A second purpose for which predictive methods can be useful is in helping governments correctly identify the citizens who are most vulnerable to automation, so as to better target social policies. In this respect, despite the proliferation of methodologies predicting the distributional impacts of new automation technologies, none of the methods proposed so far have been externally validated (that

²⁵ Even recent, more academically-oriented attempts to quantify the scale of the impact of machine learning on developed economies, such as Brynjolfsson and Mitchell (2017), clearly draw on many aspects of the methodology first developed by Frey and Osborne.

is, had their accuracy tested). Lacking real data on workers being replaced by machines, an entire collection of literature has been built on the subjective predictions of the machine learning experts gathered by Frey and Osborne. As discussed extensively in this paper, this is not the strongest of foundations, although it is an original effort to make use of existing data. In the absence of suitable training and test data, we cannot know the accuracy of machine learning and other predictive methods, including linear regression. The compilation of such data can be a suitable area for public policy, as also proposed by MGI (2017b) and Mitchell and Brynjolfsson (2017), and it would greatly facilitate the task of assessing and redressing the distributional risks posed by new automation technologies.

Unfortunately, the two uses of quantitative predictive methods described above are unlikely to be equally suitable for developing countries as they are for developed ones. This is not to say that they are not useful, since we would certainly like to assess these countries' relative exposure to automation and to correctly target social policy beneficiaries. The problem is that, more often than not, developing countries have poorly developed systems for social policy, and the state capacity necessary to deal with the challenging implications of new technologies is lacking. Moreover, although some upper-middle-income economies in Latin America seem far removed from the image of dynamic structural change usually associated with 'developing' or 'emerging' economies, many economies in Asia and Africa do fit this picture. In these dynamic economies, it may be harder to form predictions and comparisons of the effects of automation, since growth rates tend to be quite volatile; rapid structural transformation makes it harder to form counterfactuals for employment in the absence of technological change; and poor data, which often covers only the formal sector, may hinder efforts at empirical analysis.

Due to the difficulties of performing the same quantitative analysis usually done for developed countries (though seldom in a reliable way), and to the limited reach of social policy, studies in the style of Frey and Osborne or the MGI are unlikely to be as useful in developing countries as in developed ones. So what is to be done? When discussing the information shown in Table 3, I noted that most policy prescriptions for dealing with technological change fall under social protection, skills development and growth promotion. I have already mentioned that social protection is unlikely to fill the same role in developing countries as in developed ones. In skills development, although there are some important policy questions to be understood, and which are discussed at length in the draft WDR 2019, it is unlikely that quantitative analysis of the sort discussed here can provide any policy insights for developing countries.

This leaves us with the third area, economic growth. Although this is a relatively contentious point, one could plausibly argue that growth is likely to be the main policy area where developing country governments will confront the issue of technological change. To be sure, economic growth is never far from the concerns of any government in the world, and government policies certainly play a role in ensuring that firms in developed countries properly harness the productivity potential of automation and digital technologies. But the fact that firms in developing countries lack the productive capabilities of their developed country counterparts, and that governments usually have an outsized role in shaping their structural transformation, means that they will have to worry much more about how to adapt their development strategies to the new technological environment. This will involve thinking about how to reach and maintain high growth *rates* and, crucially, given the adverse distributional impacts of new technologies postulated by most analysts, thinking about the *type* of growth.

Thinking about development strategies in the new technological environment is likely to be a more arduous task for developing economies than developed ones because the effects of automation and digital technologies are still ill-understood. By contrast, the task model of the labour market appears to have taken hold in discussions on the effects of automation on developed countries, as attested by the repetitiveness in the analysis and prescriptions of texts on the subject. If we accept that the task model is a good description of the forces linking technological change and labour market developments – although it obviously is just one model and does not explain every aspect of the labour market – then it seems like there already exists a good understanding of the main challenges faced by developed countries, and of the policies that can be used to face these challenges. In fact, some of the current work in this area of research is in thinking about the policy dimension, particularly because it is bound to be intertwined with political developments (Levy 2018; Frey, Berger and Chen 2018). Meanwhile, in developing countries, we are far from the stage where we can worry about the political economy of policy implementation, since we still lack an adequate understanding of the economic forces linking technological changes and developing countries' growth prospects, including both potential sources of growth and their distributive implications.

2. The economics of technology in developing countries

In the previous section, we discussed how understanding the effects of new automation and digital technologies on developing economies calls for an exploration of the underlying economic forces driving them. In the canonical task model of the labour market, the economic force driving the distributional outcomes is the falling cost of automation, in interaction with the distribution of occupations and skills (Autor et al. 2003; Autor and Dorn 2013). Later developments of the task model have also come to incorporate variables such as trade opportunities, the introduction of new tasks, the rate of technological progress, types of automation, the long-run rental rate of capital, savings decisions, the elasticity of substitution between labour and capital, and demographics (Acemoglu and Autor 2011; Acemoglu and Restrepo 2018a; Acemoglu and Restrepo 2018b; Acemoglu and Restrepo 2018c; Caselli and Manning 2018; Susskind 2017). These interact with the falling cost of automation to bring about economywide outcomes.

Based on this example, we can ask a similar question for developing countries: what factors matter for determining how declines in automation costs bring about economic outcomes in middle- and low-income countries? While it is true that the factors identified for developed countries are also likely to play a role in developing countries, the relative importance of factors will probably be different. For instance, technological progress is likely to be exogenous in developing countries, barriers to trade are higher, and some technologies may be harder to introduce due to the absence of complementary physical and human capital. These particularities mean that some models will be more relevant than others for understanding the economic impact of new technologies. Thus, rather than debating the validity of different models in the abstract, the question boils down to figuring out how to choose the most appropriate model.²⁶ With this in mind, we can interpret the two reports discussed in this section as offering alternative strategies for understanding the economic impacts of automation on developing countries.

The UNCTAD Trade and Development Report 2017

The chapter on Robots, Industrialization and Inclusive Growth in UNCTAD (2017,p.39) discusses:

...whether the use of industrial robots can be expected to radically change the types of jobs that will be available in the future, how, where and by whom they will be done, and what impact this would have on possibilities for inclusive growth, in terms of declining income inequality both between and within countries.

Note that the report focuses exclusively on industrial robots, and not on the entire range of new automation technologies. Still, the analysis remains pertinent, since industrial robots are one of the most important automation technologies. Although the report is not set up in terms of isolating

²⁶ See Rodrik (2015) for the thesis that good economic practice requires being able to choose among competing models for the model that best fits the context we are interested in.

the economic mechanisms at play, it makes a number of relevant considerations. In this section, I attempt to reconstruct its arguments so as to single out the factors seen as mediating the effects of technological change on developing economies.

The central economic mechanism posited by the report follows directly from the task model of the labour market, applied at an international level. The replacement of workers performing routine tasks by robots has growth and distributional effects in the country where it occurs, as it can spill over internationally through trade. A falling cost of automation increases the comparative advantage of countries endowed with some complementary factor of production, notably skilled labour. This can enable that country to steal the international or even domestic market share of a third country. Following this logic, the report repeatedly mentions the danger that manufacturing becomes increasingly concentrated in developed countries due to the falling cost of automation, raising the barriers for developing countries to enter more skill-intensive manufacturing sectors, where automation is more economically feasible for the moment.

This process can have implications for growth and inequality in developing countries. On the growth side, it can limit developing countries' ability to build manufacturing sectors, which have traditionally been the most reliable 'engines' of economic growth (Rodrik 2013; 2016). Alternatively, by decreasing the labour intensity of manufacturing and other dynamic sectors, it may create a choice between entering these dynamic sectors or providing large-scale employment for the population (UNCTAD 2017, p. 58). The reduction in the share of workers in some automatable sectors, either through within-country automation or spill-overs from the automation of other countries, can also have distributional consequences. The report mentions that these distributional consequences depend on "a country's stage in structural transformation, its position in the international division of labour, demographic developments, and its economic and social policies" (UNCTAD 2017, p. 39), but shies away from explicitly delineating the exact relationship between these factors and income distribution.

The lack of clearer, or more formal, explanations of the economic forces driving the processes described reduces the clarity of the UNCTAD report. The report does make some original points, particularly with regards to the need to think of the global distributional implications of automation and the many ways they can affect developing countries, but at times, a formal illustration of the mechanisms at work would have helped in understanding under what conditions different scenarios will be manifested. Although their reasoning derives mostly from the task model, it would be interesting to consider what other factors come into play once we think about automation at the international level. What is the role of trade barriers? What other factors besides labour compensation determine technology adoption? And how does automation in one country affect income distribution in another country?

Formal modelling would be particularly important for some of the more original points made by the report. For instance, it is mentioned in the report that:

- Labour compensation cannot be continuously reduced over a prolonged period because it would lead to a fall in consumption and thus the incentive to invest.

- In the absence of expansionary global and domestic macroeconomic conditions, investment growth will be stifled, shifting employment away from technologically dynamic sectors.
- Increased productivity due to robots can be beneficial to poor countries by leading to a global decline in the prices of skill-intensive manufactured goods relative to labour-intensive manufactures and primary products.
- The distributional consequences of robot deployment depend not only on changes in wages and employment, but also on the distribution of ownership of the robots and the distribution of intellectual property embodied in them. Robots tend to increase the capital share of income.
- Reshoring of production from developing to developed countries may not occur where foreign direct investment (FDI) targets local demand, or where a dense supplier network has been built.

It would be interesting to know the parameters that determine the conditions where automation leads to income stagnation, how macroeconomic policies determine the pace and extent of automation, or what international price changes automation can lead to. More generally, formal modelling can clarify the logical structure behind ideas and allow us to consider their validity. In the absence of a formal model, particularly for the high-level mechanisms described in the UNCTAD report, it is very hard to know what to expect in terms of the relative magnitudes of the different economic forces. Of course, it might not be the role of a policy-oriented publication to build such formal models, but it certainly would be a promising area of work for academic economists.²⁷

Another area where more rigorous work would have been helpful is in the evaluation of the evidence for the hypotheses put forward. All the data analysis in the text is purely correlational. There certainly is value in some of these correlational analyses, such as in showing different sectors' routine task intensity, which is an innovation in comparison to other work on automation. However, in other parts of the text, some stronger claims are made based solely on this correlational data. For example, we are made to believe that deviations from average labour compensation in manufacturing are more important than routine-task intensity in determining robotisation, but it is also mentioned that differences in technical feasibility between sectors may be driving this relationship. In the absence of a full set of controls for factors that may influence the observed relationships, we cannot know exactly how each factor matters, which factors are most important, and if there is a truly statistically significant relationship with robot deployment. Similarly, they make several claims on the relationship between robot deployment and real wage growth based purely on correlational data (pp 52-54). Lacking a more systematic inquiry into the relationships between different economic variables, it is hard to believe their claims with much conviction.

²⁷ As it is already in the work of Daron Acemoglu and Pascual Restrepo (2018a; 2018b; 2018c) who have been building more complex models of automation taking into account a greater variety of economic factors.

The future of manufacturing-led development strategies

The UNCTAD (2017) report, despite dedicating only one chapter to a discussion of the effects of automation in developing countries, is quite encompassing, and mentions many ways in which these effects may come about. One of them is by changing the desirability of manufacturing-led development strategies. This is the theme of a World Bank report by Mary Hallward-Driemeier and Gaurav Nayyar (2017) entitled *Trouble in the Making? The Future of Manufacturing-Led Development*. The report is an in-depth discussion of the idea that the advent of new technologies might lead to a change in the “pro-development” characteristics of manufacturing. These are the characteristics of the manufacturing sector that have made it traditionally considered special, by virtue of the added benefits associate with entering that sector.²⁸ Reviewing the relevant economics literature, they identify the following pro-development characteristics of manufacturing:

1. Scope to employ unskilled workers
2. Sector's share of labour in the overall economy (that is, labour intensity)
3. (Higher) labour productivity
4. Tradedness
5. Scope for innovation and diffusion.

Importantly, they mention that these pro-development characteristics will vary across countries and over time, due to differences in the set of tasks that may be used to produce the same good. Usually poorer countries will use more labour-intensive methods than richer ones.

An additional source of variation in production methods is the part of the value chain in which firms are active (Hallward-Driemeier and Nayyar 2017, p. 24).²⁹ Based on these characteristics, the authors build a typology of manufacturing sectors according to their intensity along the five dimensions.

The central part of the report then discusses the implications of new technologies for developing countries. They argue that historically, each industrial revolution (that is, Industry 1.0, Industry 2.0, and so on) and associated technological innovations have changed countries' comparative advantages and, consequently, specialisation patterns and the international division of labour (see their discussion on pp. 40-44). To analyse the impact of Industry 4.0, they focus on: “process technologies” such as robotics (particularly artificial intelligence-enabled); digitalization and internet-based systems integration, including sensor-using “smart factories” (that is, the Internet of Things); and 3D printing. They cite some of the changes that may be brought about by the advent of these new technologies, though noting that we are still at the early stages of their deployment:

- New ICTs can reduce trade and co-ordination costs, facilitating further global fragmentation of production.

²⁸ See Dercon, Lippolis and Peel (2018) for a study on how sectoral specialisation matters for countries' development.

²⁹ One could arguably consider this a classification problem, since different production processes are being classified in the same way. This is one of the pitfalls of goods-centric classification methods. For an example of how industry classifications may mask differences in the nature of goods and of their production methods, see Schott (2003). For a more extensive discussion of how the goods produced by countries vary according to their levels of development, see Dercon et al. (2018).

- The internet of things may make it more efficient to rebundle activities in “smart” factories, as keeping all parts of the value chain at close distance can dramatically increase efficiency. This would reduce the importance of lower labour costs in determining comparative advantages.
- Smart factories increase the importance of complementary, high-skilled services, which are more likely to be present in more advanced economies, strengthening their comparative advantage in manufacturing.
- 3-D printing can lead either to more or less fragmentation in manufacturing production, depending on the presence of complementary services and the protection of intellectual property rights.

So what are we to make of these changes in terms of the pro-development characteristics of manufacturing? Technological advances that permit the automation of new tasks, together with the rising importance of complementary services, may reduce manufacturing's labour intensity and its scope to employ unskilled workers. The tradedness of manufacturing goods will not be affected, but the introduction of smart factories and 3D printing might shorten global value chains. Thus, although the manufacturing sector will continue to maintain high levels of productivity, this will come at the cost of employment. The benefits of manufacturing production may increasingly be enjoyed only by developed countries. Finally, due to the disjuncture between old and new production processes, the potential for spill-overs and dynamic gains in manufacturing may be reduced, further “raising the bar” for developing countries.

Hallward-Driemeier and Nayyar (2017) argue that manufacturing sectors vary in their exposure to these changes according to the extent of their automation, export concentration, services intensity, and international tradedness. Accordingly, automation poses more of a challenge for “medium-skill innovators” (for example, electrical equipment, machinery and equipment not elsewhere classified) and “high-skill global investors” (such as computing, electronics and optical equipment, pharmaceutical products), due to their high scores in each of these dimensions. At the opposite end of the spectrum, textiles and apparel will continue posing less of a challenge, due to their low propensity for automation and the lesser importance of complementary services. Other sectors are at intermediate levels of exposure. Meanwhile, with the growing importance of services in manufacturing – that is, the “servicification” of manufacturing – and the possibilities opened up by technological change, services are increasingly acquiring some of the pro-development characteristics of manufacturing, although they might involve a starker trade-off between the provision of employment and dynamic productivity gains.

The report calls for a number of policy measures to help attenuate the negative impacts of automation, distributed in three areas (the “three Cs”): improving competitiveness (understood as the quality of the business environment); firm and worker capabilities; and ‘connectedness’, defined as the quality of logistics and the ease of conducting international trade. New technologies will be more demanding in these areas, with the relative weights of each area depending on the same sectoral characteristics determining their vulnerability to automation. They also claim that the new scenario poses a new kind of trade-off between vertical and horizontal interventions: given the speed

at which technology is advancing and the uncertainty surrounding the global economic landscape, sector-specific policies are now riskier than before. At the same time, the increased requirements for production in terms of the three Cs increase the difficulty of creating an economic environment that is competitive in every respect, thus stressing the need for prioritisation. The authors of the report suggest that the right mix between the two kinds of intervention will be context specific. They conclude the report with questions on the continued relevance of sector-specific approaches, and on whether the target of industrial policy should switch to job provision, rather than the dynamic productivity gains that have historically been central to thinking on industrial policy.

This summary of Hallward-Driemeier and Nayyar (2017) shows that they touch on many of the same themes as UNCTAD (2017), but discuss the role of manufacturing in much greater depth. Other themes are completely new, such as the growing importance of services, while distributional issues are comparatively under-emphasised, although they are by no means ignored. The report cites the relevant literature extensively and is well grounded in economic theory and data when making its points. These strengths notwithstanding, it is unclear where some of the central analytical distinctions come from.

This is the case in the assertion that the determinants of the effects of new technologies on developing countries are their degree of automation, their export concentration, their services intensity, and the extent of their international tradedness. Although this certainly sounds plausible, it is also natural to expect these sectoral characteristics to change along with technological developments; perhaps a greater focus on the changing pro-development characteristics of different manufacturing sectors would have given the report more analytical coherence. Similarly, the proposed dimensions of reform – the three Cs – appear too general as categories and do not differ very much from standard World Bank policy prescriptions,³⁰ although the report does recognise their lack of novelty. In future work, it would be best to seek more detailed and nuanced policy prescriptions, although these might still be hard to discern, given that we are at the early stages of Industry 4.0.

From our perspective, we can observe that the arguments put forward in the report call on a variety of economic mechanisms to drive the connection between new automation technologies and their impacts on developing countries. In contrast to UNCTAD (2017), these are mostly at the micro level. Moreover, they do not draw as much from the task model of the labour market; while the task model is used to explain the economic impacts of technological progress, other reasons are cited, including technological complementarities, co-ordination costs, or the changing nature of dynamic productivity spill-overs in different economic sectors. The key feature shared by these factors is their reliance on multiple new technologies. For example, while the internet can facilitate the distribution of production along global value chains by lowering communication costs, by embedding the internet in other production technologies, smart factories could possibly increase the concentration of global manufacturing production. Similarly, 3D printing, combined with communication technologies, could make manufacturing production more accessible. However, if communication is inhibited – either due to restrictions in intellectual property rights or a given region's reduced connectivity – this could lead to a further concentration of production.

³⁰ See Neilson (2014) for an analysis of how the incorporation of a new approach in international development, Global Value Chains, did little to change standard policy prescriptions.

In the UNCTAD report (2017), these complementarities are not evident, since they deal only with the replacement of workers by industrial robots. Although this is an acceptable analytical stance given that the report only dedicates one chapter to the issue of automation, it evidences how new technologies need to be considered in conjunction, and how studying only one technology in isolation can be misleading. This same problem can be encountered in the other texts reviewed here, where the fixation with predicting the share of jobs that will be automated prevents them from adopting a more comprehensive perspective. Seen in this light, Hallward-Driemeier and Nayyar's (2017) approach of appraising the impacts of Industry 4.0 as a whole is commendable. However, it is ripe with analytical challenges, since the task model cannot provide the most useful angle for thinking about every aspect of the implications of new technologies. In the next section, I conclude with considerations on some possible directions for future research on the topic.

3. Conclusions

This paper has reviewed key publications in the literature on the effects of new automation technologies on developing countries, with a view to assessing the strengths and limitations of different analytical strategies. My starting point was an in-depth review of Frey and Osborne's (2017) seminal study, where I outlined the rationale of their approach and highlighted the key strengths and drawbacks in their methodology and interpretation. I then proceeded to review similar studies that attempt to quantify job losses due to technology in developing countries. These studies combine some of the weaknesses of Frey and Osborne's approach with others of their own making. Nevertheless, they do make some valid points concerning the complex relationship between technical feasibility and job automation, identifying key variables that mediate this effect. This quantitative stream of research culminates in the two reports by the MGI, which combine a number of models and assumptions to predict countries' exposure to the productivity and employment effects of automation technologies. I argue that, despite their laudable effort to seriously consider technological changes in all their complexity, and to recognise the uncertainty associated with their estimates, this uncertainty is just too great for the estimates to serve as an indication of the actual number of jobs in peril. Moreover, there are reasons to suspect these estimates to be less accurate for developing countries. Still, I note that they can have a useful comparative function and can help governments identify workers at risk.

In developing countries, automation is unlikely to consist of just a distributive issue, improved simply through social policy and improved education. However, automation might affect these countries' prospects of catching up with the income levels of rich countries. This dynamic and international dimension is neglected in work seeking to quantify exposure to automation, but its importance is forcefully argued in the UNCTAD (2017) report. It also brings in a macroeconomic flavour to the topic, which was conspicuously missing in previous work. This topic is also relatively neglected in Hallward-Driemeier and Nayyar (2017), which goes much deeper into the question of how technological innovations will affect growth prospects and growth policy in developing countries. The authors identify the changing "pro-development" characteristics of sectors, particularly manufacturing, as one of the main channels through which technological change will require rethinking development strategies.

Modelling the effects of new technologies

My reading of the literature suggests that, to gain a fuller understanding of the effects of new technologies on developing country job markets, it is necessary to overcome the tendency to worry only about potential job displacement by machines.³¹ Although this is one of the ways in which technology's effects may be felt, developing countries' typically large adoption lags, their low levels of technological intensity (Comin and Mestieri 2018), large levels of productivity dispersion (Hsieh and Klenow 2009; Herrendorf et al. 2014; Gollin et al. 2014), and other structural characteristics (Dao

³¹ Caselli and Manning (2018) note that most studies on the topic ignore the general equilibrium effects of automation.

et al. 2017) are likely to limit the direct impact of technologically induced labour displacement. As a result, researchers would be well advised to open their eyes to the full range of economic mechanisms mediating the impact of new technologies on developing countries, and in trying to assess their respective quantitative significance. For instance, recent economic models have brought to light the possible effects of new technologies on market structures, changes in internal firm organisation, and even the production of new ideas (Autor et al. 2017; Aghion, Jones and Jones 2017; Brynjolfsson, Mitchell and Rock 2018). One naturally wonders to what extent these results apply to developing countries, and what variables condition their applicability. This could form part of a larger research programme seeking to understand how the particularities of developing countries may affect the results obtained in the existing literature. Unfortunately, there can be no quick fix here, and rigorous academic research will be required to advance our understanding of the subject. Without this research, assessment and policy studies will continue to be largely speculative.

Another primary area that needs to be tackled is the international dimension of automation. Although the effects of technological change on concentrating global production are often discussed, little rigorous academic work has been done on the matter. Rodrik (2016) and Dao et al. (2017) are among the few academic publications to have seriously grappled with the macro-level effects of technological progress on developing economies. There is a lot of scope for research on changing terms of trade and the changing geography of production caused by automation, both theoretically – by adapting existing trade models – and empirically. New technologies must also be understood from different perspectives. For example, UNCTAD (2017) brings up the relationship between automation and macroeconomic policies. The distributional effects of falling automation costs have also been studied in a prominent macroeconomic study by Karabarbounis and Neiman (2014). Moreover, following from the work of Hallward-Driemeier and Nayyar (2017), it would be interesting to understand how the “pro-development” characteristics of manufacturing and other sectors change with technological progress. Building on the literature that contributed to the identification of these characteristics,³² it would be possible to adapt existing models to the new changes in production processes and their implications for the scope for innovation and diffusion.

Thinking holistically about technological paradigms

The future research agenda should also include recognising the diversity of new technologies and thinking holistically about their effects across the economic, political and social domains. Labels such as “new automation” often belie diverse technologies that in practice interact in myriad ways to determine economic outcomes. The distinction between production and communication technologies seems particularly relevant, even if “Industry 4.0” often leads to a blurring of the lines between the two in technologies such as the Internet of Things. On the one hand, the growing automation of production appears to pose serious competitiveness challenges for many poor countries, leading to greater difficulties in entering global value chains, onshoring of production, and the invasion of their markets by goods produced in more developed economies, although Hallward-Driemeier and Nayyar (2017) do mention the possibility that 3D printing might foster regional goods markets.

³² See Hallward-Driemeier and Nayyar (2017, pp. 9-17).

Many commentators are optimistic about the possibilities engendered by the diffusion of the internet and the rise of new ways of performing economic transactions, such as mobile banking, online platforms, and other new business models enabled by improvements in connectivity. However, it can be misleading to see these trends in separation when the interaction between different kinds of new technology can give rise to entirely new ways of performing work tasks – not seeking merely to emulate human practices, but harnessing machine capabilities to perform them in fundamentally different ways (Susskind and Susskind 2015). This is an additional reason why focusing the debate simply on job replacement can be seriously misleading, as it distracts us from the world of possibilities opened up by the combinations of new technologies.

Besides changes in production methods and their effects on labour shares, employment and growth, a case can be made for also paying attention to the effects of new technologies in areas beyond the usual scope of economic analysis, such as the functioning of the state or patterns of political competition. As argued by Susskind (2018), the advent of these new technologies might even give rise to new ways of *thinking* about politics. Susskind gives various examples of how technology can facilitate the enforcement of state policies and how, in partnership with tech firms, the state might be able to increase the projection of its power through greater scrutiny of citizens' lives and control of their perceptions (Susskind 2018, pp. 122-160). The treatment given in the book suggests we should be concerned with these developments because they might result in an over-extension of state power over our personal lives. However, things might be different from the point of view of developing countries, where state weakness has often been blamed for disappointing developmental outcomes. Although developing world states should obviously not be treated as purely benevolent actors, it is true that digital technologies harbour a great potential for application in areas such as contract enforcement, tax collection, or auditing of the use of public funds, with positive consequences on economic activity. These innovations, in conjunction with changes in production technologies, market structure, firm organisation, and international trade, can lead to large-scale alterations in power dynamics, with repercussions on the ways economies function.³³ It thus become essential to study the emerging technological paradigm holistically, as focusing on only one aspect, such as labour replacement, can blind us to important emerging patterns.

Finally, the feasibility of a research agenda is predicated on the availability of adequate data. While data from the International Federation of Robotics has recently come to researchers' attention (for example, UNCTAD 2017; Graetz and Michaels 2018), it only covers "automatically controlled, reprogrammable, and multipurpose" machines, and therefore: "A robot dedicated to a single purpose – for example, machinery that automatically assembles circuit boards – is not counted, even though it may substitute for human workers" (Levy 2018, p. 409). Our knowledge on the diffusion of other technologies in Industry 4.0 is still very limited, and it is clear that, to better understand the drivers of technology adoption and the impacts of technology on developing economies, more micro-level data is needed. This could be obtained through entirely new data collection projects, or by adding questions to widely administered surveys such as the World Bank's Enterprise Survey. Either way, increasing the amount and variety of data collected in developing (and developed) countries is an invaluable, if somewhat obvious, requirement for helping us keep track of the technology diffusion and to test the theories developed to explain its causes and effects.

³³ See, for example, the provocative idea of a "Zero-Sum Economy" put forward by Adair Turner (2018a; 2018b). Even if we do not fully buy the economic logic behind this idea, the notion that technology can fundamentally change patterns of economic activity and political competition should not be underestimated.

My relatively anodyne conclusion – “do more research” – can be justified by the still speculative nature of much of the research on this topic in developing countries. As shown in this paper, the debate on the effects of new technologies on developing countries is still in its infancy, and few people can confidently claim that they understand the implications of Industry 4.0 for developing countries. This is understandable, since technological paradigms are multifaceted phenomena. To understand their effects on the economy as well as their social and political implications, one must embrace this complexity and refrain from mechanistic approaches to prediction. As noted in the Marx quote in the epigraph, our understanding is likely to develop in parallel with the evolving effects of new technologies. The success of this research stream will therefore hinge on the ability to reach a solid understanding of the forces at play in a timely enough fashion to avoid the most pernicious effects of new technologies. Fortunately, it seems that we still have time.

References

- Acemoglu, D and Autor, D. (2011). 'Skill, tasks and technologies: implications for employment and earnings' in Card, D and Ashenfelter, O (Eds.), *Handbook of Labor Economics* (Vol. 4, pp. 1043-1171).
- Acemoglu, D and Restrepo, P. (2018a). 'The race between man and machine: implications of technology for growth, factor shares, and employment'. *American Economic Review*, 108(6), 1488-1542.
- Acemoglu, D and Restrepo, P. (2018b). 'Low-skill and high-skill automation'. *Journal of Human Capital*, 12(2), 204-232.
- Acemoglu, D and Restrepo, P. (2018c). 'Modeling automation'. NBER Working Paper No. 24321.
- Aghion, P, Jones, BF and Jones, C I. (2017). 'Artificial intelligence and economic growth'. NBER Working Paper No. 23928.
- Armstrong, S and Sotala, K. (2012). 'How we're predicting AI – or failing to' in Romportl, J et al. (Eds.), *Beyond AI: Artificial Dreams* (pp. 52-75). Pilsen: University of West Bohemia.
- Arntz, M, Gregory, T and Zierahn, U. (2016). 'The risk of automation for jobs in OECD countries: a comparative analysis'. *OECD Social, Employment and Migration Working Papers* (189).
- Athey, S. (2018). 'The impact of machine learning on economics'. In: A. Agrawal, J. Gans and A. Goldfarb, ed., *The Economics of Artificial Intelligence*. Chicago: University of Chicago Press.
- Autor, DH and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5), 1553-1597.
- Autor, D, Dorn, D, Katz, LF, Patterson, C and Van Reenen, J. (2017). 'The fall of the labor share and the rise of superstar firms'. Mimeo.
- Autor, D, Levy, F and Murnane, R J. (2003). 'The skill content of recent technological change: an empirical exploration'. *Quarterly Journal of Economics*, 118(4), 1279-1333.
- Bjorkegren, D and Grissen, D. (Forthcoming). 'Behavior revealed in mobile phone usage predicts credit repayment'. *World Bank Economic Review*.
- Blinder, AS. (2009). 'How many US jobs might be offshorable?' *World Economics*, 10(2), 41-78.
- Blumenstock, J, Cadamuro, G and On, R. (2015). 'Predicting poverty and wealth from mobile phone metadata'. *Science*, 350(6264), 1073-1076.

- Brynjolfsson, E and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W.W. Norton & Company.
- Brynjolfsson, E and McAfee, A. (2016, July/August). 'Human work in the robotic future: policy for the age of automation'. *Foreign Affairs*, 95(4).
- Brynjolfsson, E and Mitchell, T. (2017). 'What can machine learning do? Workforce implications'. *Science*, 358(6370), 1530-1534.
- Brynjolfsson, E, Mitchell, T and Rock, D. (2018). 'What can machines learn and what does it mean for occupations and the economy'. *AEA Papers and Proceedings*, 108, 43-47.
- Caselli, F. (2005). Accounting for Cross-Country Income Differences. In P. Aghion, & S. Durlauf (Eds.), *Handbook of Economic Growth* (Vol. 1, pp. 679-741).
- Caselli, F and Manning, A. (2018). 'Robot arithmetic: new technology and wages'. *American Economic Review: Insights*, Forthcoming.
- Chang, JH and Huynh, P. (2016). *ASEAN in Transformation: The Future of Jobs at Risk of Automation*. International Labour Organization.
- Comin, D and Hobijn, B. (2010). 'An exploration of technology diffusion'. *American Economic Review*, 100(5), 2031-2059.
- Comin, D and Mestieri, M. (2018). 'If technology has arrived everywhere, why has income diverged?' *American Economic Journal: Macroeconomics*, 10(3), 137-178.
- Cowen, T. (2010). *The Great Stagnation*. New York: Dutton.
- Dao, MC, Das, M, Koczan, Z and Lian, W. (2017). 'Why is labor receiving a smaller share of global income? Theory and empirical evidence'. IMF Working Paper.
- Dercon, S, Lippolis, N and Peel, S. (2018). *How do the Goods a Country Produces Affect its Growth? Background Paper for Programme on Rethinking African Paths to Industrial Development*. Oxford: Blavatnik School of Government and Centre for the Study of African Economies, University of Oxford.
- Domingos, P. (2012). 'A few useful things to know about machine learning'. *Communications of the ACM*, 55(10), 78-87.
- Frey, CB, Berger, T and Chen, C. (2018). 'Political machinery: did robots swing the 2016 presidential election?' *Oxford Review of Economic Policy*, 34(3), 418-442.
- Frey, CB and Osborne, MA. (2017). 'The future of employment: how susceptible are jobs to computerisation?' *Technological Forecasting and Social Change*, 114, 254-280.

- Gollin, D, Lagakos, D and Waugh, ME. (2014). 'The agricultural productivity gap'. *Quarterly Journal of Economics*, 129(2), 939-993.
- Gordon, RJ. (2016). *The Rise and Fall of American Growth: the US Standard of Living Since the Civil War*. Princeton: Princeton University Press.
- Graetz, G., & Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753-768.
- Halevy, A, Norvig, P and Pereira, F. (2009). 'The unreasonable effectiveness of data'. *IEEE Intelligent Systems*, 24(2), 8-12.
- Hallward-Driemeier, M and Nayyar, G. (2017). *Trouble in the Making? The Future of Manufacturing-Led Development*. Washington, DC: World Bank.
- Herrendorf, B, Rogerson, R and Valentinyi, A. (2014). 'Growth and structural transformation' in P. Aghion, P. and Durlauf, S.N. (Eds.), *Handbook of Economic Growth* (Vol. 2B, pp. 855-941). Elsevier.
- Hsieh, CT and Klenow, PJ. (2009). 'Misallocation and manufacturing TFP in China and India'. *Quarterly Journal of Economics*, 124(4), 1403-1448.
- James, G, Witten, D, Hastie, T and Tibshirani, R. (2013). *An Introduction to Statistical Learning: with Applications in R*. New York: Springer.
- Kahnemann, D. (2011). *Thinking, Fast and Slow*. London: Allen Lane.
- Karabarbounis, L and Neiman, B. (2014). 'The global decline of the labor share'. *Quarterly Journal of Economics*, 129(1), 61-103.
- Levy, F. (2018). 'Computers and populism: artificial intelligence, jobs, and politics in the near term'. *Oxford Review of Economic Policy*, 34(3), 393-417.
- Makridakis, S, Spiliotis, E and Assimakopoulos, V. (2018). 'Statistical and machine learning forecasting methods: concerns and ways forward'. *PLoS ONE*, 13(3).
- Marx, K. (1976). *Capital: A Critique of Political Economy* (Vol. I). (B. Fowkes, Trans.) London: Penguin Classics.
- McKinsey Global Institute (MGI). (2017a). *A Future that Works: Automation, Employment, and Productivity*. MGI.
- McKinsey Global Institute (MGI). (2017b). *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation*. MGI.

- Mitchell, T and Brynjolfsson, E. (2017). 'Track how technology is transforming work'. *Nature* (544), 290-292.
- Mokyr, J. (2014a). 'Secular stagnation? Not in your life' in Teulings, C. and Baldwin, R. (Eds.), *Secular Stagnation: Facts, Causes, and Cures*.
- Mokyr, J. (2014b). 'The next age of invention: technology's future is brighter than pessimists allow' in Anderson, BC. (Ed.) *City Journal*, pp. 14-20.
- Mokyr, J, Vickers, C and Ziebarth, NL. (2015). 'The history of technological anxiety and the future of economic growth: is this time different?' *Journal of Economic Perspectives*, 29(3), 31-50.
- Mullainathan, S and Obermeyer, Z. (2017). 'Does machine learning automate moral hazard and error?' *American Economic Review: Papers & Proceedings*, 107(5), 476-480.
- Mullainathan, S and Spiess, J. (2017). 'Machine learning: an applied econometric approach'. *Journal of Economic Perspectives*, 31(2), 87-106.
- Nedelkoska, L and Quintini, G. (2018). 'Automation, skills use and training'. *OECD Social, Employment and Migration Working Papers* (202).
- Neilson, J. (2014). 'Value chains, neoliberalism and development practice: the Indonesian experience'. *Review of International Political Economy*, 21(1), 38-69.
- PwC. (2017). *UK Economic Outlook 2017*.
- PwC. (2018). *Will Robots Really Steal our Jobs? An international analysis of the potential long term impact of automation*.
- Restuccia, D and Rogerson, R. (2017). 'The causes and costs of misallocation'. *Journal of Economic Perspectives*, 31(3), 151-174.
- Rodrik, D. (2013). 'Unconditional convergence in manufacturing'. *Quarterly Journal of Economics*, 128(1), 165-204.
- Rodrik, D. (2015). *Economics Rules: Why Economics Works, When it Fails, and How to Tell the Difference*. Oxford: Oxford University Press.
- Rodrik, D. (2016). 'Premature deindustrialization'. *Journal of Economic Growth*, 21(1), 1-33.
- Santos, I, Monroy, S and Moreno, M. (2015). *Technological Change and Labor Market Disruptions: Evidence from the Developing World*.
- Schlogl, L and Sumner, A. (2018). 'The rise of the robot reserve Army: automation and the future of economic development, work, and wages in developing countries'. *Center for Global Development Working Paper 487*.

- Schott, PK. (2003). 'One size fits all? Heckscher-Ohlin specialization in global production'. *American Economic Review*, 93(3), 686-708.
- Susskind, J. (2018). *Future Politics: Living Together in a World Transformed by Tech*. Oxford: Oxford University Press.
- Susskind, R. (2017). 'A Model of technological unemployment'. University of Oxford Department of Economics Discussion Paper Series.
- Susskind, R and Susskind, D. (2015). *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford: Oxford University Press.
- Turner, A. (2018a). 'Capitalism in the age of robots: work, income and wealth in the 21st century'. Lecture at School of Advanced International Studies, Johns Hopkins University, 10th of April 2018.
- Turner, A. (2018b). *The Zero Sum Economy*. Retrieved from Project Syndicate: <https://www.project-syndicate.org/commentary/zero-sum-economy-fuels-high-inequality-by-adair-turner-2018-08>
- UNCTAD. (2017). *Trade and Development Report 2017 – Beyond Austerity: Towards a Global New Deal*.
- Varian, HR. (2014). 'Big data: new tricks for econometrics'. *Journal of Economic Perspectives*, 28(2), 3-28.
- Vijg, J. (2011). *The American Technological Challenge: Stagnation and Decline in the 21st Century*. New York: Agora Publishing.
- World Development Report. (2016). *Digital Dividends*. Washington, DC: World Bank.
- World Development Report. (2019). *The Changing Nature of Work*. Washington, DC: World Bank.

