CHARLES UNIVERSITY
FACULTY OF SOCIAL SCIENCES
Institute of Economic Studies

Technology-driven unemployment: A meta-analysis

Bachelor's thesis

Author: Ondřej Zelený
Study program: Economics and Finance
Supervisor: Mgr. Petr Polák, M.Sc., Ph.D.
Year of defense: 2023

Declaration of Authorship

The author hereby declares that he or she compiled this thesis independently, using only the listed resources and literature, and the thesis has not been used to obtain any other academic title.

The author grants to Charles University permission to reproduce and to distribute copies of this thesis in whole or in part and agrees with the thesis being used for study and scientific purposes.

Prague, July 22, 2023

Ondrej Zeleny

Abstract

Will technological progress increase unemployment? Despite numerous attempts by researchers to answer this question, a consensus has yet to be reached since the findings provide contradicting results. To address this issue, we collect 516 estimates from 43 different studies and distinguish them on 31 characteristics to find the true net effect of technology advancements on employment. We observe almost negligible underlying effect based on multiple linear tests while discovering strong negative publication bias. Moreover, based on the Bayesian Model Averaging method, we identify eight factors significantly influencing the estimates of the effect - instrumental variable regression, group of other technology indicators, regional data, trends, journal impact, developed country, manufacturing and high-skill labour.

JEL Classification	E24, O31, O32, O33
Keywords	AI, robots, unemployment, technology
Title	Technology-driven unemployment: A meta-analysis

Abstrakt

Zvýší technologický pokrok nezaměstnanost? Přestože se mnoho výzkumníků snažilo na tuto otázku odpovědět, dosud nebylo dosaženo konsensu, neboť autoři poskytují protichůdné výsledky. Pro vyřešení tohoto problému jsme shromáždili 516 odhadů z 43 různých studií, které odlišujeme na základě 31 charakteristik, abychom zjistili skutečný efekt technologického pokroku na zaměstnanost. Na základě různých lineárních testů pozorujeme téměř zanedbatelný efekt, zatímco odhalujeme silný negativní vliv publikační selektivity. Navíc, s využitím bayesovského průměrování modelů, identifikujeme osm faktorů, které významně ovlivňují odhady tohoto efektu - regrese instrumentální proměnné, skupinu dalších ukazatelů technologie, regionální data, trendy, dosah časopisu, rozvinutou zemi, výrobu a pracovní sílu s vysokými dovednostmi.

Klasifikace JEL	E24, O31, O32, O33
Klíčová slova	umělá inteligence, roboti, nezaměstnanost, technologie
Název práce	Technologicky podmíněná nezaměstnanost: Metaanalýza

Acknowledgments

I am deeply grateful to my supervisor, Mgr. Petr Polák, M.Sc., Ph.D., for introducing me to the world of meta-analysis and providing invaluable guidance throughout the time of making this thesis. His constructive feedback greatly contributed to the overall quality of my research. Additionally, I extend my heartfelt appreciation to my good friend Bc. Petr Čala for his unwavering support, willingness to assist me with any uncertainties, and valuable insights that shaped the structure and content of this work.

Typeset in LaTeXusing the IES Thesis Template.

Bibliographic Record

Zelený, Ondřej: *Technology-driven unemployment: A meta-analysis*. Bachelor's thesis. Charles University, Faculty of Social Sciences, Institute of Economic Studies, Prague. 2023, pages 50. Advisor: Mgr. Petr Polák, M.Sc., Ph.D.

Contents

List of Tables

List of Figures

Acronyms

ICT Information and Communication Technology

TFP Total Factor Productivity

R&D Research and Development

SE Standard Error

FAT Funnel Asymmetry Test

PET Precision Effect Test

OLS Ordinary Least Squares

FE Fixed-Effects

RE Random-Effects

IV Instrumental Variable

BMA Bayesian Model Averaging

GMM Generalized Method of Moments

PIP Posterior Inclusion Probability

MCMC Markov Chain Monte Carlo

VIF Variance Inflation Factor

FMA Frequentist Model Averaging

AI Artificial Intelligence

Bachalor’s Thesis Proposal

Author	Ondřej Zelený
Supervisor	Mgr. Petr Polák, M.Sc., Ph.D.
Proposed topic	Technology-driven unemployment: A meta-analysis

Motivation The main question I intend to study is whether technological progress creates more jobs than it takes, hence what the net employment effect of technological change is.

New technologies may replace human labour, but at the same time, technology can create new jobs, because workers are needed to operate and guide these new technologies. The debate whether and how technological change can create more jobs than it destroys dates back to the 18th century and it has been accompanied by a surge of economic research on interactions between technology, labour, and the economy. Hotte *et al.* (2022) finds about 130 studies focusing on that topic.

Contribution Hotte *et al.* (2022) made a literature review of available studies and used an overview to show, that the support of the labour replacement effect is more than offset by the number of studies that support the labour-creating/reinstating and real income effects. Next to it, the simple literature review suggests the net impact of technology on labour to be rather positive than negative. The study does not follow up-to-date method for summarizing empirical literature but counts the positive and negative results. We intend to use this study as a starting point and by using meta-analytical techniques make an empirical estimation of net employment effect. This will allow to estimate not only if the effect is positive or negative, but also determine the magnitude of technological change on net employment. Given the review, there are almost 90 studies with such focus, which is more than enough for meta-analysis.

Methodology The methodology will follow standard meta-analytical method such as Gechert *et al.* (2022) The estimates from the primary research will be collected together with their precision and study design characteristics, dataset properties and other relevant metrics. Regression analysis then explains the heterogeneity of the outcomes of primary studies and also allows for the calculation of effect of the technological change on net employment without bias.

Outline

- Abstract
- Introduction
- Literature review and hypothesis
- Methodology
  - Relevant description of data
  - How tests were performed
- Results
  - Rejecting/not rejecting hypothesis
  - Interpretation of results

Core bibliography

Casey, G. (2018). Technology-Driven Unemployment. In 2018 Meeting Papers (No. 302). Society for Economic Dynamics.

Feldmann, H. (2013). Technological unemployment in industrial countries. Journal of Evolutionary Economics, 23(5), 1099-1126.

Bordot, F. (2022). Artificial intelligence, robots and unemployment: Evidence from OECD countries. Journal of Innovation Economics Management, 37(1), 117-138.

Blanchflower, D. G., and Burgess, S. M. (1998). New technology and jobs: comparative evidence from a two country study. Economics of Innovation and New Technology, 5(2-4), 109-138.

Postel–Vinay, F. (2002). The dynamics of technological unemployment. International Economic Review, 43(3), 737-760.

Diaz, M. S., and Tomas, F. J. Q. (2002). Technological innovation and employment: Data from a decade in Spain. International Journal of Production Economics, 75(3), 245-256.

Acemoglu, D and P Restrepo (2019), Automation and New Tasks: How Technology Displaces and Reinstates Labor, Journal of Economic Perspectives 33(2): 3-30.

Baldwin, R (2019), The globotics upheaval: Globalization, robotics, and the future of work, Oxford University Press.

Baldwin, R, J I Haaland, and A J Venables (2021), Jobs and technology in general equilibrium: A three-elasticities approach, CEPR Discussion Paper 15739.

Bessen, J (2020), Automation and jobs: when technology boosts employment, Economic Policy 34(100): 589-626.

Blanas, S, G Gancia, and S Y Lee (2019), Machines and workers: How different technologies affect different workers, VoxEU.org, 10 October.

Brynjolfsson, E and A McAfee (2014), The second machine age: Work, progress, and prosperity in a time of brilliant technologies, WW Norton and Company.

Hotte, K, M Somers, and A Theodorakopoulos (2022), Technology and jobs: A systematic literature review, Oxford Martin Working Paper Series on Technological and Economic Change No. 2022-2.

Mokyr, J, C Vickers, and N L Ziebarth (2015), The history of technological anxiety and the future of economic growth: Is this time different?, Journal of Economic Perspectives 29(3): 31-50.

Gechert Sebastian, Tomas Havranek, Zuzana Irsova, and Dominika Kolcunova (2022), Measuring Capital-Labor Substitution: The Importance of Method Choices and Publication Bias. Review of Economic Dynamics 45, 55-82

Chapter 1

Introduction

Whether technology is capable of replacing workers in their jobs has been a subject of debate, especially since the Industrial Revolution. Until now, the number of jobs lost to invention seemed to be proportionally balanced by the number of jobs created, taking into account the growing population. However, in the last few decades, we have witnessed an unprecedented surge in technological advancements in the form of computers, robots and the internet. Moreover, since the end of 2022, with the introduction of advanced Artificial Intelligence (AI) technologies, such as the ChatGPT chatbot, people have been more concerned than ever due to its impressive capabilities and widespread adoption. The implications of technological progress on employment are not only a subject of academic research but also a matter of public concern and policy debate. Understanding its relationship is crucial for policymakers, businesses, and individuals in order to make informative decisions, market policies, or strategies for workforce reskilling and preparation. In recent years many researchers have tried to answer the fundamental question: Will technological progress increase unemployment? However, the conclusive answer remains unclear.

This thesis aims to summarise existing empirical literature analysing the effect of technological progress on employment and offer valuable insights into the complex relationship between technology, employment, and the drivers that influence their interaction. Through a meta-analysis and the utilisation of various statistical approaches, the objective is to offer insight that might contribute to informed discussions and evidence-based policymaking. Meta-analysis is a method that is able to combine all relevant studies to empirically estimate the net effect on employment and identify potential drivers. Furthermore, additional statistical tests can help uncover any publication bias resulting from the selectivity of published studies based on their results. Our meta-analysis includes a range of linear tests to investigate potential publication bias and applies Bayesian Model Averaging (BMA) to explore heterogeneity. Our findings reveal a negligible effect of technological advancements on employment, accompanied by significant publication bias. At the same time, the estimates exhibit differences from the control variables, suggesting shifts in the labour market.

The remaining thesis sections are organized as follows: Chapter 2 provides an overview of the topic and meta-analysis background. Chapter 3 outlines our data collection methodology and presents initial analyses. Chapter 4 applies different statistical tests to examine publication bias and presents the corresponding results. Chapter 5 introduces the BMA method, presents its outcomes, and explains the variables used. Finally, Chapter 6 offers a conclusion to the thesis.

Chapter 2

Literature Review

2.1 Explaining technology

The literature addressing the displacement of workers due to automation and technological advances varies in many respects. This variability can be attributed to the complexities involved in adequately capturing and defining technological progress, as finding a universal measure to fully explain this process proves to be quite difficult.

One approach is to capture the increase in the use of robots. Acemoglu & Restrepo (2020), for example, used industry-level variation in the usage of robots, and Borjas & Freeman (2019) chose the number of industrial robots shipped to each industry as a benchmark. The impact of robots is easy to interpret because their implementation leads to pure automation, which directly affects the labour market.

Another popular method uses Information and Communication Technology (ICT) as a benchmark. The impact is always present, whether it is the expansion of broadband internet access (e.g. Atasoy 2013) or the investments in digitisation (e.g. Balsmeier & Woerter 2019). Although it is somewhat ambiguous how net employment changes in the context of ICT, it is still the most widely used method overall.

Innovation is also a helpful indicator. In general, there are two types of innovation: process and product innovation. The former is usually negatively correlated with employment, while the latter tends to affect it positively. Some researchers include both types in their studies (Capello & Lenzi 2013; Falk 2015). When both effects are present, they cancel each other out, resulting in a minor net employment change. However, some studies describe only one of them, usually process innovation, as could be found in Dachs *et al.* (2016).

The last major group of studies focuses on productivity. Total Factor Productivity (TFP) captures the impact of technological progress, innovation and other factors that increase the efficiency of the production process. It was popularised by Solow (1956) and is still present in today's papers. One of the authors dealing with unemployment based on this method are Autor & Salomons (2018).

There are also other approaches explaining technological progress that do not fall into any of the previous categories, such as measuring Research and Development (R&D) activities (e.g. Bogliacino & Pianta 2010) or the number of patents registered or cited (e.g. Autor & Salomons 2018).

2.2 Understanding the impact

The idea of technology-driven unemployment suggests an overall decline in job opportunities available. However, the impact of technological progress could be divided into two components - job creation and job destruction. These two tend to work against each other, and it is the result of their combination that interests us most. Nevertheless, it is not easy to capture both in a general regression. That is why many studies distinguish between the skills of the workers (e.g. Balsmeier & Woerter 2019; Autor *et al.* 1998), their age (e.g. Blanas *et al.* 2019), the industry in which they work (e.g. Cirillo 2017; Breemersch *et al.* 2019) or their gender (e.g. Borjas & Freeman 2019; Blanas *et al.* 2019).

These distinctions are essential for the explanation of shifts in the labour market. The categorisation makes it possible to capture the decline in job opportunities on the one hand and the increase on the other. The results usually show that job destruction is more common among the low-skilled group, which is logical if we assume that high-skilled workers are more adaptable. It is also consistent with the industry in which they work. The negative impact is more pronounced in manufacturing than in any other sector, where low-skilled workers are more prevalent. Based on the observations, women also seem to be a group that is more negatively affected by technological progress. However, some researchers come to the opposite conclusion (Faber 2020; Fu *et al.* 2021).

Another important factor that might influence the results of the studies is the level of analysis. Acemoglu & Restrepo (2020) examine data at the employee level, while Piva & Vivarelli (2004) are more interested in the firm level. The studies also include some other levels, such as region, country or occupation. Although it may not be clear what implications each level of analysis has, it should be taken into account. Interestingly, when comparing 37 different countries, de Vries *et al.* (2020) found that the introduction of robots reduces the employment share of routine manual jobs in high-income countries but not in emerging economies. Comparing the two, Fu *et al.* (2021) found that the developed countries had a more positive effect on the employment share, so it seems that the effect on developing countries is not that significant either way.

Some researchers argue that the negative impact on net employment is only in the short run because the displacement effect, unlike the creation effect, is immediate and is offset by the latter effect only in the long run. Nevertheless, Goaied & Sassi (2019) provide evidence to the contrary.

2.3 Meta-analysis

It is clear from the previous section 2.2 that views on impact, including the effect of individual factors, are quite diverse. Therefore, collecting all the numerical results of studies with varying survey methods and conclusions would make for interesting research. Fortunately, meta-analysis does just that, including explaining why conclusions vary. The method was developed in the late 1970s and early 1980s by Gene V. Glass, a social scientist, and Jacob Cohen, a statistician. It typically involves several steps, including identifying relevant studies, extracting data, assessing the quality of the studies, analysing the data and interpreting the results. The method allows for quantitatively synthesising findings from multiple studies, providing greater statistical power and precision than a single study. A potential limitation of meta-analysis is the risk of publication bias. It occurs when studies with significant findings are more likely to be published, leading to overestimating the effect size. However, practices such as funnel plots and sensitivity analyses can be used to identify and correct publication bias.

A meta-analysis, as a research methodology, has been experiencing up-to-date trends and continues to evolve in various disciplines. With the advancements in statistical techniques and the increasing availability of research data, meta-analysis has become an essential tool for synthesising and integrating research findings. Researchers increasingly recognise the value of meta-analysis in providing robust evidence and guiding decision-making. Furthermore, meta-analyses are commonly published in high-quality academic journals. Due to the rigorous methodology and comprehensive approach, meta-analyses are often regarded as influential studies within their respective fields. Many top-tier journals actively encourage and publish meta-analyses as they contribute to accumulating knowledge and provide valuable insights into the research landscape.

As Nordmann (2012) suggests, the popularity of meta-analysis has been on the rise in recent years. Researchers across various disciplines use meta-analytic approaches to address research questions and produce more reliable and generalisable results. The growth in popularity can be attributed to several factors, including the increasing emphasis on evidence-based practice, the need for systematic reviews of research evidence, and the recognition of meta-analysis as a powerful tool for synthesising findings from multiple studies.

In terms of recent uses, meta-analysis has been applied in diverse areas. For example, in healthcare and medicine, meta-analyses are conducted to assess the effectiveness of treatments, evaluate the safety of interventions, and explore the impact of risk factors on health outcomes (Dong *et al.* 2011; Itani *et al.* 2017). In psychology and social sciences, meta-analyses are utilised to investigate the effectiveness of interventions, examine the strength of associations between variables, and identify moderators and mediators of effects (Cuijpers *et al.* 2016; Curran *et al.* 2015). Additionally, meta-analyses are increasingly being conducted in fields such as education, economics, environmental sciences, and more (Havranek *et al.* 2016).

Overall, meta-analysis continues to evolve and adapt to changing research trends. Its widespread use, publication in high-quality journals, increasing popularity, and diverse applications across disciplines highlight its significance in evidence synthesis and decision-making processes. As researchers continue to refine the methodology and address methodological challenges, meta-analysis will likely remain a prominent approach for integrating research evidence and advancing scientific knowledge.

Chapter 3

Methodology

3.1 Literature search

One of the more used practices in meta-analysis recommends using Google Scholar to search for primary studies relevant to the research question based on title, abstract, keywords, and number of citations. We base our literature search on a very recent literature review conducted by Hotte *et al.* (2022) that already identified relevant studies but did not employ the meta-analysis.

Hotte *et al.* (2022) presents 127 studies published between 1988 and 2021 that reveal technological change and its impact on the labour market. To avoid unintended heterogeneity in the data, they divide the studies into five groups according to the effect they describe similarly as in 2.1. These groups are (1) information and communication technology (ICT); (2) robot diffusion; (3) innovation surveys; (4) productivity; and (5) a category that includes all alternative indicators. They also report three different impacts of technology on the labour market. The first is the substitution effect, which is the most direct and captures the ability of a firm to reduce its workforce after introducing a particular technological advancement. The second is the reinstatement effect, which occurs when a technological change corresponds to creating jobs associated with that improvement. These two effects usually go against each other, and the question is which one is more substantial. The last one that Hotte *et al.* (2022) examine is the income effect. The latter is indirect and harder to interpret. I, therefore, exclude from the analysis studies that describe only the income effect, but this only applies to ten of them.

Another condition for the inclusion of the study in the meta-analysis was the use of an empirical method that also provides the precision of estimates.

The literature review by Hotte *et al.* (2022) also includes papers using either descriptive or alternative methods (mostly simulations). These two were not included since they did not offer estimates and their corresponding standard errors. Thus, we were left with about 100 studies that were considered for inclusion in our data set.

3.2 Data collection

Based on the restrictions made in the 3.1, studies must satisfy two additional criteria. First, the dependent variables should be comparable in some way. In our case, any association with employment should suffice. However, there are some limitations, for instance, estimates such as the *automation potential* (Arntz *et al.* 2017) or the *high-skill/low-skill labour ratio* (Maurin & Thesmar 2004). The second is related to the explanatory variable, which should be an indicator of technological progress. Although this seems to be an unnecessary condition in the context of the included studies, it is possible to encounter variables such as the *routine-employed share* (Autor & Dorn 2013; Autor *et al.* 2015) or the *new goods* (Xiang 2005) that says nothing explicit about technology.

The main statistics collected apart from the estimates are standard errors and the number of observations of each regression. The baseline equation used in those regressions could be represented by one similar to the one found in Bogliacino & Pianta (2010).

$$y_{it} = \beta x_{it} + \gamma z_{it} + u_i + v_{it} \tag{3.1}$$

where y_{it} is the employment variable, x_{it} the technological indicator variable, z_{it} the vector of other regressors, u_i the individual effect and v_{it} the random disturbance, for industry i and time t.

To differentiate between studies, we collected more than 30 aspects specifying the variable definition (technological indicators, lagged variables), data characteristics (number of observations, the average year of the first and last year the data was collected, how long the period was and the level of analysis), estimation method, structural variation (gender, skill level, region) and publication characteristics (impact factor according to RePEc, number of citations).

3.3 Initial analysis

After collecting and visualising the data using a funnel plot, we also decided to exclude the work of Dekle (2020), who used *number of workers* as the dependent variable, and a study conducted by Feldmann (2013) with *unemployment rate* as the response variable. Although these two do not necessarily violate previous restrictions and provide insight into the trends of labour shifts, their estimates are very different from all the others and, therefore, considered outliers[1]. That leaves the data set with 516 comparable observations from 43 studies, enough for a meaningful meta-analysis. To limit the dominance of a few estimates with small standard error, we calculate the precision of our estimates using Standard Error (SE), which is $1/SE$ and winsorise[2] at the 5% level. This level is still acceptable for balancing between artificial intervention to the data and the stability of our results. We could choose a lower level, but the precision of some studies was high, which could drive our results to misleading conclusions. At this point, we can visualise the behaviour of the estimates.

Some box plots in the figure 3.1 are not very diverse and almost look like point estimates. These are the ones that had only less than five estimates in their study which concerns almost half of the studies included in this analysis. The effect is, on average, more negative than positive in our sample, as the red line suggests. That is confirmed by another initial evaluation method which is the mean statistics across various groups of data presented in Table 3.1.

We can make some initial remarks from a brief look at the table. First, as we established in the previous paragraph, the baseline effect has a negative mean of -0.188 that goes a little down to -0.114 when weighted by the number of estimates in each study. Some variables then display exciting results. When coming from the top to the bottom of the table, the first variable that stands out is 'IV' in *Methodology* with the mean of -0.584 that remains almost the same (-0.582) when weighted, suggesting that using the instrumental variable estimation method gives us more negative results. As we look at the *Technology indicator*, there is a clear difference mostly between the 'ICT' and all the other indicators. The 'ICT' is the only positive amongst them, with a mean of 0.022 that goes even higher when weighted (0.185). However, the most significant differences are to be seen in the *Level of analysis* as the 'Macro' level shows the

[1]Dekle's estimates are more than 100 times smaller than the rest of our data set, while Feldmann presents estimates at least ten times larger.

[2]A technique used to address extreme values by replacing them with a specified percentile of the data, minimising their impact on the statistical analysis

Figure 3.1: Estimate across individual studies

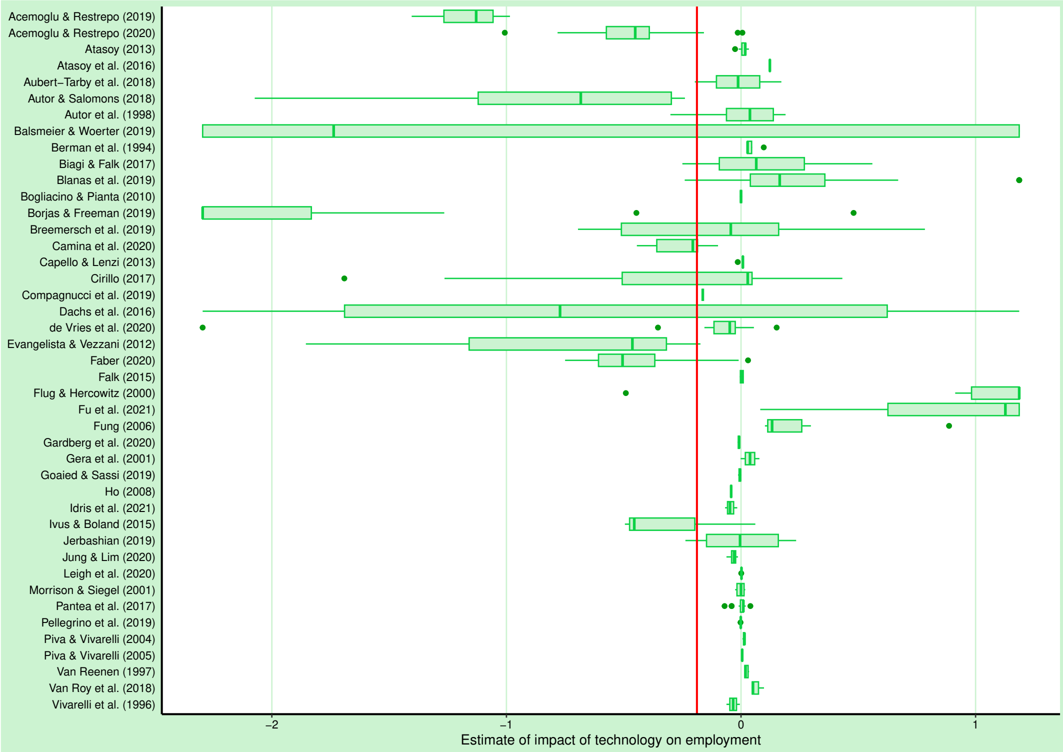

Note: This figure shows a box plot of the estimates of the effect technology has on employment across individual studies after winsorisation. The red line represents the average of all estimates.

mean of 0.201 compared to the 'Regional' (-0.606), and the difference even increases after weighting by the number of observations to (0.382 & -0.547).

Controlling for 'Manufacturing' versus 'Non-manufacturing' gives us unsurprising results described already in Section 2.2 where the 'Manufacturing' variable has a more negative mean (-0.230) than 'Non-manufacturing' (-0.083). When weighted by the number of observations, the results remain almost unchanged for 'Manufacturing' but not for 'Non-manufacturing', which drops to (-0.006). A similar story could be seen when controlling for 'High-skill' versus 'Medium-skill' and 'Low-skill' workers. As Section 2.2 already suggested, the more adaptable 'High-skill' labour has a positive mean of 0.253 compared to negative -0.289 in the case of 'Low-skill'. However, the means of variables of workers' skill change drastically when weighted to even positive mean of 'Low-skill' workers (0.208). A possible explanation for this might be the small number of observations (less than 10%). The same phenomenon is observed when controlling for the 'Female' variable. Therefore we should not draw any conclusions from them.

Furthermore, it is essential to keep in mind that those results do not provide us with any robust result but rather an insight into the distribution of our data and possible factors that are the source of heterogeneity in the estimates of the effect of technological progress on employment. We recommend reviewing Chapter 5 for more cutting-edge findings.

Table 3.1: Mean statistics across various subsets of data

	Unweighted			Weighted			
	Mean	95% conf. int.		Mean	95% conf. int.		No. of obs.
All estimates	-0.188	-0.284	-0.092	-0.114	-0.210	-0.018	516
Methodology							
OLS	-0.105	-0.235	0.025	0.109	-0.021	0.239	220
Fixed-effects	0.027	-0.022	0.077	0.025	-0.024	0.075	9
GMM	0.040	-0.008	0.087	0.028	-0.019	0.076	38
IV	-0.584	-0.974	-0.194	-0.582	-0.972	-0.192	92
Other method	-0.140	-0.249	-0.030	-0.385	-0.495	-0.276	157
Technology indicator							
ICT	0.022	-0.123	0.166	0.185	0.040	0.329	223
Innovation	-0.166	-0.431	0.100	-0.309	-0.575	-0.044	90
Robots	-0.443	-0.602	-0.284	-0.483	-0.642	-0.325	165
Other indicators	-0.365	-0.532	-0.197	-0.672	-0.840	-0.505	38
Level of analysis							
Macro	0.201	0.015	0.387	0.382	0.195	0.568	75
Meso	0.002	-0.058	0.062	0.160	0.100	0.220	191
Micro	-0.248	-0.600	0.104	-0.296	-0.648	0.057	109
Regional	-0.606	-0.756	-0.455	-0.547	-0.697	-0.396	141
Regression specifics							
Dependent lag	-0.078	-0.143	-0.012	-0.193	-0.259	-0.128	91
Independent lag	0.079	-0.024	0.182	0.183	0.080	0.287	144
Trends	0.085	-0.057	0.228	-0.094	-0.237	0.048	82
Time control	-0.442	-0.628	-0.256	-0.687	-0.872	-0.501	178
Region and journal importance							
Developed country	-0.225	-0.337	-0.114	-0.106	-0.218	0.005	428
Developing country	-0.007	-0.154	0.140	-0.199	-0.346	-0.052	88
Top 50 journals	-0.381	-0.558	-0.204	-0.132	-0.309	0.045	236
Other journals	-0.025	-0.117	0.066	-0.041	-0.132	0.051	280
Labour characteristics							
Manufacturing	-0.230	-0.465	0.004	-0.184	-0.418	0.051	109
Non-manufacturing	-0.083	-0.248	0.082	-0.006	-0.171	0.159	44
High-skill	0.253	-0.279	0.784	0.033	-0.499	0.565	46
Medium-skill	-0.154	-0.538	0.229	0.103	-0.281	0.486	46
Low-skill	-0.289	-0.799	0.221	0.208	-0.302	0.719	32
Male	0.147	-0.182	0.476	0.187	-0.142	0.516	19
Female	0.045	-0.628	0.717	0.204	-0.468	0.877	19

Note: This table shows summary statistics of employment change estimates across various data subsets. Unweighted comes from the original data set. Weighted means that the estimates are weighted by the inverse number of estimates reported by each study. For detailed explanation of the variables, see table 5.1. OLS = Ordinary Least Squares, GMM = Generalized Method of Moments, IV = Instrumental Variable, ICT = Information and Communication Technology

Chapter 4

Publication Bias

Publication bias refers to the phenomenon where the publication and dissemination of research findings are influenced by the direction or statistical significance of the results. Studies with statistically significant or positive results are more likely to be published and receive greater attention than those with nonsignificant or negative results. This selective publication of studies creates an overrepresentation of positive findings in the scientific literature, leading to a limited understanding of the actual state of knowledge in a specific field. More could be found in Stanley (2005).

Publication bias can arise from various factors, including the preference of researchers and editors to submit only statistically significant or exciting results for publication. The consequences of publication bias are significant, as it distorts the evidence base by overestimating treatment effects, intervention effectiveness, or the strength of associations between variables. This distortion can affect crucial decision-making processes in clinical practice, policy development, and resource allocation.

To mitigate publication bias, several strategies have been proposed and implemented. The most popular one is the funnel plot (Egger *et al.* 1997). In funnel plot estimates of the effect are plotted against precision which is $1/SE$ as already discussed in section 3.3. When publication bias is absent, the funnel plot is expected to display a symmetrical inverted funnel shape. Within this shape, less precise estimates are scattered around the average effect estimate, while more precise estimates tend to converge towards it. However, if publication bias exists, the funnel plot may exhibit asymmetry. This can be identified by observing a distorted or skewed plot shape or noticeable gaps, which occur when studies with insignificant effects are omitted. Such asymmetry indicates potential publication bias and warrants further investigation into the reliability and generalizability of the meta-analysis findings.

Figure 4.1 presents the obtained funnel plot, which displays a mainly symmetric shape centred around 0. While most estimates align with the expected funnel pattern, there are a few outliers, particularly on the negative side. This deviation from the expected pattern suggests the potential presence of bias that necessitates further investigation.

Figure 4.1: Funnel plot

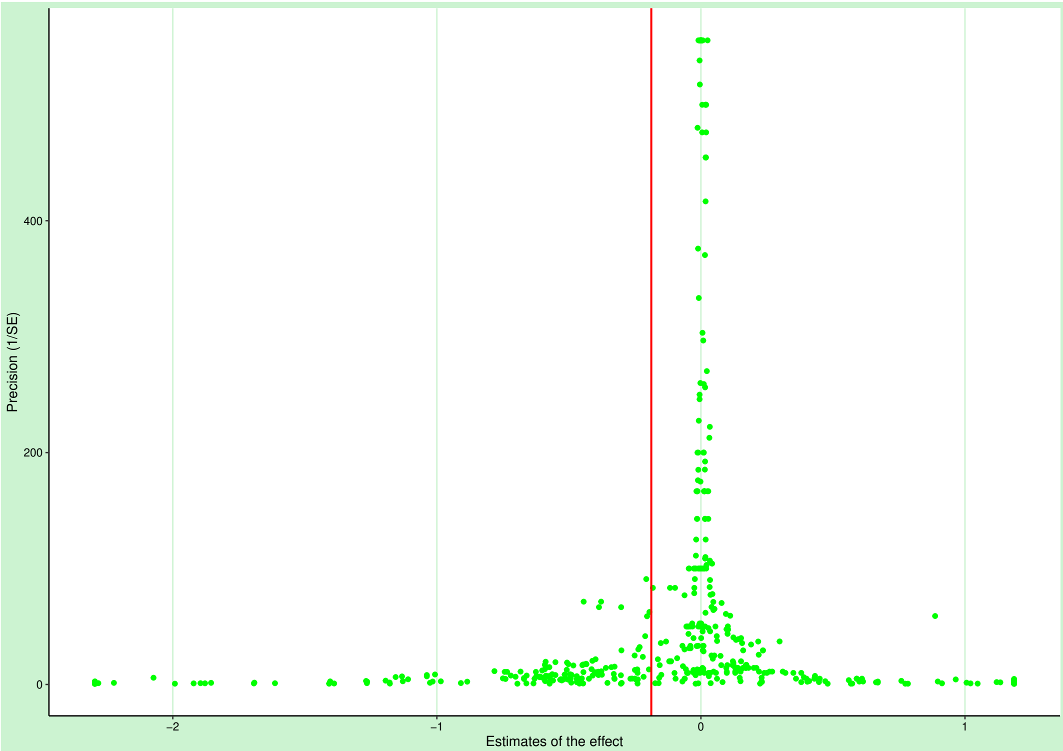

Note: This figure shows a funnel plot of the estimates as described by Egger *et al.* (1997). The plot does not display any significant asymmetry implying publication bias. The red line represents the average of all estimates.

4.1 FAT-PET

To extend Egger's test and obtain more reliable findings regarding publication bias, we conduct a Funnel Asymmetry Test (FAT) - Precision Effect Test (PET). As proposed by Stanley (2008), the FAT-PET method involves assessing the potential relationship between the estimates and their corresponding standard errors through the use of a regression model. If a correlation is observed between these variables, it indicates the presence of publication bias in our sample.

We estimate the following equation:

$$estimate_{ij} = \beta_0 + \beta_1 * SE_{ij} + u_{ij} \tag{4.1}$$

where $estimate_{ij}$ is the i-th estimate of the j-th study with its standard error SE_{ij}. β_0 denotes the size of the effect stripped of the bias (*Effect beyond bias*), β_1 represents the size of the bias itself (*Publication bias*) and u_{ij} captures the disturbance.

The Table 4.1 below presents the regression results for equation 4.1, utilizing different linear models with clustered standard errors at the study level and assuming exogeneity. We applied the standard Ordinary Least Squares (OLS) regression together with Fixed-Effects (FE) and Random-Effects (RE) models[1]. Additionally, we conducted two weighted regressions using the inverse of the number of observations per study (Study) as weights, ensuring equal influence from each study on the results. We also incorporated the inverse of the standard error (Precision) as suggested by Stanley & Doucouliagos (2017) to address heteroskedasticity.

Table 4.1: Tests for publication bias

	OLS	**FE**	**RE**	**Study**	**Precision**
SE	-0.578***	-1.148***	-0.566***	-0.423**	-1.148***
Publication bias	(0.139)	(0.050)	(0.100)	(0.160)	(0.235)
Constant	-0.039	0.006***	-0.043*	-0.033	0.006***
Effect beyond bias	(0.024)	(0.001)	(0.017)	(0.031)	(0.001)
Studies	43	43	43	43	43
Observations	516	516	516	516	516

Note: This table shows the regression results for equation 4.1. OLS = Ordinary Least Squares, FE = Fixed-Effects, RE = Random-Effects, Study = weighted by the number of observations in a study, Precision = weighted by the inverse of SE. The first row represents *Publication bias* and the second row *Effect beyond bias*. Standard errors, clustered at the study level, are included in parentheses. ***p<0.001, **p<0.01, *p<0.05

Based on the findings presented in Table 4.1, we can conclude that the baseline effect is primarily negligible, as it is very close to zero. The "Effect beyond bias" estimates from OLS and weighted by the number of observations yield relatively similar results, showing a slightly negative trend but lacking statistical significance. The Random-Effects estimates also demonstrate similar patterns but with slightly higher statistical significance. Whereas the estimates from Fixed-Effects and weighted by precision exhibit statistically significant

[1] The Hausman test could have been performed to determine the suitable model between FE and RE. While FE is typically preferred, we include both for illustrative purposes.

and positive results, albeit closer to zero. This outcome aligns with the observations from the funnel plot depicted in Figure 4.1, which centres around zero.

However, the examination of publication bias yields more intriguing insights. Across all models, we observe a substantial and statistically significant negative bias at the 1% level of statistical significance. This implies that more negative results (e.g. those leading to a negative effect on employment) are preferred in the publication process, which may generate greater excitement among potential readers.

4.2 Relaxing the exogeneity assumption

So far, we have operated under the assumption that our data set satisfies the exogeneity assumption, which stipulates that the original effect is uncorrelated with the standard errors. However, in our case, endogeneity is likely present due to variations in the estimation techniques used across studies, which can simultaneously influence both the estimates and the standard errors.

To address this issue, we employ Instrumental Variable (IV) regression in our analysis. In line with common practice, we select the instrument by transforming the number of observations. Since the number of observations is inherently related to the standard error, this instrument helps address potential endogeneity concerns while also being potentially uncorrelated with unobserved variations such as estimation techniques (Gechert *et al.* 2022). Specifically, we utilize the logarithm of the number of observations from the primary study, clustered at the study level, as this instrument performed the best based on statistical tests to evaluate its validity. The results of the IV regression are presented in Table 4.2.

The Instrumental Variable (IV) test results exhibit a similar trend to the other linear tests we conducted. However, the estimate of the effect beyond bias increased in the positive direction. Moreover, the effect remains statistically significant at the 5% level, suggesting that the true underlying effect may be positive yet still close to zero. Furthermore, the publication bias considerably increases in its magnitude while keeping its significance from the previous tests.

In summary, the tests conducted in this chapter present similar outcomes, indicating a consistent pattern of findings. The majority of these tests reveal an almost negligible impact of technological progress on employment. The effect beyond bias is minimal and mostly lacks statistical significance. Moreover, the

Table 4.2: Relaxing the exogeneity assumption

	IV
SE	-1.996***
Publication bias	(0.528)
Constant	0.327*
Effect beyond bias	(0.137)
Studies	43
Observations	516

Note: This table shows the instrumental variable regression results for equation 4.1. Standard errors, clustered at the study level, are included in parentheses. ***p<0.001, **p<0.01, *p<0.05

analysis consistently demonstrates a significant negative publication bias across all tests, indicating a higher probability of selective reporting and publication of studies with negative results. It is important to note that non-linear tests were not included in this chapter due to their complexity and being mostly beyond the scope of a bachelor's level education. Implementing non-linear tests could be a potential direction for further research.

Chapter 5

Heterogeneity of estimates

It is essential to explore the potential drivers contributing to heterogeneity across individual observations to understand the variations within our data set regarding technology-driven unemployment. We begin by examining the variables we have collected and their potential influences based on existing literature. Subsequently, we employ the Bayesian Model Averaging (BMA) method to provide a more robust analysis that accounts for model uncertainty and allows us to thoroughly investigate and control for potential factors contributing to the observed variations.

5.1 Data examination

Having collected data from primary studies, we were able to categorize the observations based on several factors. In total, we assembled 31 explanatory variables, some of which were presented in Chapter 3. However, to avoid the dummy variable trap, we only use 28[1] of them in the Bayesian model averaging. All 31 variables are listed in Table 5.1 with their brief explanation and summary statistics. To justify our selection of these specific characteristics, we will systematically review each group of variables and provide a rationale for their inclusion.

Methodology As previously mentioned, all estimates were obtained through regression analysis. The most commonly employed regression method, utilized by 43% of researchers, was Ordinary Least Squares (OLS), which is generally considered the most straightforward approach. Instrumental Variable (IV) re-

[1] Specifically, we omit variables 'OLS', 'Robots' and 'Meso'.

gression was another frequently used method, accounting for approximately 18% of the studies. Although Fixed-Effects and Generalized Method of Moments (GMM) regressions were less prevalent, we still included them as separate variables in our analysis. If a specific regression method was employed in fewer than two studies or was not explicitly stated, it was categorized as 'Other method'. For instance, Falk (2015) employed Quantile Regression, while Compagnucci *et al.* (2019) utilized the Panel VAR approach. Some researchers even employed multiple types of regressions in their studies to ensure the robustness of the results (Piva & Vivarelli 2004; Capello & Lenzi 2013).

Although we did not anticipate a significant impact of the method used on the estimates, it is still a relevant factor to consider, which we confirmed in Chapter 3 during the Initial analysis.

Technology indicator Several determinants were applied to explain the baseline effect of technological progress on employment. The most represented is the Information and Communication Technology (ICT) at 43%, which covers *broadband expansion*, *digitization* or *change in the number of mobile users*. The second most used indicator (32%) determined the effect of the use of Robots, whether it was the *exposure to foreign robots*, *change in adoption* or the raw *number* of them. Compared to ICT, we would expect the outcome to be different since the application affects more manufacturing than services. The last indicator that did not fall into the 'Other indicators' is Innovation. What is worth mentioning about this category is that it covers both process and product innovation. Those two usually displayed the opposite effect on employment. We decided to put them into the same category because they still fit under Innovation, and there are not used that often to separate into single dummy variables. Within the category of 'Other indicators', we included *R&D expenditure* or *number of patents*, but also Total Factor Productivity (TFP), which we described in section 2.1. While TFP was extensively addressed in the literature review conducted by Hotte *et al.* (2022), which served as the basis for constructing our dataset, only a limited number of studies met our inclusion criteria when it came to using this particular indicator.

Level of analysis To categorize the analysis based on its scope, we employed four different levels. The primary studies utilized two distinct classifications: one based on employee, firm, occupational, and industry levels and the other on macro, meso, and micro levels. After careful consideration, we opted to adopt the latter classification, which is more commonly used, while also incorporating the regional level. The majority of our data set was derived from the meso-level (37%), which captures industry-specific data. The regional level (27%) comprises data from various levels within a specific geographic region. The macro and micro levels correspond to country-level and firm/individual-level data, respectively. The prevalence of the meso-level can be attributed to its significance in analyzing employment and employment changes within specific sectors.

Regression specifics As we look at the regression, we identify some study-specific approaches. One is the 'Dependent lag', which accounts for approximately 18% of the observations. This variable is equal to one whether there is one or more employment variable lags, maximum being two (Atasoy 2013; Van Reenen 1997). The 'Independent lag' is set to one when the technology indicator variable is lagged. This inclusion usually notably influences the non-lagged estimate and is observed in 28% of our observations. Some researchers even utilize up to six lags in their studies (Van Reenen 1997). 'Trends' represents a variable which denotes whether the author detrended the regression. It covers not only time trends but also industry or zone trends. The last variable in this group is 'Time control', which denotes when the regression includes time-related variables as control variables in the model (35%).

Study specifications Next, we consider a set of variables that provide information specific to individual studies. The majority of these variables are self-explanatory, and there are no noteworthy characteristics or distinctive features worth mentioning., as their descriptions can be found in Table 5.1. However, one variable may be unfamiliar to readers, namely 'Journal impact'. This variable is derived from the ideas.repec.org website, which calculates the impact factor of a journal based on the number of citations it has received.

Country and labour characteristics The last group of variables focuses on country and labour characteristics. Initially, we collected information on specific countries where the studies were conducted. However, instead of creating individual dummy variables for each country, we opted to categorize them as either 'Developed' or not. This classification is based on the United Nations classification, which can be found at un.org. The data sets from various countries were categorized based on their predominant representation, which was primarily composed of developed countries.

As mentioned earlier, the meso-level of analysis, which focuses on the sector in which the study was conducted, is the most common approach. Typically, the distinction is made between 'Manufacturing' and 'Non-manufacturing'. We also included routine jobs if explicitly stated in the manufacturing category, while the non-manufacturing category covers services and management roles.

Regarding individual characteristics, we constructed variables such as 'High-skill', 'Medium-skill', and 'Low-skill' based on educational level. The 'High-skill' category represents individuals with a university degree or higher, while the 'Low-skill' category represents those with basic education. 'Medium-skill' includes individuals with education levels that fall between these two categories. Additionally, some studies distinguished between males and females, so we included gender as separate variables in our analysis.

Table 5.1: Description and summary statistics table

Variable	Description	Mean	SD
Effect	The effect of technological progress on employment	-0.188	0.726
Standard error	The standard error of the main effect	0.258	0.406
Methodology			
OLS	=1 if the authors use Ordinary least squares	0.426	0.495
Fixed-effects	=1 if the authors use Fixed-effects estimation	0.017	0.131
GMM	=1 if the authors use Generalized method of moments estimation	0.074	0.261
IV	=1 if the authors use Instrumental variables estimation	0.178	0.383
Other method	=1 if the authors use other method of estimation	0.304	0.461
Technology indicator			
ICT	=1 if the independent variable in the regression is connected to Information and Communication Technology	0.432	0.496
Innovation	=1 if the independent variable in the regression is connected to innovation	0.174	0.380
Robots	=1 if the independent variable in the regression is connected to robots	0.320	0.467
Other indicators	=1 if the independent variable in the regression is connected to other technology indicator	0.074	0.261
Level of analysis			
Macro	=1 if the study uses macro data	0.145	0.353
Meso	=1 if the study uses meso data	0.370	0.483
Micro	=1 if the study uses micro data	0.211	0.409
Regional	=1 if the study uses regional data	0.273	0.446
Regression specifics			
Dependent lag	=1 if the authors uses dependent variable lag in the regression	0.176	0.381
Independent lag	=1 if the authors use independent variable lag in the regression	0.279	0.449
Trends	=1 if the authors use trends in the regression	0.159	0.366
Time control	=1 if the authors control for time in the regression	0.345	0.476
Study specifications			
Time horizon	The number of years over which the data set was collected	12.841	7.577
Average year	The average year calculated from the time horizon	2000.653	10.029
Observations	The number of observations associated with the estimate	5455.002	12963.811
Journal impact	The journal impact factor from RePEc	32.070	28.527
Citations	The number of citations of the study	483.833	972.959
Country and labour characteristics			
Developed country	=1 if the study was conducted in a developed country	0.829	0.376
Manufacturing	=1 if the authors control for manufacturing sector in the regression	0.211	0.409
Non-manufacturing	=1 if the authors control for non-manufacturing sector in the regression	0.085	0.280
High-skill	=1 if the authors control for high-skill labour in the regression	0.089	0.285
Medium-skill	=1 if the authors control for medium-skill labour in the regression	0.089	0.285
Low-skill	=1 if the authors control for low-skill labour in the regression	0.062	0.241
Male	=1 if the authors control for males in the regression	0.037	0.189
Female	=1 if the authors control for females in the regression	0.037	0.189

Note: This table presents definitions and summary statistics of each variable in our final data set.
SD = standard deviation

5.2 Bayesian model averaging

The inclination might be to use the OLS to obtain numerical results for our meta-regression analysis. However, the model should only include some variables to avoid misspecification. The challenge lies in the fact that selecting the appropriate variables from our pool of 28 options would mean exploring a staggering number of combinations $2^{28} \approx 270000000$. Not only is this computationally intensive, but it also demands a significant amount of time. This situation is commonly known as model uncertainty. Fortunately, Bayesian Model Averaging (BMA) offers a solution by allowing for the simultaneous consideration of multiple models and providing a weight, known as the posterior model probability, to each of them. This approach allows us to assign a Posterior Inclusion Probability (PIP) to each variable by summing the posterior model probabilities in which the variable was included.

Employing the Markov Chain Monte Carlo (MCMC) algorithm makes it possible to reduce the number of models without sacrificing much information in order to reduce the runtime of the code (Gechert *et al.* 2022). In order to effectively utilize the BMA, the g-prior and model prior need to be specified. The g-prior determines the weight assigned to the prior probability of each coefficient. In our case, we employ the unit information g-prior, which assigns a weight equivalent to the information provided by a single observation. The model prior probability is then used to weight models. Typically, a uniform model prior is employed, where each model has an equal prior probability. However, due to the high number of variables in our data set, we opt for the dilution prior, which considers collinearity within the model (Bajzik *et al.* 2023). This choice is particularly favourable given the large number of variables involved.

Before proceeding with the application of BMA, it is essential to address the collinearity issue. In a meta-analysis, it is common practice to exclude variables that exhibit problematic relationships with the response variable, typically indicated by a Variance Inflation Factor (VIF) exceeding 10. Surprisingly, all our variables fall below this threshold, indicating no significant collinearity issues. Therefore, no further action is required in this regard.

Figure 5.1 depicts the outcomes of the BMA analysis. The variables are displayed on the vertical axis, arranged in order of their PIP. The models are organized on the horizontal axis based on their posterior model probabilities, with the most probable models positioned towards the left. Each column represents one model, where the coloured variables indicate their inclusion. In this representation, the green colour indicates a negative effect on the estimate, while the red color signifies a positive effect.

Figure 5.1: Bayesian model averaging results

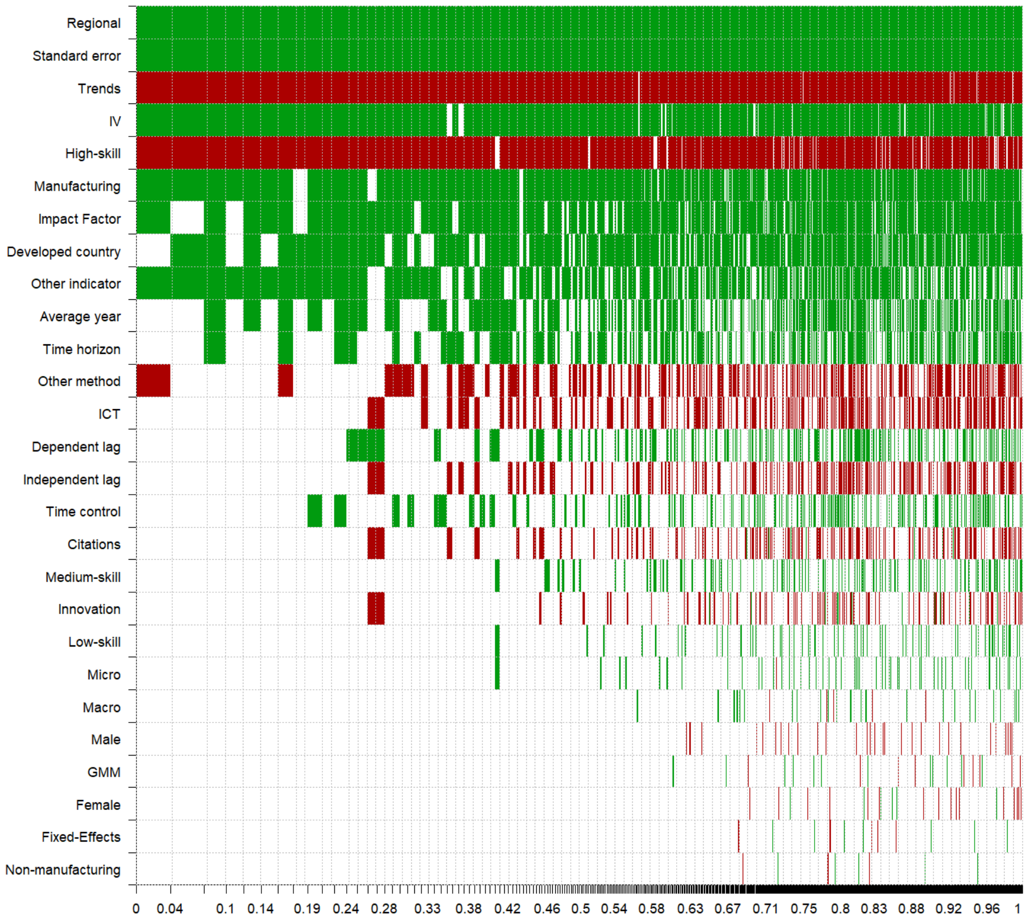

Note: The figure displays the Bayesian model averaging results with unit information g-prior and dilution model prior specification. The response variable is the employment change. Each row represents an explanatory variable as they are ranked on the vertical axis based on their PIP. Each column represents a single model as cumulative posterior model probabilities are displayed on the horizontal axis. Green colour - the effect of the variable is negative. Red colour - the effect of the variable is positive. Numerical results are presented in Table 5.2 and an explanation of the variables in Table 5.1.

The numerical results in Table 5.2 provide guidance on which variables to include in the final model and their impact on employment. Inclusion in the model is determined by the PIP, as discussed earlier, with a threshold of $PIP > 0.75$ indicating significant evidence for the effect (Kass & Raftery 1995). Ten variables in our analysis meet this threshold, including the constant and standard error. The interpretation of these two variables is similar to that discussed in Chapter 4, where the constant represents the effect beyond bias and the standard error reflects publication bias. While the constant does not provide specific information about the magnitude of the baseline effect, the posterior mean -0.386 of the standard error serves as additional evidence of the negative publication bias observed in the literature on technology-driven unemployment.

Upon examining the remaining significant explanatory variables, it is evident that IV regression exhibits a strong negative correlation compared to all other estimation methods (-0.411). As for why this is the case, we can only speculate, but the fact of the matter is that IV regression exhibits systematically different results than any other regression.

As for the technology indicator, 'ICT', 'Innovation' and 'Robots' display similar results. Only 'Other indicators' seem to affect employment more negatively (-0.313). That is interesting since using the conventional indicators of technological progress does not affect the findings. When using *TFP* or other less conventional indicators such as *R&D expenditures* or *number of patents*, we do observe a visible change.

The variable with the highest coefficient among all the variables examined is the 'Regional' level of analysis (-0.606). The rationale behind the substantial impact of this factor might be attributed to the possibility of region-specific job compositions that are not evident at the aggregate level. However, it appears reasonable to differentiate region-specific studies from the more commonly observed 'Micro', 'Meso', and 'Macro' data that tend to have less variability in general.

Out of the regression specifics, only 'Trends' has a sufficiently high PIP. Along with the 'High-skill', which we will discuss later, 'Trends' is the only variable with a significant positive coefficient (0.392). The positive sign can be attributed to the observation of a negative trend on average among our data. Consequently, detrending would likely counteract this negative effect.

Another variable worth mentioning is the 'Journal impact' with a seemingly small coefficient (-0.004). It is essential to remind the reader that unlike most of the other explanatory variables that equal either zero or one, 'Journal impact' ranges from zero to 156. Therefore, the most impactful journal implies a negative coefficient of -0.624. That may be connected to publication bias, as it implies that studies have a higher chance of being in the top journal if they present a negative correlation with employment.

Developed countries suffer from technological advances more than developing in terms of employment. This finding is consistent with the research by de

Table 5.2: Bayesian model averaging results and robustness check

Response variable:	Bayesian model averaging			OLS		
Employment change	Post. mean	Post. SD	PIP	Coef.	SE	p-value
Constant	12.597	NA	**1.000**	0.367	0.076	0.000
Standard error	-0.386	0.092	**1.000**	-0.413	0.072	0.000
Methodology (OLS)						
Fixed-effects	0.000	0.021	0.010			
GMM	0.000	0.019	0.018			
IV	-0.411	0.141	**0.960**	-0.445	0.081	0.000
Other method	0.092	0.132	0.395			
Technology indicator (Robots)						
ICT	0.075	0.130	0.306			
Innovation	0.027	0.093	0.112			
Other indicators	-0.313	0.203	**0.768**	-0.424	0.100	0.000
Level of analysis (Meso)						
Macro	-0.003	0.030	0.026			
Micro	-0.008	0.041	0.053			
Regional	-0.606	0.094	**1.000**	-0.639	0.062	0.000
Regression specifics						
Dependent lag	-0.058	0.110	0.262			
Independent lag	0.053	0.106	0.236			
Trends	0.392	0.109	**0.995**	0.487	0.075	0.000
Time control	-0.031	0.066	0.219			
Study specifications						
Time horizon	-0.006	0.007	0.445			
Average year	-0.006	0.006	0.535			
Journal impact	-0.004	0.003	**0.835**	-0.002	0.001	0.018
Citations	0.000	0.000	0.186			
Country and labour characteristics						
Developed country	-0.235	0.158	**0.792**	-0.187	0.071	0.008
Manufacturing	-0.240	0.099	**0.929**	-0.239	0.069	0.001
Non-manufacturing	0.000	0.010	0.009			
High-skill	0.357	0.128	**0.957**	0.447	0.093	0.000
Medium-skill	-0.024	0.075	0.119			
Low-skill	-0.013	0.059	0.062			
Male	0.003	0.029	0.025			
Female	0.000	0.015	0.012			

Note: This table displays the results of Bayesian model averaging. In the OLS check we only include variables with PIP > 0.75. Post. mean = Posterior mean, Post. SD = Posterior Standard Deviation, PIP = Posterior inclusion probability, OLS = Ordinary Least Squares, Coef. = Coefficient, SE = Standard Error, GMM = Generalized Method of Moments, IV = Instrumental Variable, ICT = Information and Communication Technology

Vries *et al.* (2020) found out as our coefficient of 'Developed country' equals to (-0.235). Similarly, we observe a comparable effect (-0.240) with the variable 'Manufacturing' as this sector is more exposed to technological progress, mainly through the replacement of routine jobs with robots. The last important variable for the model estimation is the 'High-skill' control. As mentioned earlier, it is one of the two variables with a positive sign (0.357). As discussed in Chapter 2, highly educated workers have the advantage of better adapting to the changes, hence the positive coefficient we observe.

We can also gain insights from the variables that turn out to be insignificant. For instance, the gender variable ('Male' or 'Female') does not appear to have a significant impact on the likelihood of being affected by technological change. Additionally, the variable 'Average year' provides only weak evidence of its effect, suggesting that we do not need to be overly concerned about a significant rise in technology-driven unemployment in the future based on our data. Moreover, we observe negligible differences between 'Medium-skill' and 'Low-skill' workers, indicating that obtaining a university degree is the primary means to increase the chances of avoiding technological displacement, and the distinction between lower levels of education does not seem to matter.

With the information provided by the BMA, we can now proceed to conduct a simple OLS regression as an additional robustness check for our results. The coefficients, along with their corresponding standard errors and p-values, are presented on the right-hand side of Table 5.2. Upon examining the coefficients closely, we observe that they all exhibit the same sign and yield similar values to the posterior means obtained from the BMA. Furthermore, all variables remain highly significant, which provides strong support for their relationship with the outcome variable in our analysis.

Chapter 6

Conclusion

Our analysis aimed to gain a deeper understanding of the relationship between technological progress and unemployment. Since many researchers have already studied this effect and presented varying results, we conducted the first meta-analysis focusing on this topic. Our objective was to shed light on whether a consistent relationship exists and to identify the potential drivers that influence the impact of technology on employment.

We compiled a data set consisting of 516 observations from 43 different studies, upon which we conducted a series of statistical tests. Our findings reveal a relatively strong negative publication bias among all tests at the 1% significance level. Interestingly, we observe almost negligible (0.006) or statistically insignificant evidence of the hypothesized effect assuming exogeneity. When relaxing this assumption, the significance of the effect beyond bias remained at the 5% level while increasing considerably but still remaining close to zero.

By employing Bayesian Model Averaging (BMA) on our selected set of 28 variables, we were able to gain valuable insights into the sources of variation in the estimates. These variables covered a wide range of aspects, including methodological choices, technology indicators, levels of analysis, regression specifications, study characteristics, country attributes, and labour-related factors. Our BMA results show a significant positive correlation between employment and the variables 'Trends' and 'High-skill' labour. Conversely, the negative association with the underlying effect is observed in relation to the use of 'IV' regression, the inclusion of 'Other indicators', the utilization of 'Regional' data sets, higher 'Journal impact', studies conducted in a 'Developed country', and the presence of 'Manufacturing' as a control variable.

The findings presented in this thesis align with the existing literature to a large extent. Even though we found some negative publication bias, the majority of studies investigating technology-driven unemployment report estimates centred around zero. Similarly, our analysis did not yield significant evidence for the baseline effect. However, it is important to note that the impact of technological progress on employment likely manifests through shifts in the labour market, as indicated by our significant coefficients for 'High-skill' and 'Manufacturing' in the BMA. The overall conclusion is that the effects of job creation and job destruction tend to offset each other, resulting in a negligible net effect. That aligns with the perspective presented by Feldmann (2013), who suggests that while there may be a short-term impact, the long-term effect becomes insignificant.

It is important to acknowledge certain limitations associated with the findings presented in this thesis. Firstly, the literature search could have been enhanced by employing a Google Scholar search query and a snowballing method to ensure the inclusion of all relevant studies beyond those identified in the literature review. Additionally, conducting non-linear tests to explore potential publication bias considering the possibility of a non-linear relationship between the effect and its standard error, would provide further insights. Lastly, there are other methods, such as Frequentist Model Averaging (FMA) or best-practice estimation, that could have been applied to contribute to the topic in a more comprehensive manner but were not employed in this study.

Bibliography

Acemoglu, D. & P. Restrepo (2019): "Automation and new tasks: How technology displaces and reinstates labor." *Journal of Economic Perspectives* **33(2)**: pp. 3–30.

Acemoglu, D. & P. Restrepo (2020): "Robots and jobs: Evidence from us labor markets." *Journal of Political Economy* **128(6)**: pp. 2188–2244.

Addison, J. T., D. A. Fox, & C. J. Ruhm (2000): "Technology, trade sensitivity and labor displacement." *Southern Economic Journal* **66(3)**: pp. 682–699.

Arntz, M., T. Gregory, & U. Zierahn (2017): "Revisiting the risk of automation." *Economics Letters* **159**: pp. 157–160.

Atasoy, H. (2013): "The effects of broadband internet expansion on labor market outcomes." *ILR Review* **66(2)**: pp. 315–345.

Atasoy, H., R. D. Banker, & P. A. Pavlou (2016): "On the longitudinal effects of it use on firm-level employment." *Information Systems Research* **27(1)**: pp. 6–26.

Aubert-Tarby, C., O. R. Escobar, & T. Rayna (2018): "The impact of technological change on employment: The case of press digitisation." *Technological forecasting and social change* **128**: pp. 36–45.

Autor, D. & A. Salomons (2018): "Is automation labor share-displacing? productivity growth, employment, and the labor share." *Brookings Papers on Economic Activity* pp. 1–63.

Autor, D. H. & D. Dorn (2013): "The growth of low-skills service jobs and the polarization of the us labor market." *American Economic Review* **103(5)**: pp. 1553–97.

Autor, D. H., D. Dorn, & G. H. Hanson (2015): "Untangling trade and technology: Evidence from local labour market." *The Economic Journal* **125(584)**: pp. 621–646.

Autor, D. H., L. F. Katz, & A. B. Krueger (1998): "Computing inequality: Have computers changed the labor market?" *The Quarterly Journal of Economics* **113(4)**: pp. 1169–1213.

Bajzik, J., T. Havranek, Z. Irsova, & J. Novak (2023): "Does shareholder activism create value? a meta-analysis." .

Balsmeier, B. & M. Woerter (2019): "Is this time different? how digitalization influences job creation and destruction." *Research Policy* **48(8)**: p. 103765. The Digital Transformation of Innovation and Entrepreneurship.

Berman, E., J. Bound, & Z. Griliches (1994): "Changes in the demand for skilled labor within us manufacturing: evidence from the annual survey of manufactures." *The quarterly journal of economics* **109(2)**: pp. 367–397.

Biagi, F. & M. Falk (2017): "The impact of ict and e-commerce on employment in europe." *Journal of Policy Modeling* **39(1)**: pp. 1–18.

Blanas, S., G. Gancia, & S. Y. T. Lee (2019): "Who is afraid of machines?" *Economic Policy* **34(100)**: pp. 627–690.

Bogliacino, F. & M. Pianta (2010): "Innovation and employment: A reinvestigation using revised pavitt classes." *Research Policy* **39(6)**: pp. 799–809.

Borjas, G. J. & R. B. Freeman (2019): "From immigrants to robots: The changing locus of substitutes for workers." *RSF: The Russell Sage Foundation Journal of the Social Sciences* **5(5)**: pp. 22–42.

Breemersch, K., J. P. Damijan, & J. Konings (2019): "What drives labor market polarization in advanced countries? the role of china and technology." *Industrial and Corporate Change* **28(1)**: pp. 51–77.

Camina, E., A. Diaz-Chao, & J. Torrent-Sellens (2020): "Automation technologies: Long-term effects for spanish industrial firms." *Technological Forecasting and Social Change* **151**: p. 119828.

Capello, R. & C. Lenzi (2013): "Innovation and employment dynamics in european regions." *International Regional Science Review* **36(3)**: pp. 322–353.

CIRILLO, V. (2017): "Technology, employment and skills." *Economics of Innovation and New Technology* **26(8)**: pp. 734–754.

COMPAGNUCCI, F., A. GENTILI, E. VALENTINI, & M. GALLEGATI (2019): "Robotization and labour dislocation in the manufacturing sectors of oecd countries: a panel var approach." *Applied Economics* **51(57)**: pp. 6127–6138.

CUIJPERS, P., I. A. CRISTEA, E. KARYOTAKI, M. REIJNDERS, & M. J. HUIBERS (2016): "How effective are cognitive behavior therapies for major depression and anxiety disorders? a meta-analytic update of the evidence." *World psychiatry* **15(3)**: pp. 245–258.

CURRAN, T., A. P. HILL, P. R. APPLETON, R. J. VALLERAND, & M. STANDAGE (2015): "The psychology of passion: A meta-analytical review of a decade of research on intrapersonal outcomes." *Motivation and Emotion* **39**: pp. 631–655.

DACHS, B., M. HUD, C. KĎŻ˝HLER, & B. PETERS (2016): "Innovation, creative destruction and structural change: Firm-level evidence from european countries." *DEM Discussion Paper Series 16-26*, Department of Economics at the University of Luxembourg.

DE VRIES, G. J., E. GENTILE, S. MIROUDOT, & K. M. WACKER (2020): "The rise of robots and the fall of routine jobs." *Labour Economics* **66**: p. 101885.

DEKLE, R. (2020): "Robots and industrial labor: Evidence from japan." *Journal of the Japanese and International Economies* **58**: p. 101108.

DONG, J.-Y., K. HE, P. WANG, & L.-Q. QIN (2011): "Dietary fiber intake and risk of breast cancer: a meta-analysis of prospective cohort studies." *The American journal of clinical nutrition* **94(3)**: pp. 900–905.

EGGER, M., G. D. SMITH, M. SCHNEIDER, & C. MINDER (1997): "Bias in meta-analysis detected by a simple, graphical test." *Bmj* **315(7109)**: pp. 629–634.

EVANGELISTA, R. & A. VEZZANI (2012): "The impact of technological and organizational innovations on employment in european firms." *Industrial and Corporate Change* **21(4)**: pp. 871–899.

FABER, M. (2020): "Robots and reshoring: Evidence from mexican labor markets." *Journal of International Economics* **127**: p. 103384.

Falk, M. (2015): "Employment effects of technological and organizational innovations: evidence based on linked firm-level data for austria." *Jahrbücher für Nationalökonomie und Statistik* **235(3)**: pp. 268–285.

Feldmann, H. (2013): "Technological unemployment in industrial countries." *Journal of Evolutionary Economics* **23**: pp. 1099–1126.

Flug, K. & Z. Hercowitz (2000): "Equipment investment and the relative demand for skilled labor: International evidence." *Review of Economic Dynamics* **3(3)**: pp. 461–485.

Fu, X. M., Q. Bao, H. Xie, & X. Fu (2021): "Diffusion of industrial robotics and inclusive growth: Labour market evidence from cross country data." *Journal of Business Research* **122**: pp. 670–684.

Fung, M. K. (2006): "Are labor-saving technologies lowering employment in the banking industry?" *Journal of Banking & Finance* **30(1)**: pp. 179–198.

Gardberg, M., F. Heyman, P.-J. Norbäck, & L. Persson (2020): "Digitization-based automation and occupational dynamics." *Economics Letters* **189**: p. 109032.

Gechert, S., T. Havranek, Z. Irsova, & D. Kolcunova (2022): "Measuring capital-labor substitution: The importance of method choices and publication bias." *Review of Economic Dynamics* **45**: pp. 55–82.

Gera, S., W. Gu, & Z. Lin (2001): "Technology and the demand for skills in canada: an industry-level analysis." *Canadian Journal of Economics/Revue canadienne d'économique* **34(1)**: pp. 132–148.

Goaied, M. & S. Sassi (2019): "The effect of ict adoption on labour demand: A cross-region comparison." *Papers in Regional Science* **98(1)**: pp. 3–16.

Havranek, T., R. Horvath, & A. Zeynalov (2016): "Natural resources and economic growth: A meta-analysis." *World Development* **88**: pp. 134–151.

Ho, C.-Y. (2008): "Investment-specific technological change and labor composition: Evidence from the us manufacturing." *Economics Letters* **99(3)**: pp. 526–529.

Hotte, K., M. Somers, & A. Theodorakopoulos (2022): "Technology and jobs: A systematic literature review." *Papers 2204.01296*, arXiv.org.

Idris, Z. Z., N. W. Ismail, S. Ibrahim, & H. Z. Hamzah (2021): "The impact of high-technology trade on employment." *The Journal of International Trade & Economic Development* **30(4)**: pp. 512–529.

Itani, O., M. Jike, N. Watanabe, & Y. Kaneita (2017): "Short sleep duration and health outcomes: a systematic review, meta-analysis, and meta-regression." *Sleep medicine* **32**: pp. 246–256.

Ivus, O. & M. Boland (2015): "The employment and wage impact of broadband deployment in canada." *Canadian Journal of Economics/Revue canadienne d'économique* **48(5)**: pp. 1803–1830.

Jerbashian, V. (2019): "Automation and job polarization: On the decline of middling occupations in europe." *Oxford Bulletin of Economics and Statistics* **81(5)**: pp. 1095–1116.

Jung, J. H. & D.-G. Lim (2020): "Industrial robots, employment growth, and labor cost: A simultaneous equation analysis." *Technological Forecasting and Social Change* **159**: p. 120202.

Kass, R. E. & A. E. Raftery (1995): "Bayes factors." *Journal of the american statistical association* **90(430)**: pp. 773–795.

Leigh, N. G., B. Kraft, & H. Lee (2020): "Robots, skill demand and manufacturing in us regional labour markets." *Cambridge Journal of Regions, Economy and Society* **13(1)**: pp. 77–97.

Maurin, E. & D. Thesmar (2004): "Changes in the functional structure of firms and the demand for skill." *Journal of labor economics* **22(3)**: pp. 639–664.

Morrison Paul, C. J. & D. S. Siegel (2001): "The impacts of technology, trade and outsourcing on employment and labor composition." *Scandinavian Journal of Economics* **103(2)**: pp. 241–264.

Nordmann, A. (2012): "Meta-analyses: what they can and cannot do." *Swiss medical weekly* .

Pantea, S., A. Sabadash, & F. Biagi (2017): "Are ict displacing workers in the short run? evidence from seven european countries." *Information Economics and Policy* **39**: pp. 36–44.

Pellegrino, G., M. Piva, & M. Vivarelli (2019): "Beyond r&d: the role of embodied technological change in affecting employment." *Journal of Evolutionary Economics* **29**: pp. 1151–1171.

Piva, M. & M. Vivarelli (2004): "Technological change and employment: some micro evidence from italy." *Applied Economics Letters* **11(6)**: pp. 373–376.

Piva, M. & M. Vivarelli (2005): "Innovation and employment: Evidence from italian microdata." *Journal of economics* pp. 65–83.

Polák, P. (2017): "The productivity paradox: A meta-analysis." *Information Economics and Policy* **38**: pp. 38–54.

Solow, R. M. (1956): "A contribution to the theory of economic growth." *The Quarterly Journal of Economics* **70(1)**: pp. 65–94.

Stanley, T. D. (2005): "Beyond publication bias." *Journal of economic surveys* **19(3)**: pp. 309–345.

Stanley, T. D. (2008): "Meta-regression methods for detecting and estimating empirical effects in the presence of publication selection." *Oxford Bulletin of Economics and statistics* **70(1)**: pp. 103–127.

Stanley, T. D. & H. Doucouliagos (2017): "Neither fixed nor random: weighted least squares meta-regression." *Research synthesis methods* **8(1)**: pp. 19–42.

Van Reenen, J. (1997): "Employment and technological innovation: evidence from uk manufacturing firms." *Journal of labor economics* **15(2)**: pp. 255–284.

Van Roy, V., D. Vértesy, & M. Vivarelli (2018): "Technology and employment: Mass unemployment or job creation? empirical evidence from european patenting firms." *Research policy* **47(9)**: pp. 1762–1776.

Vivarelli, M., R. Evangelista, & M. Pianta (1996): "Innovation and employment in italian manufacturing industry." *Research policy* **7(25)**: pp. 1013–1026.

Xiang, C. (2005): "New goods and the relative demand for skilled labor." *Review of Economics and Statistics* **87(2)**: pp. 285–298.

Appendix A

List of studies

Table A.1: List of studies

Acemoglu & Restrepo (2019)	Falk (2015)
Acemoglu & Restrepo (2020)	Flug & Hercowitz (2000)
Atasoy (2013)	Fu *et al.* (2021)
Atasoy *et al.* (2016)	Fung (2006)
Autor & Salomons (2018)	Gardberg *et al.* (2020)
Aubert-Tarby *et al.* (2018)	Gera *et al.* (2001)
Autor *et al.* (1998)	Goaied & Sassi (2019)
Balsmeier & Woerter (2019)	Ho (2008)
Berman *et al.* (1994)	Idris *et al.* (2021)
Biagi & Falk (2017)	Ivus & Boland (2015)
Blanas *et al.* (2019)	Jerbashian (2019)
Bogliacino & Pianta (2010)	Jung & Lim (2020)
Borjas & Freeman (2019)	Leigh *et al.* (2020)
Breemersch *et al.* (2019)	Morrison Paul & Siegel (2001)
Camina *et al.* (2020)	Pantea *et al.* (2017)
Capello & Lenzi (2013)	Pellegrino *et al.* (2019)
Cirillo (2017)	Piva & Vivarelli (2004)
Compagnucci *et al.* (2019)	Piva & Vivarelli (2005)
Dachs *et al.* (2016)	Van Reenen (1997)
de Vries *et al.* (2020)	Van Roy *et al.* (2018)
Evangelista & Vezzani (2012)	Vivarelli *et al.* (1996)
Faber (2020)	

Appendix B

Additional figures

Figure B.1: Correlation table for Bayesian model averaging

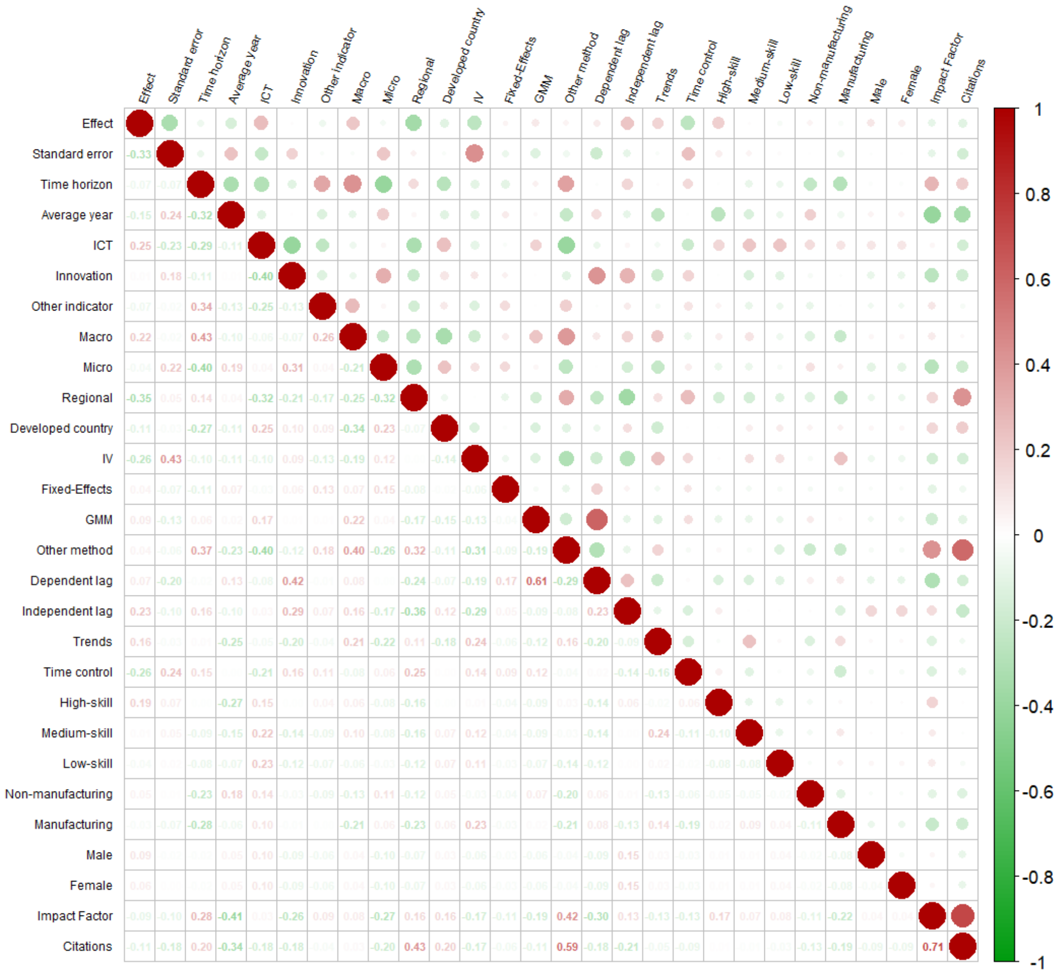

Note: This table shows the correlation table for the Bayesian model averaging with unit information g-prior and dilution model prior specification. The BMA results are presented and discussed in Chapter 5. Green colour represents negative correlation. Red colour represents positive correlation.