MASTERARBEIT | MASTER'S THESIS

Titel | Title

AI and Unemployment

verfasst von | submitted by

Teodor Yonchev

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 913
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Applied Economics
Betreut von | Supervisor:	Univ.-Prof. Dipl.-Math. Karl Schlag PhD

Abstract (EN)

This research aims to estimate the empirical effect that AI development has on the job market. The first model is focused on determining this effect on unemployment, while a deeper dive into different sectors does so by looking at employment. The panel data covers the years 2001 to 2020 and features 21 nations that are at the forefront of technology. A two-step system GMM model is employed to show that AI development is associated with a negative effect on unemployment, thus creating new jobs. In contrast, the sector-specific model has low statistical power in explaining the researched dynamic in specific sectors.

Zusammenfassung (DE)

Diese Forschung zielt darauf ab, den empirischen Effekt der KI-Entwicklung auf den Arbeitsmarkt zu schätzen. Das erste Modell konzentriert sich darauf, diesen Effekt auf die Arbeitslosigkeit zu bestimmen, während eine detailliertere Analyse verschiedener Sektoren dies anhand der Beschäftigung untersucht. Die Paneldaten erstrecken sich über den Zeitraum 2001–2020 und umfassen 21 der technologisch fortschrittlichsten Länder. Ein zweistufiges System-GMM-Modell wird verwendet, um zu zeigen, dass die KI-Entwicklung mit einem negativen Effekt auf die Arbeitslosigkeit verbunden ist und somit neue Arbeitsplätze schafft, während das sektorspezifische Modell eine geringe statistische Aussagekraft in Bezug auf die untersuchte Dynamik in einzelnen Sektoren aufweist.

Contents

1 Introduction

Fears of technological unemployment are nothing new. Automation, where workers get replaced by machines in performing certain tasks, began already with the industrial revolution. In the 18th century, steam-powered machines, as well as spinning and weaving machines, began replacing skilled artisans. Mechanized reapers and the invention of the cotton gin replaced manual labor in farming. The latest wave of automation was computerized automation in the 1970s and 1980s, which replaced humans at repetitive tasks (Acemoglu and Restrepo, 2018). However, these were all forms of *creative destruction,* a term first coined by the Austrian economist Joseph Schumpeter, meaning that some jobs were destroyed due to automation, but the increased productivity allowed for new jobs to be created, where people had a *competitive advantage* over machines. The fears of technological unemployment remained mostly unjustified because machines were very good at performing a very narrow task, while humans were very good at performing multiple tasks in various fields of expertise. However, what will happen to the labor markets if we create a technology that imitates human intelligence?

Research so far on the topic has been trying to figure out whether AI technologies are associated with a *displacement* or a *reinstatement effect*. *Displacement effect* is when automation enables capital to replace labor and is directly associated with a net negative effect on labor markets, while the *reinstatement effect* refers to the situation where new jobs are created, directly offsetting the displacement effect. Korinek and Stiglitz (2018), for example, argue that unemployment might arise in circumstances where jobs get destroyed faster than workers need in order to reeducate themselves for another occupation. Furthermore, if technological progress results in a lower marginal product of workers, the market-clearing wage might go below the cost of living of workers and result in unemployment.

Through this empirical research, I aim to contribute to the discourse on this subject by addressing the question: "What are the effects of AI development on unemployment?" This will be achieved by examining 21 of the most technologically advanced nations and analyzing the impact of AI on unemployment during the period from 2001 to 2020. The proxy that is used for AI development is the global count of AI-related patents, which is being regressed on unemployment using a two-step system General Method of Moments (GMM) regression. The results show that the development of AI leads to a slight reduction in unemployment. Moreover, by using the global AI patent count as a proxy, a much smaller negative effect on unemployment was observed than previous studies have shown. While it is interesting to see what the aggregate effect is, there is no doubt that certain sectors, as well as occupations, will be affected heterogeneously. Therefore, the research is being expanded by breaking down the dependent variable by sector. The results from the second equation seem to be statistically insignificant, and no meaningful interpretation can be derived from them.

2 Literature review

One of the main puzzles that economists have been trying to solve is whether AI will increase, decrease, or leave unemployment unchanged. Acemoglu and Restrepo (2019) empirically conclude that the recent stagnation in labor demand is due to the *displacement effect* where automation prevails over the creation of new tasks, and if automation continues to be our main source of productivity growth, it is inevitable for the labor share in the production process to continue declining. Furthermore, Acemoglu (2025) show empirically that the TFP effects due to AI in the next 10 years will be modest, with an increase of 0.71% or 0.07% annually, while GDP could be positively affected by a 0.9%-1.1% increase in the span of 10 years. Babina et al. (2024), on the other hand, developed a new method of measuring AI investments by looking at the amount of AI-skilled labor in firms. They find that firms that invest in the new technology experience an increase in sales, firm size and employment.

Another string of literature looking at the exposure levels of jobs to AI also contributes to the field. Webb (2019), for example, creates a noble task-based measure of AI exposure by pairing verb-noun pairs of patents to job descriptions. He finds that, unlike previous automation technologies, AI aims at replacing high-skilled jobs, and it endangers mostly older and highly educated individuals. Similarly, Felten et al. (2018) and Felten et al. (2019) developed a different, ability-based measure of AI exposure. They conclude that jobs with previously higher automation levels observe employment growth. Albanesi et al. (2025) take these two measures and study the employment effects on the European job market of AI depending on the exposure levels. Interestingly, they find that younger people, as well as high-skilled workers, seem to benefit from higher exposure levels to AI, meaning that moving from lower exposure to higher exposure seems to affect employment positively. My master's thesis takes a different approach to studying the heterogeneous effect of AI on employment by directly analyzing the employment effects of AI patent count on three major sectors for the period 2001-2020.

Looking at the purely empirical part of the literature, Guliyev (2023) and Guliyev et al. (2023) studied the direct effect of AI on unemployment using the Google Trend Index (GTI) as their main explanatory variable. Although their primary econometric approach is solid, I will argue in my research that utilizing AI patents could serve as a superior proxy for the development of AI technologies, since additional Google searches on the topic of AI do not directly result in a technological advancement in the field. In another study, Bordot (2022) used AI patent count as a proxy for AI development. He broke down his regression by different levels of education and got statistically significant results, showing that people at the lower education levels benefit from AI in terms of employment. Mutascu (2021) also implements AI patents as a proxy in their empirical study. Furthermore, the authors argue that the relationship between AI development and unemployment is not linear and use inflation as a determinant for this nonlinearity and conclude that under a low inflation environment, AI development reduces unemployment.

What these studies do in practical terms is look at the AI patent counts in a certain country for a certain period and regress them against the respective unemployment values for the same country and period. However, this approach might have some limitations since it assumes that only the patents released in a country influence the unemployment of that particular country. AI is a global technology and is used across borders, especially with the ever-elevating globalization and interconnectedness of our economies. It seems more realistic to think that a country's unemployment levels are influenced by the global technological advancements of AI and not only by that which is restricted to the borders of the specific country. To this end, my biggest contribution with this paper so far is that I construct a global AI patent count and regress this on each country's unemployment levels for the studied period. Another interesting contribution my research makes to the literature is that I use government expenditure on education as a control variable, thus suggesting that depending on how well a country's education system is developed might influence its unemployment levels, and ignoring this might introduce an omitted variable bias.

3 Empirical approach

In this section of my master's thesis, I will discuss several key aspects of my empirical approach to studying the effects of AI development on unemployment. I will begin by providing a comprehensive overview of the study's purpose and scope. The model that will be employed is a General Method of Moments (GMM).

3.1 Methodology

Since we are studying the effects on unemployment, one ought to be very careful, because it is a very persistent variable. This means it does not immediately adjust to economic shocks, and it takes time for everything to be factored into the data. Therefore, it is expected that unemployment from the previous period is directly impacting unemployment in the current period. Thus, it exhibits a "memory" effect, meaning that past levels are strongly correlated with current ones, proving the rigidity and dynamic nature of this variable. Things like hiring frictions, wage stickiness, and structural unemployment contribute to this phenomenon. Government policies or technological shocks and other drivers of the economy take time until they are felt in the labor market. Therefore, a dynamic panel data (DPD) model is highly suitable for our purposes (Guliyev, 2023). However, including the lag of the dependent variable on the right-hand side comes with some additional challenges since it introduces endogeneity into the model because the lag of unemployment will be correlated with the error term.

Usually, any kind of empirical research would start off from a simple Ordinary Least Squares (OLS) regression, however, a main assumption of the OLS is that the error term is not correlated with any of the independent variables. Furthermore, in panel data errors seem to be correlated with each other over time. Especially when we are studying the same countries over time, there is good reason to think that any unobserved random factors that influence the dependent variable will continue to do so also in the future, which leads to autocorrelation. Therefore, including the lag of the dependent variable on the right-hand side automatically introduces endogeneity into the model since any unobserved random factors that influence the dependent variable in the current period probably also influenced the dependent variable in the previous period and thus violating the main assumption of the OLS that the error term is not correlated with any of the independent variables. Therefore, using a more sophisticated technique, like a General Method of Moments model proposed by Arellano and Bond (1991), is necessary.

The method developed by Arellano and Bond (1991) is called a difference GMM. This regression type is very good at taking care of endogeneity concerns. The first step of the difference GMM is to take the first difference of the level equation, thus getting rid of any country-specific fixed effects. After that, we use the lagged levels of the explanatory variables as instruments for the difference equation in order to get rid of the endogeneity. However, the difference equation might suffer from weak instruments, especially when the endogenous variable is highly persistent, as is the case with unemployment. Arellano and Bover (1995) and Blundell and Bond (1998) came up with a solution to this problem. This improved GMM is called a system GMM. In the system GMM we perform a second step where we estimate the level equation, using the lagged differences of the endogenous variables as instruments. Finally, we create a system of the two equations and estimate them jointly, thus correcting for the issue of the weak instruments in the initial difference GMM proposed by Arellano and Bond (1991).

Moreover, there are two kinds of system GMM – a one-step and a two-step system GMM. The difference is that the one-step system GMM assumes that errors have a constant variance. This approach provides consistent estimates, but it is less efficient and is much more suitable for a small number of countries and a longer period of time. The two-step system GMM, on the other hand, uses a weight matrix that accounts for heteroscedasticity, thus improving the efficiency and accuracy of standard errors. The two-step approach is considered much better for a higher number of countries and a shorter time period (Hoeffler et al., 2001). Therefore, since this research studies 21 countries over the period of 19 years, it is more appropriate to use the two-step system GMM.

To conclude that the regression is estimated correctly, two essential tests must be performed: a second-order autocorrelation test and a test for the validity of the instrument.

Firstly, a Hansen J-Test is applied after the regression results are obtained. The Hansen J-test aims to determine whether the instruments of the GMM are valid, meaning that they are uncorrelated with the error term. This test is formally represented by Equation (1).

$$J = g_n(b_n)' S_w^{-1} g_n(b_n) \quad (1)$$

Where $g_n(b_n)$ is the vector sample moment conditions, S_w is the covariance matrix of the moment conditions. Furthermore, the test follows a chi-squared distribution with the degrees of freedom equal to the number of instruments minus the number of the parameters that are being estimated.

The second test that is being applied is an AR(2) test, which checks for second-order autocorrelation. The formal equation as specified by Arellano and Bond (1991) is shown in Equation (2).

$$M_2 = \frac{A}{V^{1/2}} \sim N(0,1) \quad (2)$$

Where M_2 is the test statistic for the second-order serial correlation, A is the sample moment condition based on the second-lagged residuals, V is the asymptotic variance of the sample moment conditions and $N(0,1)$ shows that the test follows a normal distribution with a mean of 0 and a variance of 1.

3.2 Dataset and variables

This research is done using panel data from 21 high-tech and developed countries (i.e. Austria, Australia, Canada, Switzerland, Germany, Denmark, Estonia, Finland, France, United Kingdom, Ireland, Israel, Iceland, Japan, South Korea, Luxembourg, Netherlands, Norway, New Zealand, Sweden, United States of America) over a period of 19 years (i.e., 2001-2020). AI technologies have received a lot of attention since the beginning of the 21st century, particularly in the last few years. One limitation to using patents as a proxy, though, is that the most recent data is until 2020, which is also why this research extends until then. Ultimately, it would have been best if the recent boom in the technology could also be included, which was caused by the release of ChatGPT on November 30, 2022. However, datasets on patents take time to get updated, and the data availability is simply limited. Countries are picked based on data availability and their high ranking in the Global Innovation 2024 index (Dutta et al., 2024). The logic behind picking the most technologically advanced countries is that one would expect them to be most active in creating and implementing AI technologies.

Data is obtained from several different data sources. These include *the UNESCO Institute for Statistics (UIS)* and *the OECD database*. Sources are linked appropriately in the Bibliography section. All the data for all the variables is in annual terms. Data for unemployment is lacking for Switzerland for the period 2001-2009, for Iceland for the period 2001-2002, and for Luxembourg for the period 2001-2002. All other data points on unemployment for the other countries are available. Furthermore, most of the data is taken from *the OECD database* with the exception of the dataset on government spending on education, which was taken from *the UNESCO database*. The annual data for unemployment is calculated using averages of infra-annual estimates.

The ***dependent variable*** is unemployment, measured as the percentage of people who are part of the total labor force that are unemployed and actively looking for a job. In a second regression, global AI patents will be regressed against employment in order to study the heterogeneous effects on different sectors. The reason for using employment this time is that categorizing a certain worker by a sector poses significant challenges when he is not employed. Theoretically, one could categorize them by the previous sector that they were in. However, data for this is still not available; therefore, using unemployment on the left-hand side is simply not feasible. Nevertheless, this minor adjustment poses little concern, as it essentially represents the converse aspect of the same concept.

Unlike unemployment, employment is a stock variable in my dataset, which changes the econometric approach a little bit. Different countries have different population sizes, and one needs to control for that. Otherwise, the effect that the population size has on unemployment will be assigned to the AI variable and will give an upward bias to the results. Therefore, all the variables were transformed into logarithmic terms. Figure 1, available in the appendix, gives us some additional descriptive statistics on the behavior of the variable. Four big economies were chosen to visually describe how this variable evolves over the studied time period; these include two from Europe, one from Asia, and one from North America.

The ***interest variable*** is global AI-related patents for the period 2001-2020. Patents are a good proxy for measuring the development of any technology. One could argue that R&D would be a good proxy for AI, however, there doesn't exist a macroeconomic dataset on R&D that is especially devoted to AI. Moreover, R&D is a measure of input into the technology and not output, meaning that an extra dollar invested in AI does not directly result in a better and more developed technology, while patents represent exactly improvements in the field (Bordot, 2022). An unavoidable limitation of patents is the data availability mentioned earlier. However, patents still remain the best option compared to other proxies proposed by the literature. One of the main novelties presented by this research is that I will construct a global AI patent count by summing up all the patents that were released by all the countries in a particular year. Thus, regressing the global AI development on a country's specific unemployment rate, as motivated earlier in the literature review. There is some limitation to this approach as I am allowing the AI variable to vary only in between years but not across countries, thus assuming that all countries are affected equally by the technology. While this is still far from the truth, my argument is that it is much closer to reality than assuming that only the technology developed in a certain country affects the unemployment of that particular country. These technologies are a global phenomenon, and it would be naïve to think that their use is only restricted to the countries in which they were created.

When constructing the AI patent variable, there were a few key decisions that had to be made. There are three major patent offices in the world. The European one is called "European Patent Office", the US one is called "United States Patent and Trademark Office", and the third one is "World Intellectual Property Organization. In order to not lose any observations, all the patents by all the patent offices were summed up, because a certain technology may be protected only by one office. However, now there is the potential of double-counting patents, because the same patent could be issued, for example, ones under the US patent office and ones under the European patent office. To tackle this potential issue, the dataset was filtered by patent *priority date*. The *priority date* is the first time a technology has been filed for patenting and is usually used when disputes over a technology arise. For example, if two different entities claim the rights over the same technology, the legislators look at who has the earliest priority date, meaning who claimed first rights over the innovation. There can be only one single *priority date* for a single technological innovation, breakthrough, or small improvement. Thus, all patents are retained in the dataset, and no double-counting occurs. Moreover, a min-max normalization was applied to the interest variable for easier readability of the results. The reason for this choice is that a single increase in the patent count would result in a minimal change in the unemployment variable. By using a min-max transformation, one is studying the effect that is observed by going from the minimum to the maximum value of AI patents in the dataset. Thus, the results afterward can be read by dividing the coefficient, for example, by four and seeing what effect a 25% increase in the patents has on unemployment.

Some additional descriptive statistics on the AI variable are presented in Figure 2. The increased attention and popularity that this technology has seen recently is clearly visible. The AI patent count has an upward trend up until 2019 and slightly descends afterward.

An additional figure is included in the appendix to visualize the development of unemployment and the AI patents released in the USA. Figure 4 shows this development for USA, since it is considered the country that sits at the technological frontier in this field and thus provides some intuition on what to expect from the results that will be produced later.

The ***control variables*** are the lag of unemployment, again measured as the percentage of people that are part of the labor force and are unemployed but actively looking for a job, inflation, GDP growth and government spending on education. The use of the lag of unemployment is necessary since this is a very persistent variable and the previous period's value has some considerable explanation power for the current period's value, as motivated earlier.

Inflation is considered also a very important determinant of unemployment, as discussed in the theoretical literature (Phillips, 1958). Phillips (1958) argued that under low unemployment levels, employers compete for fewer workers, thus are incentivized to offer higher wages. Higher wages result in higher costs for businesses, which they pass onto the consumer. Workers on the other hand also have higher purchasing power due to the increased wages, thus also driving up prices. Therefore, the famous Phillips curve entails that unemployment and inflation are inversely related. The inflation variable is in annual terms and is calculated using the Consumer Price Index (CPI).

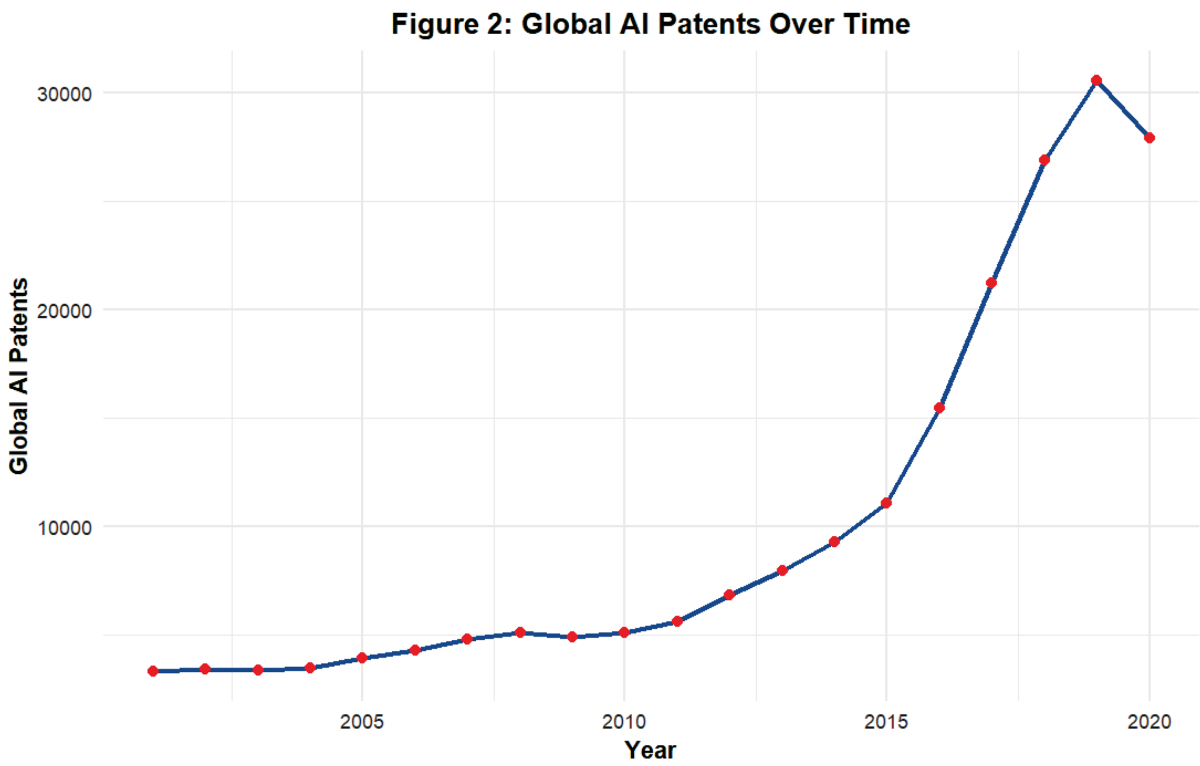

A second very important control variable is GDP growth. Economic growth and unemployment are two very important economic metrics followed by both the public and policymakers. The dynamic and link between these two variables have been widely studied by both theoretical and empirical research. This very prominent macroeconomic relationship is often referred to in the literature as the famous "Okun's law". Okun (1963) argued that when unemployment increases, actual GDP tends to fall below its potential value. Moreover, he came to the conclusion that this relationship is not linear, meaning that a small increase in unemployment led to a larger drop in GDP. The economic reasoning is straightforward: when employment levels are lower, fewer goods are produced; furthermore, higher unemployment levels mean lower overall demand for goods by consumers. Thus, it is expected that the results in this research will indicate towards an inverse relationship between the two variables. Some additional descriptive statistics are presented in Figure 3, available in the appendix. Four big economies were chosen to visually describe the dynamic of this variable: two from Europe, one from Asia, and one from North America.

Government spending on education, expressed in PPP (millions), is another variable included in the regression. The dynamic between education and unemployment has been widely studied by the previous literature, and empirical studies so far have shown an inverse relationship between them (Riddell et al., 2011; Ashenfelter et al., 1979). Onuoha et al. (2019), using a two-step GMM regression, show, for example, that expenditure on education reduces the unemployment rate; thus, omitting this variable might introduce a bias into the regression. The logic is simple: governments that invest more in their education systems and thus have a higher overall education quality could benefit from lower unemployment levels. Therefore, an inverse relationship between the variables is expected. This small addition is also considered a novelty to this research since previous studies on the labor market effects of AI have not considered the effects of education in their regressions.

3.3 Econometric model and results

This study aims to assess the effect that the development of AI technologies has on unemployment in 21 high-tech developed countries. It does so by running a two-step system GMM regression estimating the following dynamic panel data model:

$$UN_{it} = \alpha_i + \beta_1 UN_{it-1} + \beta_2 AIG_t + \beta_3 IR_{it} + \beta_4 GDPGW_{it} + GSE_{it} + \epsilon_{it} \quad (3)$$

Where UN_{it} is the unemployment rate, UN_{it-1} is the unemployment rate from the previous year. AIG_t is the global AI patent count, IR_{it} is the inflation rate. $GDPGW_{it}$ is the GDP growth rate and GSE_{it} is government spending on education. Lastly, α_i is the country-specific fixed effects that get eliminated in the first step of the GMM estimation, where first-differencing is applied and ϵ_{it} is the error term. Furthermore, the countries picked for this study are denoted by (i), and the years between 2001 and 2020 are denoted as (t). The output of Equation 3 is represented by Table 1. A two-step system GMM is implemented. As mentioned earlier, the GMM uses the lags of the endogenous variables as instruments. In this research, the first and second lag are used for all the endogenous variables, and the second and third lag are used for the previous period unemployment UN_{it-1}, which is an explanatory variable included on the right-hand side.

At the bottom of Table 1 are the results from the diagnostic tests. A Hansen J-Test developed by Hansen (1982) is applied in order to test the validity of overidentifying restrictions. An instrument is considered valid if it satisfies the *Relevance Condition* and the *Exogeneity Condition.* Meaning that the instrument must be correlated with the endogenous variable and uncorrelated with the error term. The Null Hypothesis (H_0) of the Hansen J-Test is that all instruments are valid. We accept H_0 in case the *p-value* is high, for example, above 0.10. If the *p-value* is low, for example, below 0.05, we need to reject the null hypothesis and accept that at least one of the instruments is invalid. From Table 1, we can see that the *p-value* of the performed Hansen J-Test is 0.6215. Thus, we accept our H_0 and conclude that all instruments used in Equation (3) are valid.

When the regression model is being estimated, the standard errors assume that observations are independent. However, this may not hold for the way this empirical research is being conducted, since unemployment levels in the dataset may be correlated across years within

Table 1: Regression Results with Diagnostic Tests

	Dependent variable:
	Unemployment Rate
Unemployment Rate (Lagged)	0.8877***
	(0.0334)
Global AI Patent	-0.7517***
	(0.2426)
Inflation Rate	-0.0807
	(0.0556)
GDP Growth Rate	-0.3102***
	(0.0484)
Gov. Spending on Education	-0.000001
	(0.000001)
Constant	1.6358***
	(0.3354)
AR(2) Test	0.4631
Hansen J-Test	0.6215
Number of Observations	283
Note:	*p<0.1; **p<0.05; ***p<0.01

the same country. Thus, standard errors may be too small and we may observe an incorrect statistical significance in the results. This problem is solved by clustering standard errors and thus correcting for the within-country correlation. By doing so, we allow for errors to be independent across countries but correlated within countries. Clustering will produce more accurate and larger standard errors. That is also why a Hansen J-Test better suits this research than a Sargan test. The Sargan test assumes constant variance across observations, while the Hansen J-Test allows for heteroscedasticity and autocorrelation.

Furthermore, a second-order autocorrelation Arellano–Bond's (1991) test is performed. In a GMM regression type, first-order autocorrelation is normal and does not introduce any problems into the results. However, second-order autocorrelation should not be present. The Null Hypothesis (H_0) of the AR(2) test is that there is no second-order autocorrelation. If the *p-value* of the AR(2) test is above 0.05, we accept the H_0 and conclude that there is no second-order autocorrelation. If the *p-value* is below 0.05, we reject H_0 and conclude that there is a problem in the regression that we need to resolve. As shown in Table 1, the AR(2) test has a *p-value* of

0.4631, which is well above the 0.05 threshold. Thus, we accept the H_0 and conclude that there is no second-order autocorrelation detected in our model.

The results presented in Table 1 show that the lag of unemployment, global AI patents as well as GDP growth appear to be all statistically significant at the 1% level. The results indicate that the logic behind the dynamic nature of the unemployment variable was correct, thus including it on the right-hand side as a regressor proved to be reasonable, as argued by previous literature (Guliyev, 2023; Bordot, 2022; Damioli et al., 2023). Therefore, the regression results imply that a 1% increase in the previous period's unemployment results in a 0.8877% increase in the current year's unemployment rate. Next, we observe that GDP growth inversely affects unemployment rates as argued by Okun (1963). This means that good economic conditions have a positive effect on labor markets, or more specifically, a 1% increase in GDP results in a 0.3102% decrease in unemployment. Unfortunately, no meaningful conclusions can be made on the effect that inflation and government spending on education have on unemployment rates since the results are statistically insignificant.

Lastly, the main variable of interest, Global AI Patents, is highly significant and has an inverse relationship with the dependent variable. Given that the AI variable has undergone min-max normalization, the effect observed in our regression output reflects the anticipated outcome of transitioning from the minimum to the maximum number of patents. Consequently, dividing the result by four yields the effect we would expect from a 25% increase in the number of patents. Thus, a 25% increase in patents is associated with a 0.1879% decrease in the unemployment rate. Although this may not sound like a big effect, the results are statistically significant at the 1% level, and we conclude that the development of AI technologies results in a decrease in the overall unemployment rate.

3.4 The heterogeneous effect of AI on different employment sectors

Nevertheless, one needs to keep in mind that the results, observed in Table 1, are on the aggregate unemployment level. However, the effects may vary, depending on the nature of different sectors. One may argue that some sectors benefit greatly from AI, while others experience detriment. The preceding section concludes that aggregating the effects from various sectors yields a net positive impact on labor markets. Thus, the previous section contributes to the literature by confirming a *reinstatement effect* of AI technologies on labor markets. However, the goal of this consecutive regression is to study the heterogeneous effect that AI has on different sectors.

Three main sectors are chosen to be further studied in depth. These include agriculture, forestry and fishery, manufacturing, and services. At first, the regression was done by splitting up the dataset into three and then doing a GMM regression for each sector separately. This method kept things simple, since the same regression specification from the previous part was used, with the only difference that on the left-hand side was the employment level of a particular sector.

Nevertheless, this led to losing a lot of observations, since the dataset is being split into three. Therefore, another method was chosen for the sector-specific regression, where a dummy variable approach was applied. Moreover, as mentioned earlier, employment by sector is being used, which is a stock variable. However, there is a potential problem with using a stock variable on the left-hand side since employment levels in each country will be influenced by the different population sizes. Thus, one needs to control for that. There are various ways to do that; however, in this research, a very simple and straightforward approach will be applied, and that is to log all the variables. By logging all the variables, the regression essentially measures percentage changes rather than absolute levels. This adjustment implicitly accounts for the differences in population size. Moreover, doing this will change the way the regression results are being read, since now the coefficients will represent the elasticity of employment with regards to AI, meaning that if AI patents increase by 1%, what will be the percentage change in employment?

Equation (4) is estimated in order to investigate the heterogeneous effects of AI on different sectors.

$$EMP_{it} = \alpha_i + \beta_1 EMP_{it-1} + \beta_2 AIG_t + \beta_3 IR_{it} + \beta_4 GDPGW_{it} + \beta_5 GSE_{it} + \beta_6 DMAN + \beta_7 DSER + \beta_8 (DMAN \times AIG_t) + \beta_9 (DSER \times AIG_i) + \epsilon_{it} \quad (4)$$

Here EMP_{it} is the employment level, AIG_t is the aggregate AI patent count, IR_{it} is the inflation rate, $GDPGW_{it}$ is the GDP growth, GSE_{it} is the government spending on education, $DMAN$ is the dummy for the manufacturing sector, $DSER$ is the dummy for the service sector, $DMAN \times AIG_t$ represents the interaction term for the manufacturing sector and $DSER \times AIG_i$ is the interaction term for the service sector. As mentioned earlier a dummy variable approach is used in order to retain as many observations as possible to keep the statistical power high.

In Equation (4) $DMAN = 1$ when an observation is from the manufacturing sector and 0 otherwise. Following the same logic $DSER = 1$ when an observation is from the service sector and 0 otherwise. There is no dummy included for the agriculture, forestry, and fishery sector since this one is used as the reference category. The reason for choosing this sector as the reference category is that the used dataset has the most observations for it. Lastly, when both, $DMAN = 0$ and $DSER = 0$ the specific observation is from the agriculture, forestry and fishery sector.

A dummy variable regression always gives us the results in comparison to the effect on the reference category. Meaning that the β_2 coefficient indicates the effect of AI on the agriculture, forestry, and fishery sector. Then β_6 and β_7 gives us the baseline difference in employment levels in manufacturing and in services accordingly. Meaning that the β_6 indicates how much higher or lower the employment levels in manufacturing are compared to agriculture, forestry, and fishery, while the β_7 indicates the same but for the service sector. Including the dummies in the regression is important since otherwise, the model will assume that the employment levels start at the same baseline level across all sectors, which is simply not realistic because different sectors naturally have different employment levels. This way, the model accounts for these differences and the results are kept unbiased.

Lastly, the interaction terms indicate the effect that AI has on employment levels in the manufacturing and the service sector compared to the agriculture, forestry and fishery sector. The coefficient that the regression produces for β_8 will indicate how much more or how much less manufacturing is affected compared to the reference category and β_9 will indicate the same for the service sector. Therefore, β_2 indicates the effect just on the agriculture sector, while $\beta_2 + \beta_8$ is the effect for the manufacturing sector and $\beta_2 + \beta_9$ is the effect for the service sector.

The results from the second regression are to be found in Table 2.

Table 2: GMM Regression Results with Diagnostic Tests

	Dependent variable:
	Employment
Employment (Lagged)	0.99969***
	(0.02630)
Global AI Patents	-0.02708
	(0.11832)
Inflation Rate	0.74814***
	(0.17558)
GDP Growth Rate	0.55729***
	(0.08463)
Gov. Spending on Education	0.00383
	(0.02541)
Manufacturing Sector	-0.34941
	(1.28816)
Services Sector	-0.55701
	(2.42616)
AI x Manufacturing	0.03808
	(0.13539)
AI x Services	0.06329
	(0.25373)
Constant	0.17855
	(1.21991)
AR(2) Test	0.25789
Hansen J-Test	0.69509
Observations	695
Note:	*p<0.1; **p<0.05; ***p<0.01

After the GMM was run, a standard procedure is to test for the validity of the instruments that were used as well as to check for second-order autocorrelation issues. Therefore, a Hansen J-Test along with an AR(2) test were performed to check for any potential issues with the regression.

As mentioned earlier, the Hansen J-Test essentially checks whether the instruments are uncorrelated with the error term. The null hypothesis is that all instruments are valid and will be accepted in case the *p-value* is above 0.10. The estimated *p-value* for the second regression lies at 0.70. Therefore, it can be safely concluded that all of the instruments used in this regression are valid, thus uncorrelated with the error term.

For similar reasons as with the estimation of Equation (3), the standard errors must also be clustered when estimating Equation (4). This adjustment is necessary, because the model assumes independence across observations. However, when doing panel data regressions, it is expected that the observations within the same country across the studied time period may be correlated. This could lead to invalid *p-values*, confidence intervals, and standard errors. Thus, standard errors are clustered by country to account for this and to keep the statistical significance realistic.

Furthermore, an AR(2) test is performed to check for second-order autocorrelation. In a GMM regression, first-order autocorrelation is not a problem and is expected due to the first-differencing. However, a second-order autocorrelation is considered problematic since it suggests that the instruments are invalid, meaning that they are correlated with the error term. The results from the AR(2) test are printed in Table 2. They indicate a *p-value* of 0.26. The Null Hypothesis (H_0) of the AR(2) test is that there is no second-order autocorrelation and generally a *p-value* above 0.05 lets us accept it. Thus, we conclude that for our second regression, estimated by Equation (4), there is no second-order autocorrelation.

The first thing that is clearly visible from Table 2 is that the lagged employment explanatory variable that we used in our dynamic panel data model is highly significant. This result is expected, as discussed by previous literature, taking into account that employment is a highly persistent variable by nature (Guliyev, 2023; Guliyev et al., 2023; Damioli et al., 2023; Bordot, 2022). Interestingly, the coefficient for the previous period's employment seems to be extremely high compared to the results produced by Equation (3). The reason for this is that there is a log transformation applied to Equation (4). Thus, the coefficients that the model produces are essentially elasticities between the dependent and independent variables. The log transformation effectively compresses large values in the panel data, which leads to reducing extreme variations. Thus, we observe a much higher coefficient of the lagged employment variable. In this case, the log transformation is necessary because of the differences in population sizes between countries, as discussed earlier. If not controlled for, the effect of the population size might bias our estimate and undermine our results.

In the regression output from Equation (4), it is visible that the inflation variable is now statistically significant, and it has a positive effect on employment. This result is in line with Phillips (1958) since the "Phillips curve" implies an inverse relationship between unemployment and inflation. The observed coefficient is in line with economic theory since the economy is stimulated under a higher inflation environment. Thus, the results from Table 2 indicate that a 1% increase in the inflation rate leads to a 0.748% increase in employment. This is anticipated to enhance the demand for labor and create new employment opportunities. When considering the reasons behind the statistical significance of inflation in the second regression, as opposed to its absence in the first, several plausible explanations may be identified. Equation (3) and Equation (4) have a similar intuition. However, there are some small differences that might contribute to the difference in the statistical explanation power of the inflation variable in the two equations. One such difference is that Equation (3) measures the effect of AI on aggregate unemployment across all sectors within the countries included in the study. Equation (4), on the other hand, measures the effect of AI on the employment levels of three main sectors, which are manufacturing, services, and agriculture, forestry and fishery. It might be so that inflation has a higher statistical power in explaining the variation of these three specific sectors, compared to the effect that it has on the aggregate level. For example, inflation might have a positive impact on some sectors and a negative one on others, and thus, the effects on the aggregate level might cancel out. On the other hand, if these three specific sectors all observe a positive effect from inflation, the variable might appear statistically significant in the second regression output.

The last statistically significant variable is the GDP growth rate. As expected in Equation (4) the coefficient has a positive sign, compared to Equation (3) since this time the dependent variable is employment. This result also confirms the previously discussed "Okun's law," according to which there should be a positive effect of GDP growth on the labor market since good economic conditions lead to new jobs created, higher demand for workers, and overall better job market. Therefore, according to our model, the results indicate that a 1% increase in GDP growth is associated with a 0.557% increase in employment.

Lastly, it is important to comment on the results from our AI patent variable. As seen from Table 2, the AI patent variable, as well as the two dummy variables and the two interaction terms, appear statistically insignificant. One possible reason for this result is that the model is essentially splitting up the effect of AI between multiple sectors. The regression is estimating the effects of AI separately on manufacturing and services relative to agriculture, fishery and forestry. However, there could be a potential problem if AI has a very similar impact across all three sectors, because this will increase the standard errors. If the effect of AI on manufacturing and services is very similar, the regression struggles to tell the difference, which leads to high standard errors for both. The logic is simple: high standard errors lead to wide confidence intervals, which makes coefficients insignificant.

To test this, an additional test developed by Wald (1943) was run. Table 3 shows us the results from the first test, which checks whether the effect that AI has on the manufacturing sector relative to the agriculture, forestry and fishery sector is different from the effect that AI has on the service sector relative to the agriculture, forestry and fishery sector. The Null Hypothesis (H_0) is that the coefficients of the two interaction terms $DMAN \times AIG_t$ and $DSER \times AIG_i$ are not statistically different from each other. If the *p-value* is above 0.05, we cannot reject the H_0.

The formulation of the Wald test can be found in Equation (5).

$$W = \frac{(\beta_1 - \beta_2)^2}{VAR(\beta_1 - \beta_2)} \qquad (5)$$

Here β_1 and β_2 are the coefficients that we want to test whether they are significantly different from each other and $VAR(\beta_1 - \beta_2)$ is the variance of the difference. If W is large enough, there is enough evidence to reject the Null Hypothesis (H_0) that $\beta_1 = \beta_2$.

Table 3: AI's Effect on Manufacturing vs. Services

Chi.sq.Statistic	p.value	Degrees.of.Freedom
0.027	0.869	1

From Table 3 it is visible that the *p-value* is 0.869 which is well above the threshold. Therefore, we can conclude that AI's effect on the manufacturing sector doesn't differ from AI's effect on the services sector.

Moreover, since the regression output from the GMM model indicates that the two coefficients from the interaction terms are not statistically significant, meaning that the model cannot conclude that the effect of AI on manufacturing and services relative to agriculture, forestry, and fishery is different from zero. Therefore, AI might be affecting both sectors similarly.

This finding, along with the Wald test, which showed us that also β_8 and β_9 are not statistically different from each other, leads us to the conclusion that we cannot infer a heterogeneous effect of AI on the studied sectors. This does not mean that AI has no effect on employment levels; it just means that AI does not have a heterogeneous effect across the three sectors, that were studied.

Another possible hypothesis is that the effect between the three sectors has an opposite sign and gets canceled out, which leads to insignificant results for the interaction terms. However, this will be visible in the Global AI patent variable coefficient. Therefore, since this variable is also not statistically significant, this leads us to think that AI's effect on agriculture, forestry, and fishery is not statistically different from zero. Moreover, since the effect on the manufacturing sector and the service sector is also not significantly different from the effect on the reference category, we might think the effect of AI is similar across sectors.

4 Conclusion

This study examines the effect that artificial intelligence has on the labor market in 21 high-tech developed countries for the period 2001-2020 using a dynamic panel data estimation, which incorporates the lag of our dependent variable as an explanatory variable, considering the high persistence of unemployment and employment data. The controls that were used to isolate the effect of AI development were inflation, GDP growth, as well as government spending on education. The proxy used for the main variable of interest is also a novelty to this research. To this end, the global AI patent count was used to explain variation in unemployment and employment data. As argued throughout this paper, AI is a global phenomenon, and technologies developed in one country may affect unemployment in another country. A Blundell-Bond (1998) two-step system GMM was used to run two separate regressions.

The first regression was aimed at studying the aggregate effect of AI technologies on unemployment. Thus, contributing essentially to the debate on whether this new wave of technology is causing a *displacement effect* or a *reinstatement effect.* The results from the first regression proved the importance of including the lag of unemployment on the right-hand side, since it showed to be highly statistically significant. No evidence was shown to support the Phillips's (1958) curve, as well as the importance of government spending on education. Nevertheless, evidence towards the inverse relationship of economic growth and unemployment was found, as argued by Okun (1963). Lastly, the main variable of interest, global AI patents, showed to explain variation in the unemployment rate. AI technologies seem to have a negative effect on unemployment, thus validating a *reinstatement effect.*

The second regression aimed at establishing what the heterogeneous effect of AI is on three different employment sectors, which were manufacturing, services, as well as agriculture, forestry and fishery. A similar regression setup was used, with the only difference of having employment as a dependent variable and using dummy variables to separate the effect on different sectors. The results showed again the significance of the dynamic setup and the GDP growth variable. Moreover, evidence for the Phillips's (1958) curve dynamic was found. Lastly, according to the results from the second regression, AI doesn't show a significant or different effect on sectors separately. That is true at least for the three sectors chosen for this research.

The main novelty of this paper is the use of global AI patents as a proxy, which showed that the effect of AI on the labor market is positive but considerably smaller than previous studies have shown. The direction of the dynamic that was described may look optimistic for the future of job markets and productivity around the world. However, there are a few factors that policymakers as well as innovators should stay conscious about when dealing with this new technological breakthrough. As argued by Acemoglu and Restrepo (2019), if automation continues to be our main source of productivity increase, job markets will stagnate. Creating a technology that aims at replicating human intelligence and skills is a dangerously unknown path to take and the argument that this will free people from repetitive, low-skill jobs and create more creative, high-skill jobs is just a false premise already proven wrong by various studies including Webb (2019), Felten et al., (2018) and Cazzaniga et al. (2024). Nevertheless, the possible positive and productivity-enhancing effects of AI on the labor market are not to be simply discarded. Thus, research needs to continue studying this topic in order to navigate policymakers as well as developers of the technology.

A few notable limitations of this study ought to be mentioned. Firstly, the data on AI patents is far from perfect, as the most recent data is up until 2020, which does not include the recent boom in Generative AI technologies like ChatGPT or the most recent Chinese competitor DeepSeek. This also leads to another consideration, which is that research usually lags behind innovation, due to data limitations. Technology is first created, and the economic implications and consequences are later found out. Therefore, policymakers ought to keep in mind the unknown nature of the technology and aim at antifragility by reinforcing the job markets and navigating technological innovation.

Appendix

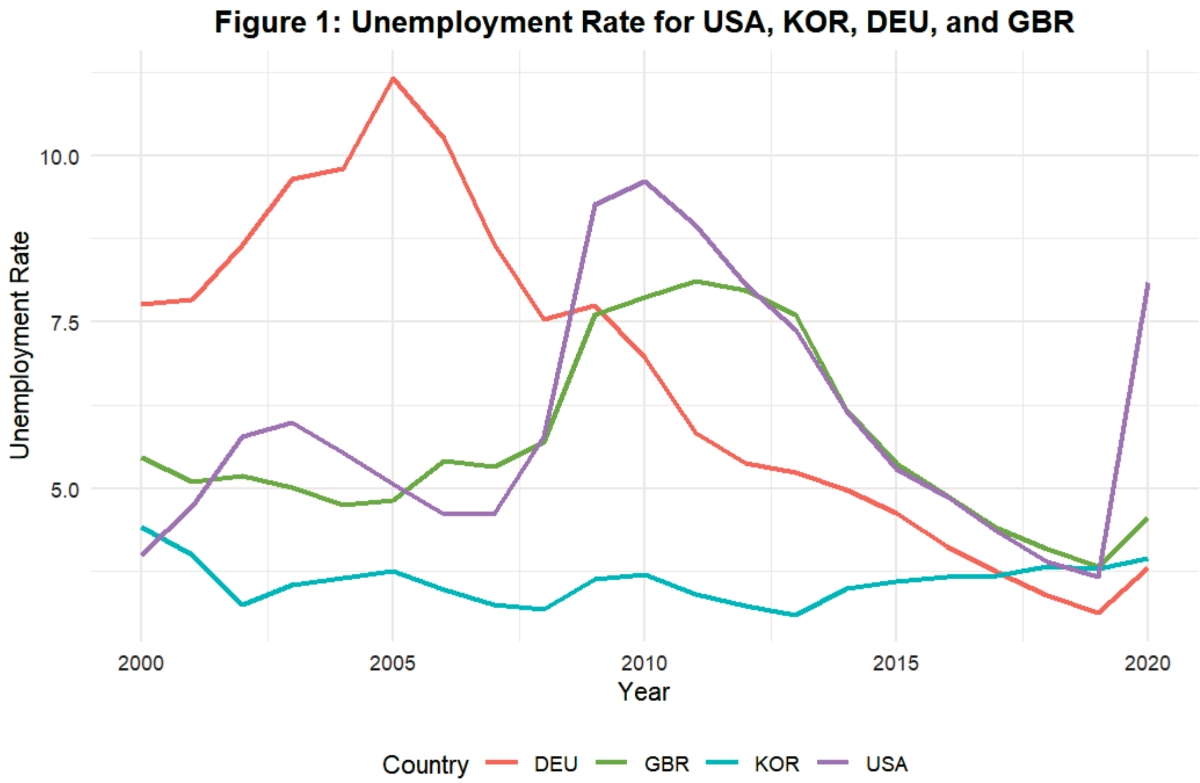
Figure 1: Unemployment Rate for USA, KOR, DEU, and GBR
10.0
7.5
5.0
Unemployment Rate
2000
2005
2010
2015
2020
Year
Country
DEU
GBR
KOR
USA

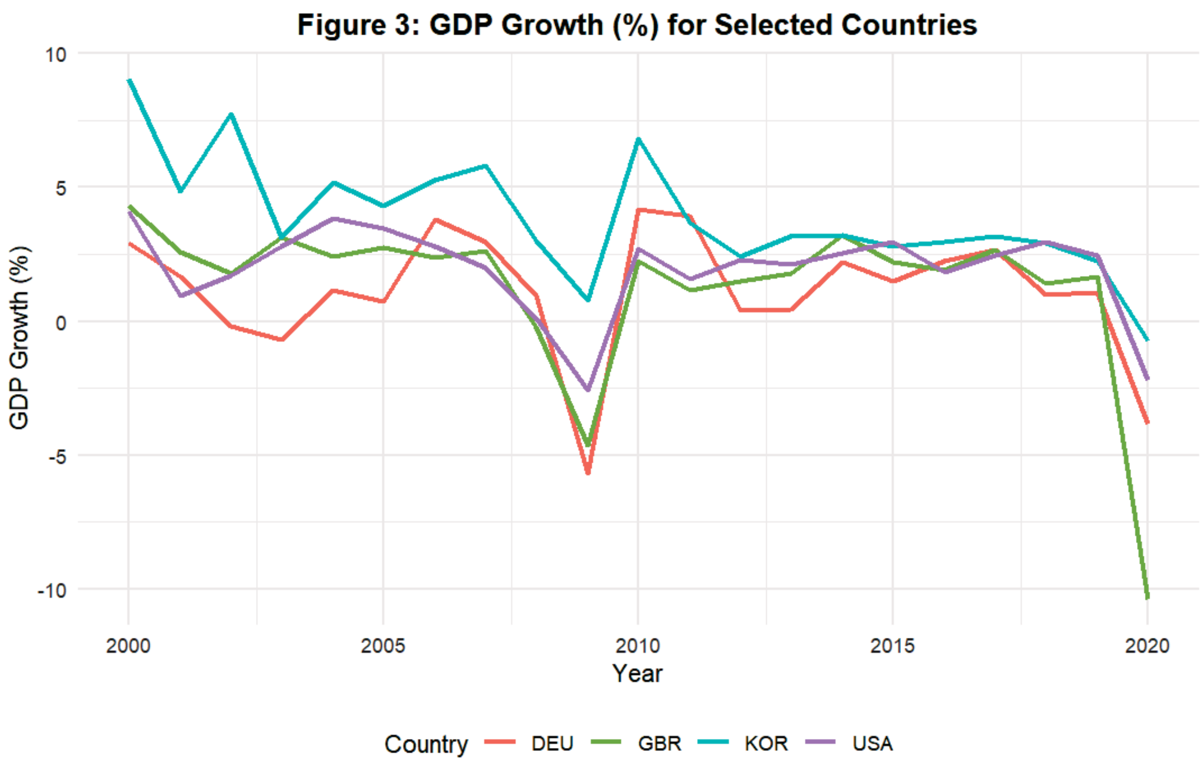
Figure 3: GDP Growth (%) for Selected Countries
10
5
0
-5
-10
GDP Growth (%)
2000
2005
2010
2015
2020
Year
Country
DEU
GBR
KOR
USA

Figure 4: Unemployment Rate in the USA

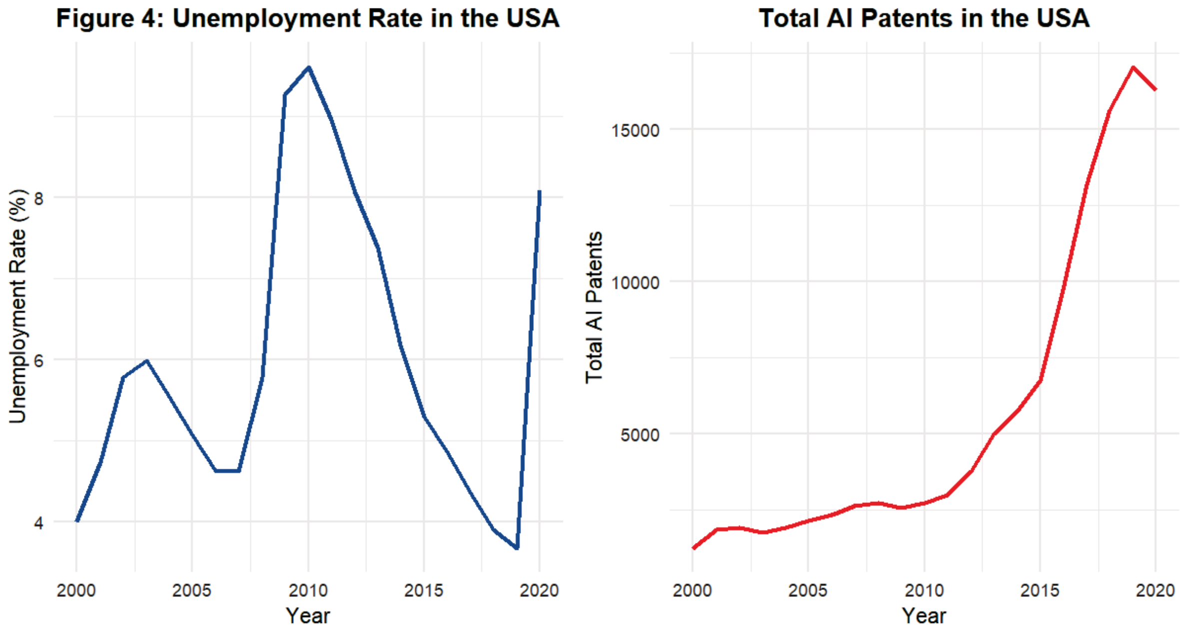

Bibliography

Acemoglu, Daron, and Pascual Restrepo. "Artificial intelligence, automation, and work." *The economics of artificial intelligence: An agenda*. University of Chicago Press, 2018. 197-236.

Acemoglu, Daron, and Pascual Restrepo. "Automation and new tasks: How technology displaces and reinstates labor." *Journal of economic perspectives* 33.2 (2019): 3-30.

Acemoglu, Daron. "The simple macroeconomics of AI." *Economic Policy* 40.121 (2025): 13-58.

Albanesi, Stefania, et al. "New technologies and jobs in Europe." *Economic Policy* 40.121 (2025): 71-139.

Arellano, Manuel, and Olympia Bover. "Another look at the instrumental variable estimation of error-components models." *Journal of econometrics* 68.1 (1995): 29-51.

Arellano, Manuel, and Stephen Bond. "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations." *The review of economic studies* 58.2 (1991): 277-297.

Ashenfelter, Orley, and John Ham. "Education, unemployment, and earnings." *Journal of Political Economy* 87.5, Part 2 (1979): S99-S116.

Babina, Tania, et al. "Artificial intelligence, firm growth, and product innovation." *Journal of Financial Economics* 151 (2024): 103745.

Blundell, Richard, and Stephen Bond. "Initial conditions and moment restrictions in dynamic panel data models." *Journal of econometrics* 87.1 (1998): 115-143.

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

Cazzaniga, Mauro, et al. *Gen-AI: Artificial intelligence and the future of work*. International Monetary Fund, 2024.

Damioli, Giacomo, et al. "AI technologies and employment: micro evidence from the supply side." *Applied Economics Letters* 30.6 (2023): 816-821.

Dutta, Soumitra, et al., eds. *Global Innovation Index 2024: Unlocking the Promise of Social Entrepreneurship*. World Intellectual Property Organization, 2024.

Felten, Edward W., Manav Raj, and Robert Seamans. "A method to link advances in artificial intelligence to occupational abilities." AEA Papers and Proceedings. Vol. 108. 2014 Broadway, Suite 305, Nashville, TN 37203: American Economic Association, 2018.

Felten, Edward, Manav Raj, and Robert Channing Seamans. "The effect of artificial intelligence on human labor: An ability-based approach." *Academy of Management Proceedings*. Vol. 2019. No. 1. Briarcliff Manor, NY 10510: Academy of Management, 2019.

Guliyev, Hasraddin, Natiq Huseynov, and Nasimi Nuriyev. "The relationship between artificial intelligence, big data, and unemployment in G7 countries: New insights from dynamic panel data model." *World Development Sustainability* 3 (2023): 100107.

Guliyev, Hasraddin. "Artificial intelligence and unemployment in high-tech developed countries: new insights from dynamic panel data model." *Research in Globalization* 7 (2023): 100140.

Hansen, Lars Peter. "Large sample properties of generalized method of moments estimators." *Econometrica: Journal of the econometric society* (1982): 1029-1054.

Hoeffler, Anke, Stephen Roy Bond, and Jonathan RW Temple. "GMM estimation of empirical growth models." (2001).

Korinek, Anton, and Joseph E. Stiglitz. "Artificial intelligence and its implications for income distribution and unemployment." *The economics of artificial intelligence: An agenda*. University of Chicago Press, 2018. 349-390.

Mutascu, Mihai. "Artificial intelligence and unemployment: New insights." *Economic Analysis and Policy* 69 (2021): 653-667.

OECD 2024, Annual labour force survey, summary tables, https://data-explorer.oecd.org/vis?fs[0]=Topic%2C1%7CEmployment%23JOB%23%7CUnemployment%20indicators%23JOB_UNEMP%23&pg=0&fc=Topic&bp=true&snb=19&df[ds]=dsDisseminateFinalDMZ&df[id]=DSD_ALFS%40DF_SUMTAB&df[ag]=OECD.SDD.TPS&df[vs]=1.0&dq=.LF.._Z._T....A&lom=LASTNPERIODS&lo=13&to[TIME_PERIOD]=false, accessed in September 2024

Okun, Arthur M. Potential GNP: its measurement and significance. Cowles Foundation for Research in Economics at Yale University, 1963.

Onuoha, Favour C., and Agbede Moses Oyeyemi. "Impact of disagregated public expenditure on unemployment rate of selected African countries: A panel dynamic analysis." Journal of Economics, Management and Trade 24.5 (2019): 1-14.

Phillips, Alban W. "The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861-1957." *economica* 25.100 (1958): 283-299.

Riddell, W. Craig, and Xueda Song. "The impact of education on unemployment incidence and re-employment success: Evidence from the US labour market." *Labour economics* 18.4 (2011): 453-463.

UNESCO Institute for Statistics (UIS) database, https://databrowser.uis.unesco.org/browser/EDUCATION/UIS-EducationOPRI/gov-exp

Wald, Abraham. "Tests of statistical hypotheses concerning several parameters when the number of observations is large." *Transactions of the American Mathematical society* 54.3 (1943): 426-482.

Webb, Michael. "The impact of artificial intelligence on the labor market." *Available at SSRN 3482150* (2019).