Working Paper No. 209

"This time it's different" Generative Artificial Intelligence and Occupational Choice

Daniel Goller, Christian Gschwendt and Stefan C. Wolter

Universität Zürich
IBW – Institut für Betriebswirtschaftslehre

UNIVERSITÄT BERN

Working Paper No. 209

“This time it’s different” Generative Artificial Intelligence and Occupational Choice

Daniel Goller, Christian Gschwendt and Stefan C. Wolter

November 2023

The Swiss Leading House on Economics of Education, Firm Behavior and Training Policies is a Research Program of the Swiss State Secretariat for Education, Research, and Innovation (SERI).

www.economics-of-education.ch

“This time it’s different”

Generative Artificial Intelligence and Occupational Choice

Daniel Goller[1], Christian Gschwendt[1], Stefan C. Wolter[1,2,3]

[1] Centre for Research in Economics of Education, University of Bern
[2] Swiss Coordination Centre for Research in Education, Aarau
[3] CESifo & IZA

Abstract: In this paper, we show the causal influence of the launch of generative AI in the form of ChatGPT on the search behavior of young people for apprenticeship vacancies. There is a strong and long-lasting decline in the intensity of searches for vacancies, which suggests great uncertainty among the affected cohort. Analyses based on the classification of occupations according to tasks, type of cognitive requirements, and the expected risk of automation to date show significant differences in the extent to which specific occupations are affected. Occupations with a high proportion of cognitive tasks, with high demands on language skills, and those whose automation risk had previously been assessed by experts as lower are significantly more affected by the decline. However, no differences can be found with regard to the proportion of routine vs. non-routine tasks.

Keywords: Artificial intelligence, occupational choice, labor supply, technological change

JEL classification: J24, O33

Address for correspondence:

Daniel Goller, Centre for Research in Economics of Education, University of Bern, Schanzeneckstrasse 1, Bern, Switzerland, daniel.goller@unibe.ch

Acknowledgments: The authors are grateful to the SDBB | CFSO for the provision of the data on queries on the national apprenticeship vacancy platform and acknowledge the generous financial support by the Swiss State Secretariat for Education, Research and Innovation through its Leading House on the Economics of Education (ECON-VPET). The usual disclaimer holds.

1 Introduction

The announcement and public deployment of the generative artificial intelligence (AI) ChatGPT version 3.5 led to a shock shift in perceptions of how AI can impact future workplaces. In professions that felt safe from being partially or entirely replaced, the risk of automation has suddenly increased considerably, because unlike in the past, when the risk of automation was equated with the replacement of manual and routine labor, the emerging generative artificial intelligence is targeting non-routine and cognitive tasks.

In this paper, we show how the unexpected launch of ChatGPT in November 2022 had an immediate, significant, and long-lasting effect on young people's search behavior for apprenticeship vacancies. The impact on search behavior for specific occupations is of great economic importance as it is a high-stakes decision. Young people apply for apprenticeships in very specific occupations and, if they get the apprenticeship of their choice, invest a large part of the next three to four years of their education in learning occupation-specific skills that enable them to perform specialized tasks. If technology can easily perform these tasks now or in the near future, this investment will be lost and must be replaced by further investment in human capital. It can, therefore, be assumed that expectations about such shifts in demand for specific occupations caused by technological changes are already considered as far as possible by the (labor) supply side, i.e., potential apprentices searching and deciding for a training occupation.

In the past, observations of skill shifts in industrialized labor markets showed, first and most, that digitalization was skill-biased (Katz & Murphy, 1992), shifting demand from less skilled to more skilled workers and leading to a skill upgrading of the labor force in developed economies (Berman, Bound, & Machin, 1998). Second, digital technology increasingly substituted workers in routine-intensive jobs while complementing workers executing non-routine tasks (Autor, Levy, & Murnane, 2003). As a result, many labor markets experienced a decline in routine employment and steady growth in non-routine cognitive jobs, which predominantly require higher levels of education, further contributing to the ongoing upgrading trends in developed economies (e.g., Goos & Manning, 2007; Goos, Manning and Salomons, 2009; Gschwendt, 2022). Similarly, advances in digital technologies can be considered a key driver of widespread skill upgrading (Katz & Autor, 1999; Acemoglu & Autor, 2011; Oesch and Rodríguez Menés 2011; Balsmeier and Woerter, 2019, as well as Pusterla and Renold, 2022), and, in turn, were more beneficial for skilled than for unskilled workers, as occupations became more demanding and complex and thus requiring longer, higher, and better training.

Considering the labor market effects of generative AI, however, initial observations and estimates point in a different direction than previous trends with regard to specific skills and tasks. First empirical studies show that the use of (generative) AI leads to increased productivity in more complex and less ambiguous tasks, with a greater impact on low-skilled than on highly skilled workers (Acemoglu et al., 2022; Haslberger, Gingrich, & Bhatia, 2023; Brynjolfsson, Li, & Raymond, 2023). While current analyses of job ads do not yet show a high prevalence of jobs in which AI is used (OECD, 2023a), early estimates suggest that many occupations will be affected by AI in the near future, with higher educated workers disproportionally exposed (Eloundou, Manning, Mishkin, & Rock, 2023).

Generative AI particularly improved in tasks demanding literacy/reading, less so in math skills, which is relevant for potential labor market effects, because while only between 27% and 44% of workers daily perform numeracy tasks at work, having numeracy proficiency below or at the level of AI, almost 60% of the workforce using literacy skills daily have a proficiency comparable to or below that of computers (numbers for countries having participated in the PIAAC survey, see OECD, 2023b).

In a first study on the impact of the introduction of ChatGPT, DALL-E-2, and Midjourney on labor demand, Hui, Reshef, and Zhou (2023) investigate on a large online platform the demand for freelancers providing work that AI tools are specialized in. The authors showed that the availability of these AI tools significantly decreased the employment and earnings of exposed freelancers, with a greater negative impact on freelancers who did higher-quality work before the tools were introduced. This could be seen as a first indication of how disruptive generative AI can be for labor markets. Our study is complementing this study in two ways. First, by investigating the supply-side effects of the new generative AI, i.e., how it affects the perception of tomorrow's workforce and how it influences their occupational choices. Second, by investigating this for a high-stakes and long-term decision, that is, to learn specialized occupation-specific skills for three or four years.

We find that with the launch of ChatGPT the supply of potential apprentices, in the form of search queries for apprenticeship positions, decreases substantially by, on average, about 8 percent. While this speaks for a general uncertainty about which skills will be made redundant by generative AI in the future, we see substantial heterogeneity for different occupations according to cognitive requirements and task types performed in these occupations. Occupations using primarily cognitive tasks are particularly affected by the shock in perception, with the largest decline in supply. Similarly, those occupations with the highest language requirements experienced the sharpest decrease in supply – an area of expertise in which, at first glance, a large language model such as ChatGPT is most likely to replace jobs. All of this contrasts previous waves of automation, which we can also see from the fact that for the occupations with the highest automation risk attributed in previous studies, we see the smallest reduction in supply: This time, it's different.

The remainder of the paper is structured as follows: In the next section, we show the interplay between the launch of ChatGPT and how the apprenticeship market works. In Section 3, we describe the data used, and in Section 4, the empirical strategy. Section 5 shows the empirical results, and Section 6 concludes with a discussion of the results.

2 The launch of ChatGPT and the Swiss apprenticeship market

With the announcement and public release of version 3.5 of ChatGPT[1] on November 30th, 2022, the public at large has been granted access to highly capable generative artificial intelligence for the first time. The opportunity for experimentation has led both the media and individuals to experience a major shift in their perception of the possibilities of AI from one day to the next.

This increased attention to the possibilities of AI hit Swiss adolescents in the midst of making a far-reaching, high-stakes decision about their educational and professional future. After compulsory school, the majority of young people in Switzerland start a three- or four-year apprenticeship and choose among about 240 different occupations. In an apprenticeship, apprentices learn general skills but, most importantly, occupation-specific skills. While some of these occupation-specific skills can be transferred to other occupations (Eggenberger, Janssen, & Backes-Gellner, 2022), there is a considerable risk that the significant time and financial investment of an apprenticeship may become partially worthless if some of the skills learned can be easily substituted by AI now or in the near future.

The apprenticeship market works in many aspects comparable to a usual labor market. Companies demand apprentices by opening vacancies; the supply side consists of individuals who apply for those vacancies.[2] If both sides agree, a work and training contract is signed. All apprenticeships start after the summer school holidays in a given year. Potential trainees begin looking for an apprenticeship approximately a year before the starting date of the apprenticeship and start signing training contracts as early as October/November of the preceding year, with a

---

[1] When referring to «ChatGPT», we always mean ChatGPT version 3.5.

[2] The interested reader is referred to, e.g., Wolter and Ryan (2011) or Muehlemann and Wolter (2020), for an overview of the economics of the apprenticeship market.

peak in the spring preceding the next school year. In other words, the launch of ChatGPT 3.5 coincided with the time when most young people interested in an apprenticeship were very concretely considering their choice of a specific occupation.

3 Data

For our research question, precisely tracking young people's career choice behavior over time is crucial. This is possible because we can use data from the national platform for apprenticeship vacancies. The majority of training companies operating in Switzerland advertise their apprenticeship vacancies on this online platform.[3] Interested young people can use this platform to search for specific professions within a certain geographical perimeter and thus obtain information on which apprenticeships are still available from which companies. If they are interested, they can also find the training companies' contact details so they can apply. The search behavior on this platform is, therefore, more than just a search for information about occupations (for which other platforms exist), but shows a concrete interest in finding an apprenticeship in a specific occupation. All search queries on this platform are recorded and used in this paper in an aggregated form on a daily basis, showing the number of queries by occupation.[4]

[3] This proxy for labor supply was first introduced and used by Goller and Wolter (2021). The national online vacancies platform is berufsberatung.ch (de; fr: orientation.ch; it: orientamento.ch). This is the official platform of the Swiss Conference of Cantonal Ministers of Education for information about and search for apprenticeship positions. There are also other, private, platforms, but the national platform, collects data on vacancies from employers in all cantons, and is likely to be the platform that the majority of apprenticeship seekers rely on during their search and application process.

[4] Goller and Wolter (2023) show that search queries are linked to changes in the equilibrium in the apprenticeship market. They show that increasing search queries caused by an exogenous reputation shock to some occupations translated into significantly more signed apprenticeship contracts in the year following the shock.

For this research, we extracted the number of search queries from log files from the 1st of January 2021 until the 30th of June 2023, resulting in more than 4 million observations. The outcome variable represents the number of search queries by occupation, canton, and day and is used in a logarithmic transformation for easier interpretation. The treatment variable is constructed as a binary variable indicating the time before (=0) and after (=1) the public launch of ChatGPT version 3.5 on the 30th of November 2022.

Other covariates are information on bank holidays and school vacations in each canton. The number of open vacancies is added for each day, canton, and occupation. We also classify occupations according to different dimensions of work content and skills used. First, for each occupation, we use the official information from schools and career counselors on the cognitive job requirements according to four dimensions: the language of the region, foreign languages, math, and science.[5] Second, we classified occupations along the two dimensions of cognitive vs. manual and routine vs. non-routine tasks (Mihaylov & Tijdens, 2019). Third, we added the measures on automation risk and exposure to software, robots, and AI in a particular occupation from Frey and Osborne (2017) and Webb (2019). Because there are no profiles of job requirements available for a few occupations that can be learned through apprenticeship in Switzerland, the subsample analyses involving this information are on a reduced sample. All available variables can be found in the descriptive statistics in Appendix A, Table 4.

4 Empirical Strategy

When OpenAI released ChatGPT version 3.5 on the 30th of November 2022 for public use, it was a surprise to most people. The evidence for this is that while AI as a generic topic had already been dealt with in the media earlier, media coverage of ChatGPT and what this tool

[5] Received from Anforderungsprofile.ch.

can achieve only started on the day of its launch. The media coverage afterward was rapid and very broad, as the data in Figure 1 shows, and many articles dealt in detail and with practical examples of what work could be substituted by ChatGPT.[6] Crucial to our analysis are the two observations that before the launch of ChatGPT, only very few insiders knew about the possibilities of this and similar AI tools, but after the launch of ChatGPT, information about it spread rapidly. The launch of ChatGPT on November 30, 2022, can, therefore, be seen as an exogenous information shock with a clear before and after. Therefore, the variation we exploit in this study comes from the natural experiment of the unanticipated public launch of ChatGPT.

To investigate the causal effect in this event study, we must ensure that no other patterns related to the timing of the public launch and search patterns in the labor market contaminate our estimated effects. For this, we take several potential confounders into account. First, during a year, searches for apprenticeship vacancies follow some cyclical pattern that results from the clearing mechanisms of the apprenticeship market. Second, the maximum supply is determined by yearly changing cohort sizes that potentially enter the apprenticeship market. Third, each profession differs in terms of popularity. To account for all those potential concerns, we use month, year, and occupation fixed effects. Moreover, we control for days of the week, school vacations, and public holidays because previous work (Goller & Wolter, 2021) showed a high and robust statistical correlation between them and our measure of the supply of apprentices. This is particularly important in this study because of the Christmas holidays and accompanying school vacations shortly after the announcement of ChatGPT, which otherwise could lead to confounded effects.

[6] The same can be observed when analyzing search queries on the Google search engine (Figure 3 in the Appendix).

Figure 1

Media articles mentioning AI and ChatGPT before and after the introduction of ChatGPT 3.5

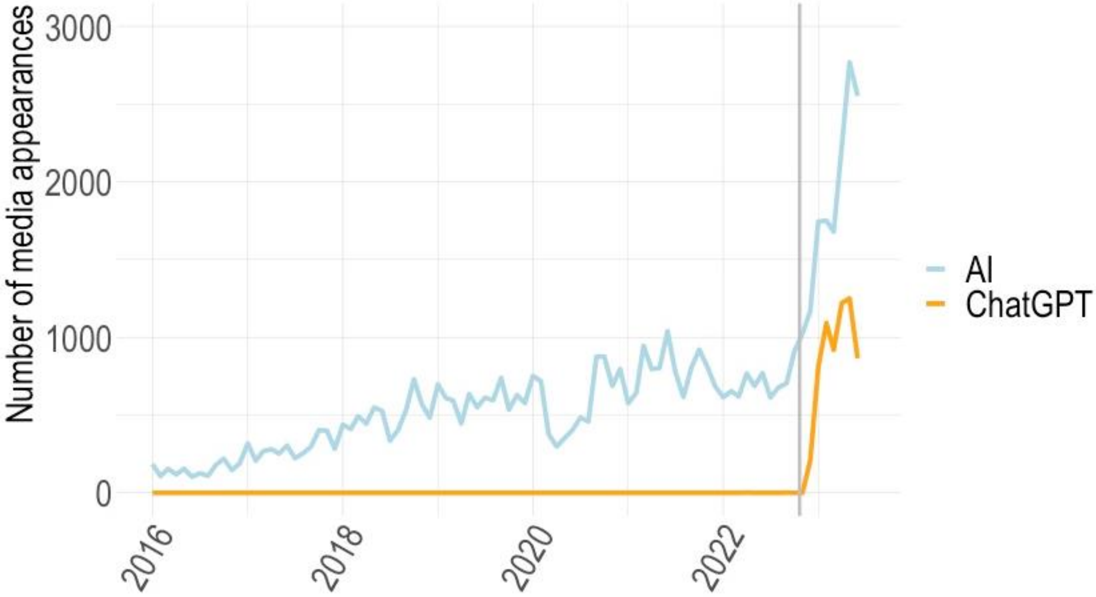

Notes: Number of media articles by month that include the terms “AI” (in German: “Künstliche Intelligenz”) and “ChatGPT” in Switzerland. Only media in the German language is considered. Among “media,” there are mainly newspaper articles and radio or TV shows for which the transcript is archived in the media archive. The grey line shows the introduction of ChatGPT.

We assess the credibility of this specification to estimate causal effects in multiple ways. First, we assess the sensitivity toward including different covariates in the model. Second, we use different samples and different time periods to estimate the effects in robustness checks. Third, we investigated if the demand side influences the effect, i.e., by offering more or fewer vacancies. Fourth, in a placebo treatment exercise, we investigate a pseudo-treatment for the days/weeks before the public deployment of ChatGPT.

5 Empirical Results

5.1 Before-After Analysis

Table 1

Before-After Analysis of the Introduction of ChatGPT on Labor Supply; Sensitivity of Effects

	(1)	(2)	(3)	(4)
ChatGPT	-0.058***	-0.075***	-0.079***	-0.078***
	(0.012)	(0.010)	(0.010)	(0.010)
School vacation			-0.095***	-0.083***
			(0.008)	(0.007)
Public holiday			-0.107***	-0.100***
			(0.010)	(0.009)
Month FE		X	X	
Year FE		X	X	X
Occupation FE		X	X	
Day of the week			X	X
Canton FE				X
Occupation by month FE				X
Sample	01.01.2021 – 30.06.2023			
N	4,611,880	4,611,880	4,611,880	4,611,880

Notes: Linear regression. The outcome is the log of the number of search queries by day, canton, and occupation. Standard errors are clustered on the occupation level. Every regression includes a binary indicator for the days the platform was the victim of a hacker attack. *, **, and *** indicate statistical significance at the 10, 5, and 1% level, respectively.

Table 1 shows the impact of the shock on perceptions about the capabilities of AI substituting human tasks, measured by the introduction of ChatGPT on the supply of apprentices, with and without the inclusion of potential confounders. Controlling for the most important confounders, like month, year, and occupation fixed effects (column 2), we find a coefficient of -0.075, which can be interpreted as 7.5 % fewer search queries following the public launch of ChatGPT. The result is robust to including additional controls in models 3 and

4.[7] In further robustness checks (Appendix B.2, Table 5), we can show that the effect is robust to different sample specifications and to controlling for the number of vacancies (the demand side). Moreover, in contrast to the supply of potential apprentices, we do not find a statistically significant impact of the introduction of ChatGPT on the demand for apprentices.

Figure 3 in Appendix B.1 shows the temporal variation of the impact in the weeks that immediately followed the introduction of ChatGPT. The effect is immediately visible, and the magnitude of the effect increases over the first weeks. Importantly, for the weeks before the launch, conducting a placebo-treatment estimation, we find a zero effect of the pseudo-treatment, i.e., no anticipation effect.

5.2 Task-type intensity

The influence on the aggregated searches for apprenticeship vacancies shows an effect that can be understood as a general shock because young people looking for an apprenticeship do not know what to do. On the other hand, it is evident that the possibilities opened up by generative AI, such as Chat GPT, with regard to the substitution (or complementation) of work tasks affect individual occupations very differently. For this reason, we turn to the potential heterogeneity of the effect according to different categories and types of occupations. To investigate this, we conduct our analyses separately for professions that predominantly contain certain task types that are more or less in danger of skill obsolescence due to generative AI.

Table 2 shows the estimates for different subsamples of occupations classified as predominantly routine versus non-routine and manual versus cognitive jobs. While the effects on the search for predominantly routine vs. non-routine jobs are almost identical, we find remarkable differences when comparing the impact on the search for more manual vs. more

[7] We use model 3 as our baseline specification in the following analyses.

cognitively demanding jobs, with the largest decline in searches for routine cognitive jobs by far.

Table 2

Subsamples; Automationrisk and Task-types

	(1)	(2)	(3)	(4)
	Routine	Non-Routine	Manual	Cognitive
ChatGPT	-0.077***	-0.080***	-0.054***	-0.142***
	(0.023)	(0.010)	(0.005)	(0.033)
N	1,552,980	3,058,900	3,317,730	1,294,150
	Routine Manual	Routine Cognitive	Non-Routine Manual	Non-Routine Cognitive
ChatGPT	-0.030***	-0.184**	-0.065***	-0.118***
	(0.006)	(0.070)	(0.006)	(0.033)
N	1,082,380	470,600	2,235,350	823,550

Notes: Linear regression. The outcome is the log of the number of search queries by day, canton, and occupation. Standard errors are clustered on the occupation level. Each regression uses the same specification as in Table 1, column (3). *, **, and *** indicate statistical significance on the 10, 5, and 1% level, respectively.

Figure 2 shows how the effects for manual and cognitive jobs evolve in the short term over the first ten weeks after the public launch of ChatGPT (vertical line). Although both categories of jobs experience a decrease in the search activity for vacancies, there is a statistically significant difference in the effect size after the 4th week. The overall negative effect can be seen over the entire time span of ten weeks following the launch of ChatGPT. Before the introduction of ChatGPT, both estimates were close to zero. This further indicates that there was no anticipation effect before the event, and both types of jobs followed similar pre-trends.[8]

[8] In a differences-in-differences estimation in Appendix C, Figure 4, we show that there are significant differences between cognitive and manual jobs. While both groups of jobs follow similar pre-trends, we cannot claim manual jobs to be completely unaffected by the event. This might be due to general uncertainty about which skills

Figure 2

Event Study (Before-After Analysis) by Manual- vs. Cognitive-type Jobs

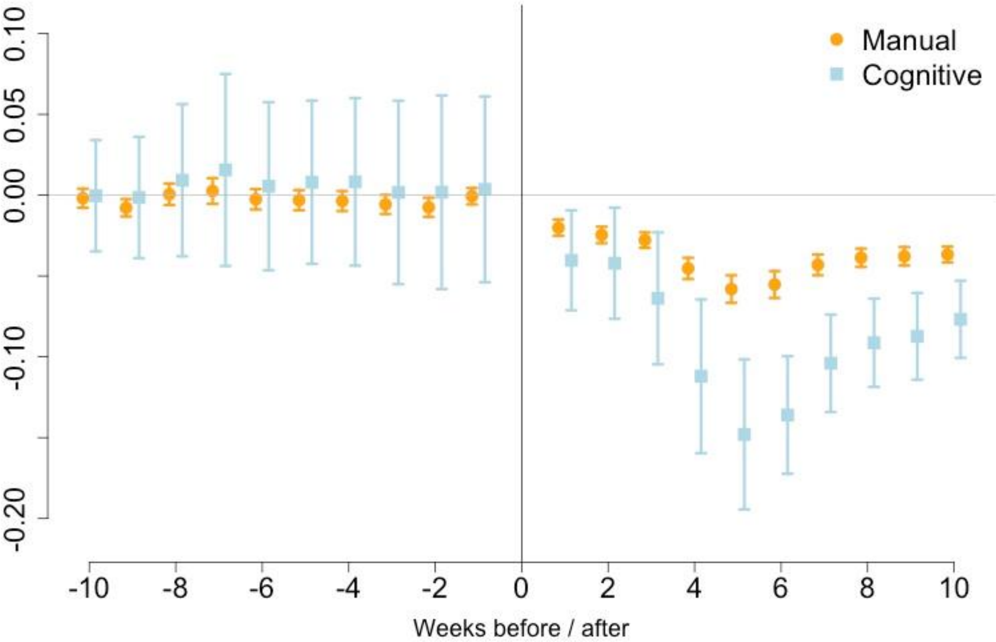

Notes: The vertical line represents the introduction of ChatGPT version 3.5 on November 30th, 2022. Occupations are divided into jobs that primarily feature manual or cognitive tasks (based on Mihaylov & Tijdens, 2019). On the left side, placebo treatment effects for the *x* weeks before the introduction are estimated. Each point refers to a point estimate of a separate linear regression (using column (3) model specification, Table 1). The whiskers mark the 90% confidence intervals, with standard errors clustered on the occupation level.

5.3 Occupational requirements

To complete the picture of which occupations are most affected, we classify occupations according to their cognitive requirements in four different school-type subjects. We differentiate the impact of ChatGPT along the four quartiles of requirements in the local language, foreign languages, math, science, and a combined measure of the four competencies (overall requirements).

---

become obsolete with generative AI. Consequently, the assumptions for a difference in difference model are arguably not fulfilled, and the analysis is relegated to the appendix.

Table 3

Subsample effects by quartiles of occupational requirements

Requirements...	(1) Lowest	(2) low	(3) high	(4) highest
		...overall		
ChatGPT	-0.058***	-0.065***	-0.073**	-0.142***
	(0.010)	(0.011)	(0.029)	(0.039)
		... in local language		
ChatGPT	-0.053***	-0.049***	-0.081***	-0.151***
	(0.009)	(0.007)	(0.013)	(0.046)
		... in foreign language		
ChatGPT	-0.055***	-0.056***	-0.062***	-0.160***
	(0.007)	(0.008)	(0.015)	(0.045)
		... in math		
ChatGPT	-0.127***	-0.103**	-0.048***	-0.061***
	(0.029)	(0.040)	(0.008)	(0.011)
		... in science		
ChatGPT	-0.112***	-0.084***	-0.058***	-0.084***
	(0.040)	(0.029)	(0.014)	(0.012)

Notes: Linear regressions on the outcome, the log of the number of search queries by day, canton, and occupation. Here, a reduced sample of occupations is used for which we observe the occupational requirements (N=3,623,620). Subsamples are analyzed in the four columns by splitting the sample into four quartiles, each about 25% of observations according to the requirements in the respective skill. Requirements are weighted by their relevance for each profession. Standard errors are clustered on the occupation level. *, **, and *** indicate statistical significance on the 10, 5, and 1% level, respectively.

Table 3 provides the results for the estimations on those occupations in the respective requirements quarter. In line with expectations, because ChatGPT is, at first sight, useful for language-related tasks, the decrease in search queries for jobs with the highest requirements in the local and foreign languages is the largest. Because most apprenticeships have high requirements for language, this also translates into having the biggest negative impact on the occupations with the highest overall requirements. In contrast to languages, occupations with the lowest requirements in mathematics and natural sciences recorded the biggest drop in search queries.

5.4 Prevalent Automation Risk Measures

To complete the picture, we investigate how supply is affected for subsets of occupations categorized according to prevalent automation risk and exposure estimations. Table 6 in Appendix B.2 presents estimations for a breakdown of occupations according to the risk of substitution by automation, software, robots, or AI, as proposed by Frey and Osborne (2017) and Webb (2019). Even though the standard errors for the occupations that were assigned a low risk of substitution in all classifications are quite high, all effect sizes point in the direction that search activities were negatively affected by the launch of ChatGPT in these occupations in particular, generally more so than in those occupations that were previously considered to be under threat.

6 Discussion and Conclusion

In this paper, we present the first causal evidence of the influence of the availability of generative AI on the career choice behavior of young people. On the one hand, the launch of ChatGPT triggered lasting uncertainty, which led to a slump in search queries for apprenticeship vacancies in general. What is far more interesting, however, is that the individual professions were affected by the slump to very different degrees. As expected, occupations with high cognitive requirements in terms of language skills were significantly more affected than those with high requirements in terms of mathematical skills and, as was also to be expected, occupations with a high proportion of manual tasks were less affected by the decline than occupations with a high proportion of cognitive tasks. Compared with previous estimates on the substitutability of jobs by computer, robot, or AI technology, the results ultimately indicate that search queries fell less sharply in occupations traditionally considered at high risk of substitutability.

As far as the latter is concerned, two explanations are possible. On the one hand, it is possible that the risk of substitutability in occupations with a traditionally high risk was already present in the minds of young people, and therefore, no additional shock was triggered by ChatGPT. On the other hand, they may assume that generative AI will now affect occupations that were previously considered low risk, rather than those always considered high risk. In other words, generative AI is seen as a "game changer" in the interplay between digitalization and developments in the labor market.

Only time will tell whether the young people who are faced with a high-stakes decision, that is, in which job-specific skills they want to invest in the next three to four years of training, are right or wrong in their assessments. At the very least we can now see that they have radically changed their assessments of how specific professions will be affected by the possibilities of AI through the use of tools such as ChatGPT. In another paper with data on apprenticeship search (Goller & Wolter, 2023), we were able to show that shocks to the supply of young people for apprenticeships of a comparable level can have a significant impact on the equilibrium later on, i.e., more or fewer apprenticeship contracts are concluded. Therefore, future research will also be necessary to see whether the availability of generative AI will lead to major shifts in the apprenticeship market.

References

Acemoglu, D., & Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics, Volume 4* (pp. 1043-1171). Elsevier.

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics, 40*(S1), 293-340.

Autor, D., Levy, F., & Murnane, R. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics, 118*(4), 1279-1333.

Balsmeier, B., & Woerter, M. (2019). Is this time different? How digitalization influences job creation and destruction. *Research Policy, 48*(8), 103765.

Berman, E., Bound, J., & Machin, S. (1998). Implications of skill-biased technological change: international evidence. *The Quarterly Journal of Economics, 113*(4), 1245-1279.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. *NBER working paper*, No. 31161.

Eggenberger, C., Janssen, S., & Backes-Gellner, U. (2022). The value of specific skills under shock: High risks and high returns. *Labour Economics, 78*, 102187.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv:2303.10130v4*.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114*, 254-280.

Goller, D., & Wolter, S. C. (2021). "Too shocked to search" The COVID-19 shutdowns' impact on the search for apprenticeships. *Swiss Journal of Economics and Statistics, 157*(6), 1-15.

Goller, D., & Wolter, S. C. (2023). Reaching for gold! The impact of a positive reputation shock on career choice. *IZA Discussion Paper Series*, No. 16607.

Goos, M., & Manning, A. (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *Review of Economics and Statistics, 89*(1), 118-133.

Goos, M., Manning, A., & Salomons, A. (2009). Job polarization in Europe. *American Economic Review, 99*(2), 58–63.

Gschwendt, C. (2022). Routine job dynamics in the Swiss labor market. *Swiss Journal of Economics and Statistics, 158*(1), 1-21.

Haslberger, M., Gingrich, J., & Bhatia, J. (2023). No Great Equalizer: Experimental Evidence on AI in the UK Labor Market. *SSRN working paper*, No. 4594466.

Hui, X., Reshef, O., & Zhou, L. (2023). The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market. *CESifo Working Paper*, No. 10601.

Katz, L., & Autor, D. (1999). Changes in the Wage Structure and Earnings Inequality. In *Handbook of Labor Economics. Vol. 3* (pp. 1463-1555). Elsevier.

Katz, L., & Murphy, K. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics, 107*(1), 35-78.

Mihaylov, E., & Tijdens, K. G. (2019). Measuring the Routine and Non-Routine Task Content of 427 Four-Digit ISCO-08 Occupations. *Tinbergen Institute Discussion Paper* , 2019-035/V.

Mühlemann, S., & Wolter, S. C. (2020). The Economics of Vocational Training, in: Bradley, S. and Green, C. (Eds.):. *The Economics of Education, A Comprehensive Overview, Second Edition, Elsevier*, 543-554.

OECD. (2023a). *Emerging Trends in AI Skill Demand Across 14 OECD Countries.* OECD Artificial Intelligence Papers. OECD Publishing, Paris, https://doi.org/10.1787/dee339a8-en.

OECD. (2023b). *Is Education Losing the Race with Technology? AI's Progress in Maths and Reading.* Educational Research and Innovation. OECD Publishing, Paris, https://doi.org/10.1787/73105f99-en.

Oesch, D., & Jorge Rodríguez Menés, J. (2011). Upgrading or polarization? Occupational change in Britain, Germany, Spain and Switzerland, 1990-2008. *Socio-Economic Review, 9*(3), 503-53.

Pusterla, F., & Renold, U. (2022). Does ICT affect the demand for vocationally educated workers? *Swiss Journal of Economics and Statistics, 158*(1), 1–22.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *SSRN working paper*, 3482150.

Wolter, S. C., & Ryan, P. (2011). Apprenticeship, in: Hanushek, E.A., Machin, S. and Woessmann, L. (Eds.):. *Handbook of Economics of Education, Volume 3, Elsevier*, 521-576.

Appendices

Appendix A: Descriptive statistics

Table 4

Descriptive Statistics

	Mean	Std. dev.
Panel A: Full sample (N= 4,611,880)		
Log (Search Queries)	0.373	(0.793)
Public holiday	0.026	(0.159)
School vacation	0.243	(0.429)
ChatGPT	0.235	(0.424)
Automation risk	0.673	(0.255)
Routine task intensity score	-0.218	(0.682)
Routine job	0.337	(0.473)
Manual task intensity score	0.273	(0.699)
Software exposure percentage	0.511	(0.217)
Robot exposure percentage	0.755	(0.567)
AI exposure percentage	0.438	(0.193)
Panel B: Sample May 2021 – June 2023 (N= 3,893,760)		
Log (Search Queries)	0.368	(0.787)
Public holiday	0.026	(0.158)
School vacation	0.247	(0.431)
Open vacancies	8.408	(29.042)
ChatGPT	0.273	(0.446)
Panel C: Reduced sample (N= 3,623,620)		
Log (Search Queries)	0.410	(0.841)
Overall intensity	31.390	(8.428)
Own language intensity	38.976	(15.059)
Foreign language intensity	10.070	(11.718)
Math intensity	36.947	(16.889)
Science intensity	39.567	(13.431)

Notes: Log (Search Queries) is the log of the number of search queries by day, canton, and occupation. Std. dev. = Standard deviation.

Figure 3

Google Searches before and after the introduction of ChatGPT 3.5

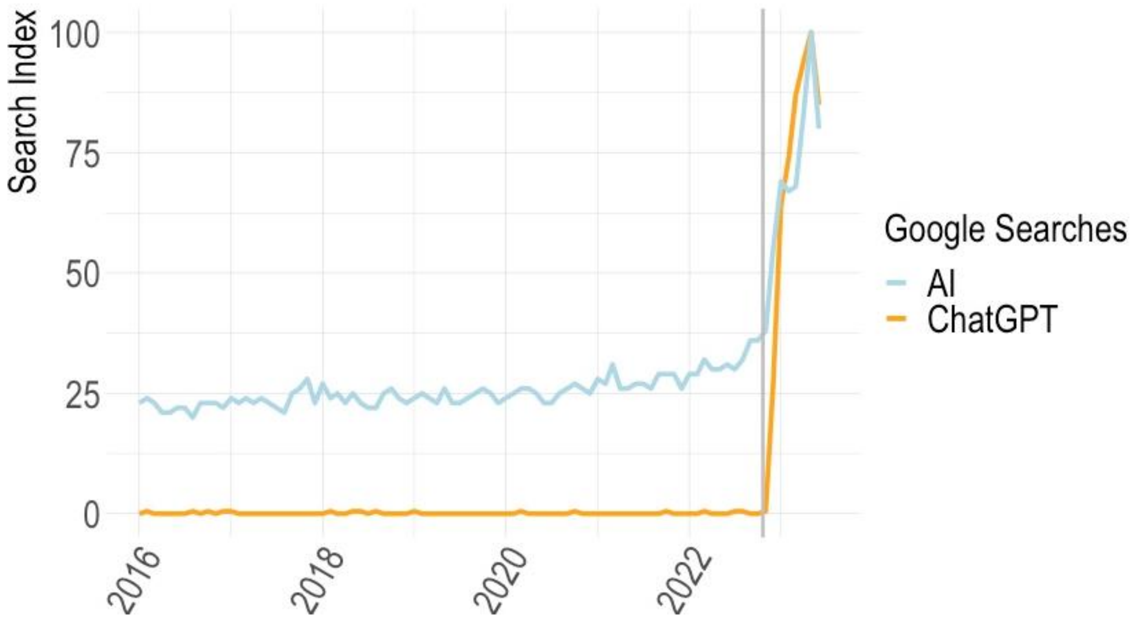

Notes: Search intensity on the Google search engine for the terms "AI" and "ChatGPT" in Switzerland. An Index of 100 shows the highest search intensity; lower values represent proportionally smaller numbers of searches. The grey line shows the introduction of ChatGPT.

Appendix B: Additional Results

Appendix B.1: Average effects over time

Figure 4

Before and after the introduction of ChatGPT 3.5, all occupations; average effect

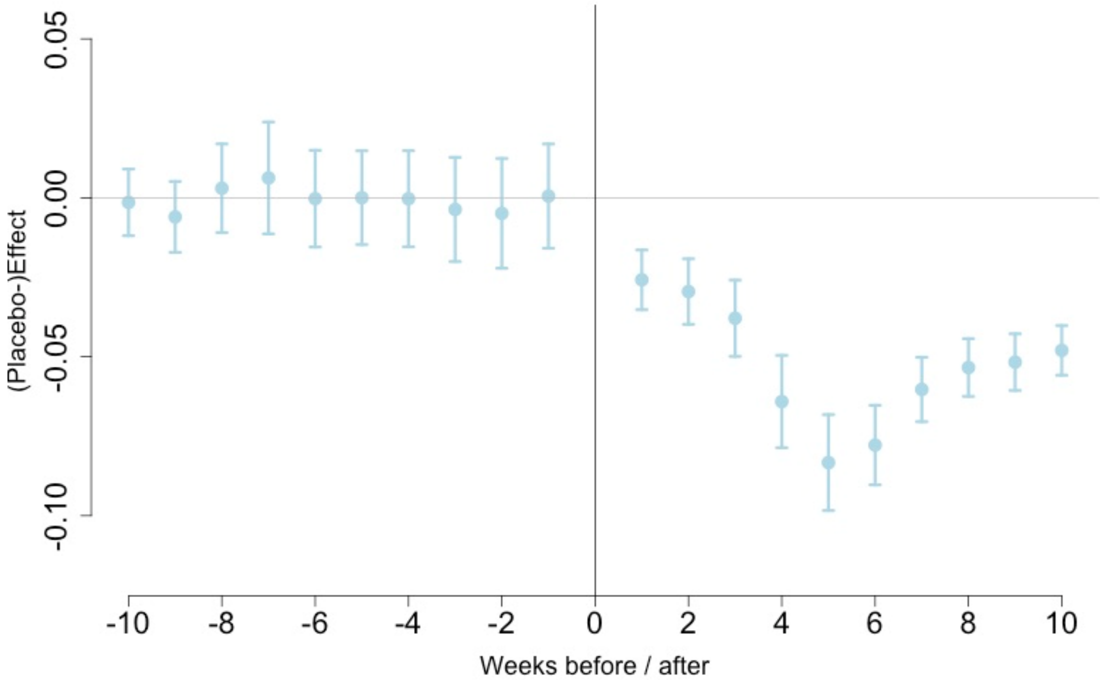

Notes: Linear regression based on Specification in Column (5) in Table 1. Effects for the *t* weeks before/after the introduction of ChatGPT 3.5 (in *t*=0). 90 % confidence intervals based on standard errors clustered on the occupation level.

Appendix B.2: Subsample results

Table 5

Before-After Analysis of the Introduction of ChatGPT; Different subsamples

	(1)	(2)	(3)	(4)	(5)
	Search Queries	Search Queries	Search Queries	Search Queries	Open Vacancies
ChatGPT	-0.079***	-0.084***	-0.079***	-0.082***	0.325
	(0.010)	(0.013)	(0.008)	(0.008)	(0.373)
School vacation	-0.095***	-0.102***	-0.092***	-0.088***	-0.412***
	(0.008)	(0.009)	(0.008)	(0.007)	(0.061)
Public holiday	-0.107***	-0.117***	-0.099***	-0.096***	-0.648***
	(0.010)	(0.012)	(0.009)	(0.009)	(0.108)
Open vacancies				0.007***	
				(0.001)	
Month FE	X	X	X	X	X
Year FE	X	X	X	X	X
Occupation FE	X	X	X	X	X
Day of the week	X	X	X	X	X
Sample (time)	01.01.2021 – 30.06.2023		06.05.2021 – 30.06.2023		
Sample (occupations)	All	Reduced	All	All	All
N	4,611,880	3,623,620	3,893,760	3,893,760	3,893,760

Note: Linear regression. Column (1) is the baseline estimation for comparison from Table 1. The outcome in columns (1) – (4) is like in the main analysis (Search Queries). The outcome in column (5) is the number of open vacancies by day, canton, and occupation. Standard errors are clustered on the occupation level. *, **, and *** indicate statistical significance at the 10, 5, and 1% level, respectively.

Table 6

Subsamples; Automationrisk and Task-types

	(1)	(2)	(3)
	Panel A: Automation risk (Frey & Osborne, 2017)		
	low	**medium**	**high**
ChatGPT	-0.113***	-0.082***	-0.069***
	(0.024)	(0.010)	(0.016)
N	400,010	1,129,440	2,705,950
	Panel B: Exposure measures (Webb, 2019)		
	B.1 Software exposure		
	low	**medium**	**high**
ChatGPT	-0.105***	-0.066***	-0.066***
	(0.030)	(0.007)	(0.008)
N	1,458,860	1,600,040	1,552,980
	B.2 Robot exposure		
	low	**medium**	**High**
ChatGPT	-0.102**	-0.078***	-0.057***
	(0.025)	(0.018)	(0.006)
N	1,458,860	1,600,040	1,552,980
	B.3 AI exposure		
	low	**medium**	**high**
ChatGPT	-0.102***	-0.062***	-0.072***
	(0.029)	(0.007)	(0.009)
N	1,529,450	1,529,450	1,552,980

Appendix C: Difference in Differences estimation

Figure 5

Difference in Differences (Cognitive vs. Manual)

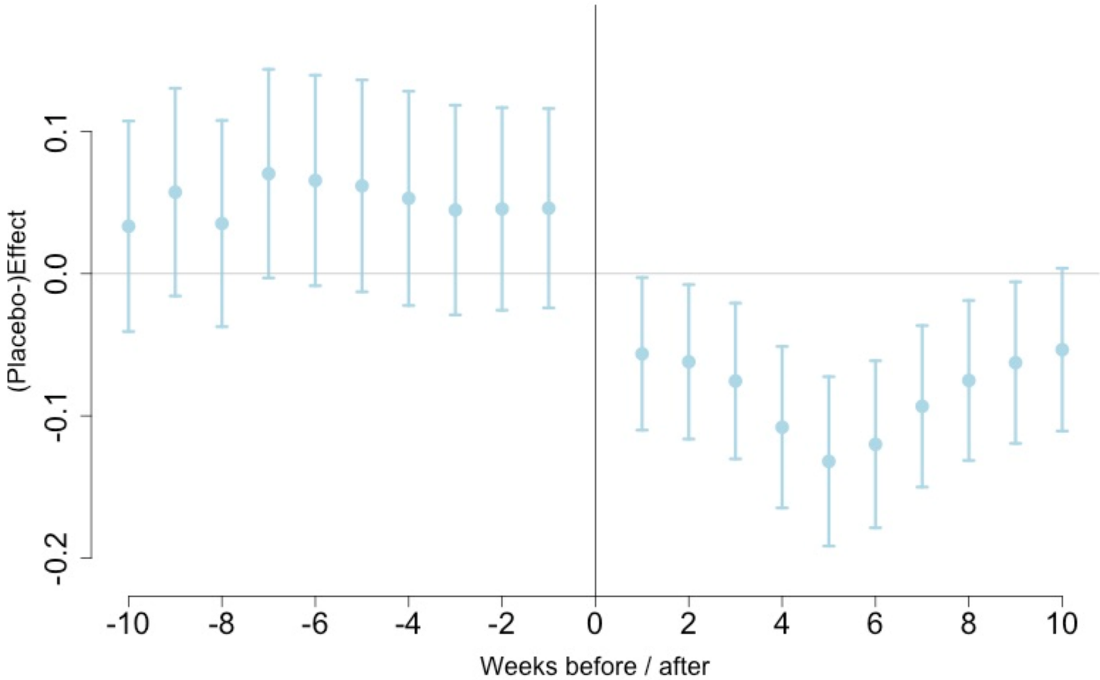

Notes: Linear regression of differences-in-differences between cognitive-type (more affected) and manual-type (not/less affected) jobs. Arguably, the assumption of not affected manual-type jobs is not valid.