Technological Change and Skill Gaps

Authors: Lukas Jonas, Anna Salomons, Cäcilia vom Baur, Ulrich Zierahn-Weilage

Reviewer: Piotr Lewandowski

Grant agreement no. 101132581

Deliverable: D4.10

Due date: 30.6.2026

Funded by the European Union

Document control sheet

Project Number:	**101132581**
Project Acronym:	SkiLMeeT
Work-Package:	4
Last Version:	19.6.2026
Issue Date:	29.6.2026

Classification

Draft	
Final	
Confidential	
Restricted	
Public	x

Keywords: **v***ocational education, technological change, skill mismatch, curriculum updates, job ads*

Legal notice

This project has received funding from the European Union's Horizon Europe Research and Innovation programme under grant agreement No. 101132581. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

Curriculum Updates Reduce Skill Mismatch*

Lukas Jonas
RWI
Ruhr-University Bochum

Anna Salomons
Tilburg University
Utrecht University

Cäcilia vom Baur
Goethe University Frankfurt

Ulrich Zierahn-Weilage
Utrecht University
ZEW

June 19, 2026

Click here for most recent version

Abstract

This report studies whether updating vocational curricula reduces skill mismatch in Germany's dual system. Linking the near-universe of training regulations to 2017–2025 job ads, it constructs an ESCO-based curriculum–job ad mismatch index and exploits curriculum revisions in a stacked difference-in-differences design. Curriculum updates reduce the share of demanded-but-not-supplied skills by about 4 percentage points. This effect is similar in occupations with high versus low technology exposure once controlling for the younger age and lower pre-update mismatch levels of high exposure occupations: Curricula which are more exposed to new technologies require more frequent updates to remain well-matched to demand relative to less exposed curricula. Combined with prior evidence that updates raise wages, the results highlight curriculum revision as a key education-policy lever for mitigating technologically induced skill obsolescence and improving match quality for vocationally trained workers.

Keywords: Vocational education, technological change, skill mismatch, curriculum updates, job ads
JEL codes: J25, J24, O33

*This report is based on the authors' ongoing work. We are currently exploring an extension of the analysis using job ad data from the Institute for Employment Research (IAB) joint with Duncan Roth and Michael Stops – preliminary results are comparable to the results from this report. We thank Ronald Bachmann and Piotr Lewandowski for invaluable feedback. We thank Victor Medina for excellent research assistance. This report is Deliverable 4.10 ("Effect of innovation on skills mismatches") of the SkiLMeeT project and builds on this projects' Deliverable 2.6 (see Salomons et al., 2025). The SkiLMeeT project is funded by the European Union's Horizon Europe Research and Innovation program under grant agreement No 101132581. The content of this report is the sole responsibility of the authors and does not necessarily reflect the opinion of the European Union.

1 Introduction

Rapid technological change is transforming the tasks performed in jobs and the skills that workers require, raising concerns that education and training systems may not adjust quickly enough to prevent growing skill mismatches (see e.g. Spitz-Oener, 2006; OECD, 2025). In many advanced economies, debates about skills gaps and mismatch have become central to discussions of productivity, wage inequality, and the distributive consequences of innovation (see e.g. Vandeplas and Thum-Thysen, 2019), yet there remains limited evidence on how the content of formal education shapes the alignment between workers' skills and labor market demand. A key open question is whether, and to what extent, updating educational content can attenuate the mismatch between the skills supplied by new cohorts of workers and the skills demanded by firms.

We study the role of educational content in reducing job mismatch in the context of Germany's dual vocational training system. Building on our prior work within the SkiLMeeT project (Salomons et al., 2025), we exploit the near-universe of legally binding vocational training curricula over roughly five decades linked to detailed data on occupational exposure to technological change. In this report, we link this database with job ads to study skill demand. The German setting is particularly informative: vocational curricula are standardized nationwide, define in detail the skills and knowledge to be acquired in each training occupation, and are periodically updated through an institutionalized process involving social partners and the federal training authority. Since vocationally trained workers account for about 70 percent of the German workforce over our period of interest, curriculum content provides a rich and policy-relevant measure of skill supply.

Our analysis builds on an explicit notion of skill mismatch that goes beyond traditional measures of underqualification. Rather than inferring mismatch from workers' formal education levels or broad skill categories, we measure the alignment between multidimensional bundles of skills supplied through vocational curricula and those demanded in job ads. We conceptualize mismatch as the share of skills that employers require in a given occupation but that are not covered in the corresponding curriculum, an underqualification margin defined at the level of individual skills.[1] This perspective recognizes that curricula shape both occupation-specific and more general skills and that match quality matters not only

[1]We refrain from comparing skills supplied by curricula to skills demanded in job postings to measure overqualification because employers may exclude skills supplied in curricula not because these are not required, but because employers expect workers to posses those skills due to their vocational training.

for workers' current jobs but also for their outside options in a changing labor market.

To implement this approach empirically, we construct skill demand profiles from a large-scale dataset of job ads for Germany covering 2017-2025, which we link to vocational training curricula at the four-digit level of the German classification of occupations, KldB *(Klassifikation der Berufe).* For each posting, we focus on the text zones describing job activities and qualification requirements, from which we extract skills using the *European Skills, Competences, Qualifications and Occupations* (ESCO) taxonomy as a common reference for both vacancies and curricula. This allows us to construct comparable occupation-year vectors of skills supplied through vocational training and skills demanded in vacancies. A central methodological challenge is to reliably identify ESCO skills in unstructured text from both vacancies and curricula. We address this using two complementary extraction strategies: a keyword-based approach and an embedding-based approach. Using the extracted skills, we define a curriculum–job ad mismatch index for each occupation and year as the share of demanded ESCO Level 2 skills that are not supplied in the corresponding vocational curriculum.

We then exploit the timing of curriculum updates to study how changes in educational content affect skill mismatch, with particular attention to the role of new technologies. Following Salomons et al. (2025), we assemble a stacked event-study dataset of curriculum update episodes at the four-digit KldB level, in which treated occupations that experience a curriculum update are compared to control occupations without updates in a symmetric window around each event.

Our main result is that curriculum updates significantly reduce skill mismatch. On average, an update lowers the share of demanded-but-not-supplied skills in the treated occupation by about 4 percentage points relative to comparable occupations that do not experience an update in the same period. This indicates that revisions of vocational training regulations are effective in bringing curriculum content closer to the skill requirements observed in the labor market.

We further find that this the update-induced decline in mismatch is slightly larger for occupations highly exposed to technological change. Technological change causes curricula to receive updates more frequently (see Salomons et al., 2025). Accordingly, occupations with a high exposure to new technologies are younger and have lower levels of mismatch when updated. Once we control for differences in curriculum age and pre-update levels of mismatch, the effect of curriculum updates on mismatch is comparable between occupations with a high and low exposure to new technologies. This indicates that technology exposure requires vocational training curricula to be updated more frequently to keep pace with technological change and to remain well-matched with demand embedded in job ads when compared to curricula with low technology exposure.

These findings have several implications. First, they provide direct evidence that adjustments in the content of vocational education can mitigate skill mismatches generated by changing technologies, complementing existing work that focuses on wages, employment, or broad qualification mismatches. Second, they highlight that the timing and targeting of curriculum updates are particularly important in technology-intensive occupations, where the payoff from aligning training content with evolving skill demand is especially high. Third, combining our results with prior evidence that curriculum updates raise wages, especially in tech-exposed curricula, suggests that reducing mismatch through educational reforms can translate into tangible gains for workers. Our evidence underscores the central role of education systems, and vocational curricula in particular, in shaping how labor markets absorb new technologies and in determining who benefits from innovation.

Our work relates to four strands of the literature. The first strand examines skill mismatch and job displacement in the context of technological and structural change, emphasizing the costs of misaligned skill bundles for workers and firms. Recent work develops new measures of mismatch between workers' skills measured via surveys to job requirements measured by job postings or occupational job descriptions. These papers show that misalignment is pervasive and a key determinant of wages, employment stability, and worker mobility (Guvenen et al., 2020; Fredriksson et al., 2018; Guo et al., 2020; Neffke et al., 2024). Related studies document that displacement and mass layoffs often push workers into jobs with worse match quality or longer non-employment spells, with consequences that depend critically on the portability and obsolescence of their skills (Andersson et al., 2018; Gathmann et al., 2018; Hanushek et al., 2025; Tian, 2025; Ziya Sen, 2025). Complementing this evidence, a growing literature links changes in skill demand over the technology lifecycle to workers' career trajectories and earnings risk, highlighting that skill mismatch can emerge both from shifting demand and from the slow adaptation of educational content (Deming and Noray, 2020; Atalay et al., 2020; Braxton et al., 2021; Agrawal et al., 2025; Storm, 2023). We contribute to this literature by providing a direct, multidimensional, occupation-level measure of curriculum–job ad mismatch in vocational skills, by exploiting the universe of occupational skills covered in ESCO, and by showing that institutionalized curriculum updates in Germany's dual system meaningfully reduce the share of skills demanded in vacancies but not supplied in training curricula. Our findings therefore connect the mismatch literature to an explicit, policy-relevant lever on the skills supply side and clarify how educational reform can mitigate the costs of technologically induced skill obsolescence and displacement.

Second, a long-standing economic literature studies the role of new technologies for the demand for skills, tasks, and occupations with a particular focus on wages and wage inequality (Katz and Murphy, 1992; Autor et al., 2003; Goldin and Katz, 2008; Acemoglu and Restrepo, 2018, 2019, 2022; Autor et al., 2020; Restrepo, 2024). New technologies reshape the demand for labor by simultaneously displacing workers from existing tasks and complementing them in others, while generating demand for new labor-using tasks. This changes the composition of demanded skills and the wage structure. We contribute by shifting attention from the demand side to the supply side of skills: rather than asking only how technology reshapes tasks performed in jobs, we study how formal educational content responds and whether curriculum updating reduces the gap between the skills workers are taught and the skills firms demand.

A third strand analyzes within-occupation changes in tasks and skill demand, often using job ads or other text-based measures of work content (Spitz-Oener, 2006; Atalay et al., 2020; Deming and Noray, 2020; Acemoglu et al., 2020; Autor et al., 2024; Hampole et al., 2025). This literature shows that a substantial share of labor-market change occurs within occupations rather than only through reallocation across occupations, and that online vacancies are a powerful source for measuring shifting skill requirements and the emergence of new tasks. This report contributes by building the supply-side counterpart to that demand-side approach. We construct multidimensional measures of skills supplied through vocational curricula, compare them directly with job ad-based demand profiles at the occupational level, and use that comparison to quantify occupation-year skill mismatch and the role of curriculum updates in reducing it.

A fourth strand examines skill obsolescence, educational content change, and worker adjustment in the face of technological progress (MacDonald and Weisbach, 2004; Janssen and Mohrenweiser, 2018; Deming and Noray, 2020; Fillmore and Hall, 2021; Kogan et al., 2023; Dillon et al., 2025). This literature finds that when technology or training content changes, incumbent workers may suffer from obsolete skills, while workers with more up-to-date skill bundles may benefit. Related recent studies show that educational programs and curricula adjust to changing labor demand and that such adjustments matter for later outcomes (Boustan et al., 2022; Biasi and Ma, 2023; Conzelmann et al., 2023; Di Giacomo and Lerch, 2023; Light, 2024; Salomons et al., 2025). Our report contributes to this debate by providing a systematic measure of curriculum-job ad mismatch across the near-universe of German dual vocational training occupations, and by showing that curriculum revisions significantly reduce the share of demanded-but-not-supplied skills. Our report identifies curriculum updating as a concrete institutional mechanism through which education systems can mitigate technologically induced skill obsolescence and mismatch.

The rest of this report is structured as follows. We first introduce our vocational training and job ad datasets in Section 2. Section 3 discusses our approach to measure skill mismatch by comparing skills supplied via vocational training content to skills demanded via job ads. Using this data and skill mismatch indicator, we study the effects of curriculum updates on skill mismatch with a particular focus on the role of new technologies in Section 4. Section 5 concludes.

2 Data

2.1 Vocational Training Curricula

The vocational training curriculum data used in this project are taken from Salomons et al. (2025), and our discussion here borrows from that work. That study constructs a comprehensive database of legally binding vocational training curricula in Germany, covering the near-universe of dual-system training occupations over roughly five decades (1971–2021). For this project, the data has been updated to cover the period up to 2025. The curricula are obtained by web-scraping the archives of the Federal Law Gazette (Bundesgesetzblatt), which publishes nationally standardized training regulations following the 1969 Vocational Training Act. These regulations define, for each training occupation, the title, program duration, detailed lists of skills and knowledge to be acquired, a training framework specifying their sequence, and the requirements for final examinations. The resulting texts are extensive, averaging more than ten pages per curriculum, and provide rich detail on the content of formal skill acquisition within the German dual system.

To build the curricula dataset, each scraped training regulation is organized at the level of "training occupations", defined by unique occupation titles, and then linked to their corresponding training occupations. These training occupations are then linked over time using the official "Index of Recognized Training Occupations", which records when curricula are updated, renamed, aggregated, or split. This matching step produces a panel of training occupations by year, 1971–2025, that indicates for each occupation when a new curriculum version is introduced and how it relates to its predecessor. Finally, training occupations are mapped to the German occupational classification KldB 2010 at the four-digit level (KldB-4)using a crosswalk provided by the Federal Institute for Vocational Education and Training (BIBB), resulting in 236 distinct KldB-4 occupations for which curriculum content can be observed and linked to labor-market and innovation measures.

The German dual vocational system makes curricula particularly suitable as a measure of skill supply at the occupational level. Vocational training typically follows secondary school and combines part-time schooling (1–2 days per week) with on-the-job training in firms (3–4 days per week), usually over a three-year period. The training content is codified in legally binding federal regulations and applies nationwide, ensuring that curricula are standardized and regularly revised through an institutionalized process involving employer associations, unions, and the Federal Institute for Vocational Education and Training. This institutional setting implies that, for almost all dual-system training programs, one observes both the timing of each curriculum update and the detailed skill content taught to each cohort of apprentices, including which "vintage" of skills each cohort acquires. Since vocationally trained workers make up around 70% of the German workforce over the period considered, these curricula capture skill supply adjustments for a large segment of the labor market.

Curriculum updates are frequent but heterogeneous across occupations and over time. On average, a given training occupation experiences a curriculum change in about 3.6% of occupation-year observations, corresponding to an update interval of roughly 20 years. Many updates concern the content of the curriculum without altering the number or structure of training occupations, while others involve renaming, merging, or splitting existing occupations. Descriptive evidence shows that updates are especially frequent for occupations in IT and scientific services and in business services, whereas personal services tend to be updated less often. Qualitative examples of specific curricula highlight substantive changes in skills over time: for instance, updates in industrial and clerical occupations shift content toward digital systems, standard and firm-specific software, networking, customer support, and teamwork, while reducing emphasis on purely mechanical or routine tasks. These patterns illustrate how curricula adjust the skill bundle supplied within occupations in response to evolving technologies and work organization.

For the purposes of this project, the vocational curricula database serves as the supply-side counterpart to the occupational skill demand profiles derived from job ad data. Each training occupation, mapped to a KldB-4 code, provides a structured text describing the skills imparted during formal training.

2.2 Job Ads

The analysis draws on a large-scale dataset of job ads for Germany covering the period 2017–2025. The primary data are collected by Palturai GmbH / Finbot AG, which systematically scrape vacancies from major job boards (such as Stepstone, Monster and stellenmarkt.de), company websites, temporary employment agencies and head-hunter agencies. Figure 1 highlights that most vacancies stem from major job boards. After collection, the raw postings are processed to improve quality and usability: duplicates are removed, geographic location is added, firms are identified, and each posting is assigned an occupation code according to the German classification of occupations KldB 2010 at the four-digit (KldB-4) level. Furthermore, the textual content of each posting is segmented into distinct zones, namely (i) job activities, (ii) qualifications, (iii) benefits, and (iv) company description, which facilitates targeted extraction of skill information.

Figure 1: Job Ads by Source

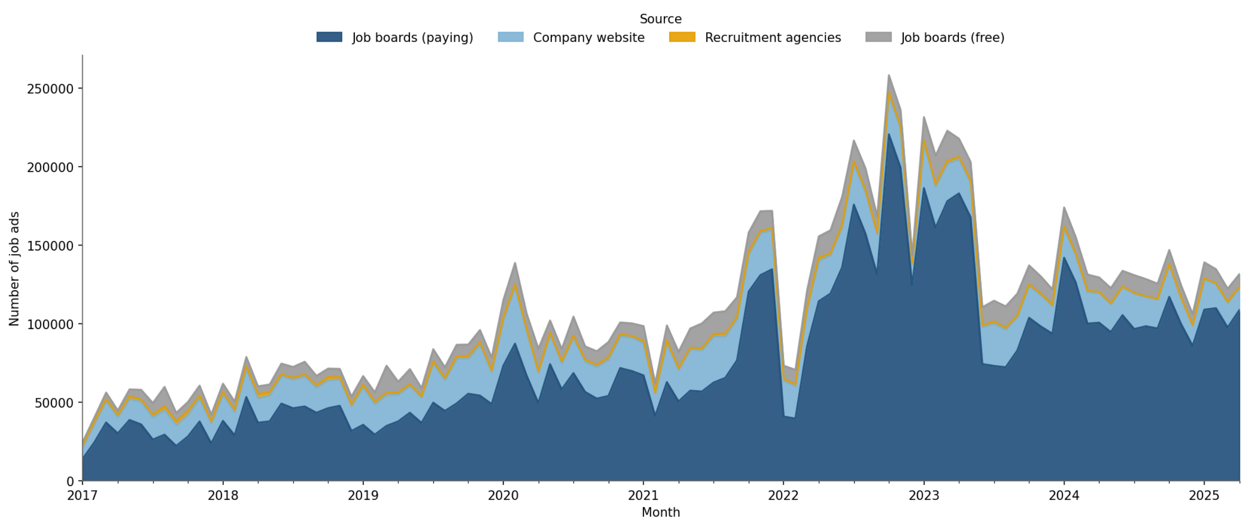

Monthly counts of online job ads by source group for the used KldB-4 occupations, 2017–2025.

For the purposes of this report, several sample restrictions are applied to align the job ad data with the available information on vocational education curricula. First, only postings related to KldB-4 occupations for which corresponding vocational curricula are observed are retained, ensuring that demand in vacancies can be systematically matched with skills supplied through training programs. Second, among these occupations, only those with at least 10 postings per year and at least 100 postings in total over the observation period are kept, which yields a working sample of 185 occupations and approximately 6 million job ads.

To focus on standard employment relationships, vacancies for apprenticeships, traineeships, internships, and other forms of irregular work, as well as postings from temporary work agencies, are excluded. When a job ad is linked to multiple KldB-4 occupations, the posting is duplicated across these occupation codes so that demand is fully represented at each relevant occupational level. Descriptive statistics in Table 1 indicate substantial variation in the number of postings across occupations as well as in average posting length (measured in tokens), which underlines the importance of working with occupation-level aggregates rather than individual vacancies.

Table 1: Descriptives on job ads

	Number of job ads per occupation	Average tokens per job ad
Min	300	83.91
p25	1,666	102.27
Median	5,808	121.95
Mean	32,270	125.67
p75	22,754	143.07
Max	655,098	226.36

Descriptive statistics on the number of job ads and the average posting length (in tokens), across the KldB-4 occupations.

In subsequent steps of the analysis, the information in the job ads is used to infer skill demand at the occupational level. Skill extraction is performed only on the "job activities" and "qualifications" zones of each posting, where skill-relevant content is most likely to appear. For each KldB-4 occupation, the extracted skill signals are aggregated across postings to construct an occupation-level demand profile that can be compared with the skill content of vocational curricula.

2.3 Technological Change Embedded in Patents

To capture occupational exposure to technological change, our analysis draws on a patent-based indicator constructed in Salomons et al. (2025). They link the German vocational training curricula to so-called breakthrough innovations identified in U.S. utility patents. The resulting measure provides, for each curriculum (and thus each linked KldB-4-digit occupation), an intensity of past exposure to transformative digital technologies. In this report, we take the indicator as given from Salomons et al. (2025) and merge it to our curricula–occupation panel. Below, we provide a description of the process and indicator for the purpose of this report.

The underlying patent data consist of U.S. utility patents, with technological breakthroughs defined following Kelly et al. (2021) as patents that are both novel and highly influential for subsequent innovation. Specifically, breakthroughs are the top 10 percent of patents by year in terms of forward-to-backward textual similarity, and are observed over 1946-2001. Salomons et al. (2025) focus on the subset of breakthroughs in digital technologies (the "Instruments/Information" class), which have expanded strongly over time and capture the core of the digital revolution. To address concerns about simultaneity and reverse causality, these patents are lagged by 20-25 years relative to the curriculum years over 1971-2021, so that current training content is related only to earlier waves of major technological advances.

To link curricula to patents, Salomons et al. (2025) use Natural Language Processing on the full text of machine-translated curriculum documents and patent descriptions. After standard text preprocessing and the construction of *Term Frequency-Inverse Document Frequency* (TF-IDF) weighted word-embedding vectors, cosine similarity is calculated for every curriculum-patent pair and normalized at the patent level by subtracting the median similarity of that patent across all curricula. For each curriculum, the 15 most similar lagged breakthrough patents are retained, and their (demeaned) similarity scores are summed to obtain a curriculum-specific patent count. This patent count, computed separately for all breakthroughs and for the subset of digital breakthroughs, is interpreted as the curriculum's exposure to technological change embodied in prior innovations.

The technology exposure measure used in this report is the digital breakthrough patent indicator from Salomons et al. (2025), aggregated to the KldB-4-digit level via the mapping between training occupations and KldB occupations described in Section 3.1. It is defined as the log count (or, in robustness checks, the level) of textually linked digital breakthrough patents in the 20-25 years preceding the curriculum year. This indicator varies both across occupations and over time. In the empirical analysis below, this patent-based technology exposure enters as an exogenous measure of long-run technological pressure on occupations; all details on its construction, validation, and robustness are taken from Salomons et al. (2025) and are not re-estimated here.

3 Skill Mismatch Approach

3.1 Overall Approach

The core idea of the skill mismatch methodology is to compare skills demanded in job ads with skills supplied in the vocational curricula associated with the same KldB-4 occupation. The key problem is identifying a common set of skills. We rely on the ESCO skill taxonomy as a reference list of skills, and aim to identify these skills in both, the job ads and the vocational training curricula. Figure 2 provides an overview of how we extract ESCO skills from curricula and job ads to generate a skill mismatch indicator. As a robustness check, we additionally restrict the sample to skills that appear in at least 5 percent of postings within a given KldB-4 occupation. This is intended to reduce the influence of rare or idiosyncratic mentions, so that the derived demand measures are less likely to be driven by noise rather than systematically required skills. Reassuringly, our main results are robust to this restriction.

Skills are identified using the ESCO skill taxonomy, drawing on both the most detailed (Level 1) and the second-most detailed (Level 2) descriptions. Skills are extracted separately from the job ad data and from the text of the vocational curricula, and then aggregated to Level 2 for the main analysis. For the job ad data, extraction is based on the job activities and qualifications zones, while for curricula the relevant sections describing the skills and knowledge in detail are used.

We rely on two complementary approaches to extract skills from text: (1) a keyword-based approach and (2) an embedding-based approach. (1) The keyword approach relies on an expanded list of ESCO-related keyphrases, where Level 1 skill phrases are enlarged through synonym substitution and simplification of longer phrases, while Level 2 phrases remain unchanged. After standard preprocessing steps (removing stopwords and lemmatization), a skill is considered present if all words of a keyphrase appear within a 5-token window (five tokens before and after the first matched word), irrespective of word order. This approach is designed to minimize false positives but may miss skills that are expressed in less standard wording, thereby raising the risk of false negatives. (2) The embedding approach instead embeds both chunks of the input text and ESCO skills using a job-specific BERT model (Gnehm et al., 2022) and identifies a skill as present in a text chunk if the similarity between the chunk and the skill embedding exceeds a threshold of 0.92. This method has the opposite strengths and weaknesses: it is more flexible and can capture paraphrased skill ex-

Figure 2: Skill Mismatch Methodology

A. Keyword approach

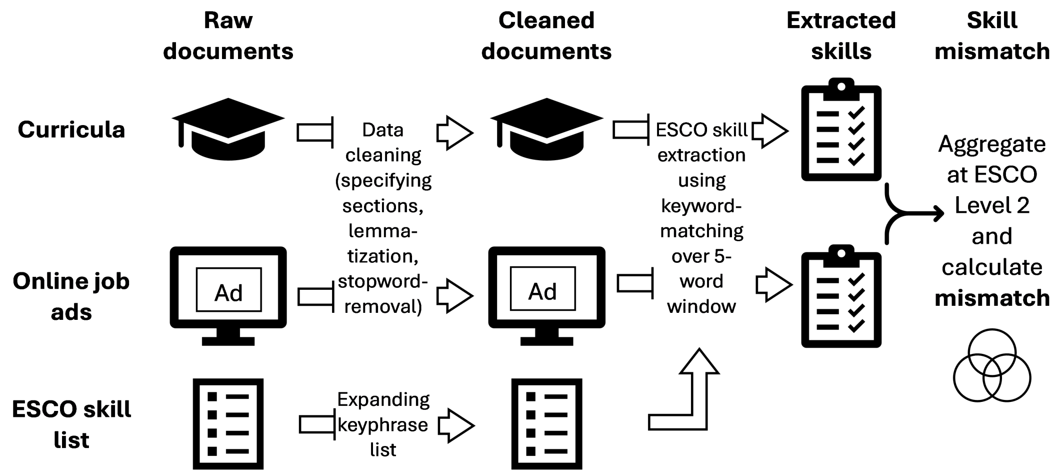

B. Embedding approach

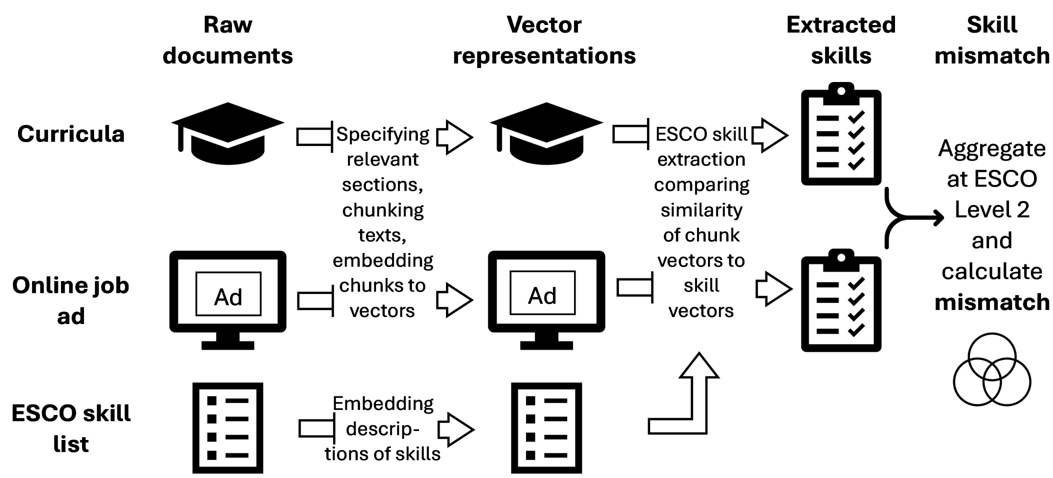

Overview of the two skill-extraction pipelines. Panel A (keyword approach) and Panel B (embedding approach) both extract ESCO skills from curricula and job ads, aggregate them to ESCO Level 2, and compute the mismatch index.

pressions, but may be more susceptible to false positives. Validation exercises indicate that, overall, the embedding approach outperforms the keyword approach in this application.

Using the extracted skills, occupation-level demand and supply vectors are constructed for each KldB-4 occupation and year. The mismatch measure focuses on underqualification, defined as skills demanded in job ads but not supplied in the corresponding vocational curriculum.[2] For each occupation and year, the underqualification index is computed as the number of demanded but not supplied Level 2 skills divided by the total number of demanded Level 2 skills in that occupation-year:

$$\text{Mismatch}_{it} = \frac{\text{Number of demanded but not supplied L2-skills}_{it}}{\text{Total demanded L2-skills}_{it}} \tag{1}$$

where i indexes KldB-4 occupations and t indexes years. This yields a mismatch indicator between 0 and 1, with higher values indicating a larger share of demanded skills that are not covered in the relevant curricula.

To support robustness and interpretation, the methodology also considers non-thresholded variants of the measures (e.g. retaining all skills, not only skills that meet minimum frequency conditions within an occupation) and evaluates mismatch at both Level 1 and Level 2 of the ESCO taxonomy.

In the rest of this paper, we will not interpret the level of our mismatch indicator, but instead focus on changes in the indicator over time or differences in mismatch between occupations. This is, among others, because firms may not mention skills that they assume workers possess because of their training; and because the level of mismatch differs depending on which level of ESCO skills is used, whether we apply thresholds to include skills, or whether we rely on a keyword versus embedding approach. We hence rely only on changes in mismatch over time and differences in mismatch between occupations, which are not sensitive to this.

3.2 Keyword versus Embedding Approach

The mismatch indices are computed separately for the keyword-based and embedding-based extraction, and summary statistics suggest relatively high shares of demanded-but-not-

[2] In principle, we could also measure obsolete skills by computing the share of skills supplied in curricula that are not mentioned in job postings. However, if firms know the content of vocational training curricula, they may simply assume that workers who completed a vocational training track possess those skills and hence not ask for it in their job postings, limiting the interpretability of such a measure.

supplied skills, with the embedding-based measures generally showing better performance in skill detection. In particular, Table 2 provides summary statistics for the matched skills per KldB-4 occupation from the job ads in the keyword approach compared to the embedding approach. As shown in column 1, the share of job ads without any extracted skill is considerably smaller under the keyword approach. Despite this, the embedding approach extracts substantially more level 1 skills (column 2), and the same holds for directly extracted level 2 skills (column 3). Column 4 reports the combined number of extracted level 2 skills, comprising both those directly extracted and those obtained by aggregating level 1 skills to level 2. The final two columns show the number of skills remaining after applying a frequency threshold, retaining only skills that appear in more than 5 percent of job ads within an occupation — at level 1 (column 5) and level 2 (column 6), respectively. On average, the keyword approach yields 31 ESCO level 2 skills per KldB-4 digit occupation, compared to 23 for the embedding approach.

Table 3 provides comparable statistics on the skill extraction from vocational training curricula at the KldB-4 digit level. The embedding approach clearly outperforms the keyword approach in identifying ESCO skills within the curricula. These results highlight that curricula typically contain a more concentrated, standardized set of skills, whereas vacancies display a broader range of skill requirements.

Table 4 provides descriptive statistics on our mismatch indicator, the share of skills demanded in the job ads but not supplied via the vocational training curricula at the level of KldB-4 digit occupations. At the more detailed level of ESCO skills (Level 1), the share of skills demanded but not supplied is much higher, as it is harder to directly identify fine-grained skills from job ads on training curricula. Mismatch is significantly lower when using a broader classification of skills at ESCO level 2. This is true for both, the keyword and the embedding approach. Mismatch is on average lower for the embedding approach, as the embedding approach does not require words to match but instead allows for more differences in languages between job ads and curricula texts. Taking into account a minimum threshold for similarity between job ads and ESCO skills further reduces skill mismatch, as we then rely only on skills that have sufficient support in the job ad data. Accordingly, for our further analyses, we rely on the embedding approach at ESCO level 2 with a minimum threshold for skill similarity.

To test the reliability of our skill mismatch indicator, we compute the cosine similarity between the skills demanded and skills supplied across all KldB-4 occupations. Figure 3 shows that skills demanded and skills supplied are more similar within occupations than

Table 2: Skill Extraction from Job Ads per KldB-4

(a) Keyword Approach

	% no skill	L1 skills	L2 skills (direct)	L2 skills (combined)	L1 skills (thr.)	L2 skills (thr.)
Min	0.014	179	5	130	11	11
p25	0.057	424	17	218	23	25
Median	0.079	706	25	269	28	30
Mean	0.087	870.5	29.65	272.4	28.18	30.56
p75	0.110	1152	40	326	33	35
Max	0.246	2956	90	419	47	49

(b) Embedding Approach

	% no skill	L1 skills	L2 skills (direct)	L2 skills (combined)	L1 skills (threshold)	L2 skills (threshold)
Min	0.089	256	16	117	0	1
p25	0.171	920	61	235	8	13
Median	0.253	1647	91	289	13	21
Mean	0.281	2049	103.7	286.4	14.14	22.41
p75	0.363	2769	137	345	19	29
Max	0.790	7144	259	431	42	56

Legend: (1) % no skill – share of job ads without any extracted skill; (2) L1 skills – number of extracted Level 1 skills; (3) L2 skills (direct) – directly extracted skills at Level 2 (i.e. exact match); (4) L2 skills (combined) – combined number of extracted level 2 skills from direct Level 2 extraction and from aggregating Level 1 to Level 2 skills; (5) L1 skills (threshold) – number of Level 1 skills remaining after applying threshold; (6) L2 skills (threshold) – number of Level 2 skills remaining after applying threshold.

Table 3: Skill Extraction from Curricula per KldB-4

(a) Keyword Approach

	L1 skills	L2 skills (direct)	L2 skills (combined)
Min	25.00	0.000	23.00
p25	43.25	2.000	40.00
Median	53.00	2.000	49.00
Mean	59.05	2.632	52.37
p75	71.00	3.000	63.00
Max	122.00	6.000	96.00

(b) Embedding Approach

	L1 skills	L2 skills (direct)	L2 skills (combined)
Min	11.0	0.00	8.00
p25	63.0	4.25	47.25
Median	109.0	9.50	72.00
Mean	154.4	13.99	83.26
p75	192.0	19.00	110.00
Max	563.0	58.00	216.00

Legend: (1) L1 skills – number of extracted Level 1 skills; (2) L2 skills (direct) – directly extracted skills at Level 2 (i.e. exact match); (3) L2 skills (combined) – combined number of extracted level 2 skills from direct Level 2 extraction and from aggregating Level 1 to Level 2 skills.

Table 4: Share of Skills Demanded but not Supplied

(a) Keyword Approach

	Level 1	Level 1 (thr)	Level 2	Level 2 (thr)
Min	0.789	0.200	0.609	0.156
Mean	0.928	0.661	0.825	0.570
Max	0.991	0.905	0.966	0.913

(b) Embedding Approach

	Level 1	Level 1 (thr)	Level 2	Level 2 (thr)
Min	0.679	0.000	0.283	0.000
Mean	0.945	0.689	0.753	0.356
Max	1.000	1.000	0.985	1.000

The table reports descriptive statistics for mismatch, defined as the share of demanded skills (in job ads) which are not supplied (in the curriculum of the corresponding occupation). The statistics are separate for mismatch defined using Level 1 versus Level 2 ESCO skills, and separate for an approach where we use all skills versus where we apply a minimum threshold “(thr)”. In the threshold version, we only consider skills that appear in a minimum share of job ads per occupation.

between occupations. This is reassuring as it indicates that we do not identify generic skills, but extract matching occupation-specific information from both job ads and vocational training curricula.

Figure 3: Skill Similarity between Supply and Demand: Within versus Between Occupations.

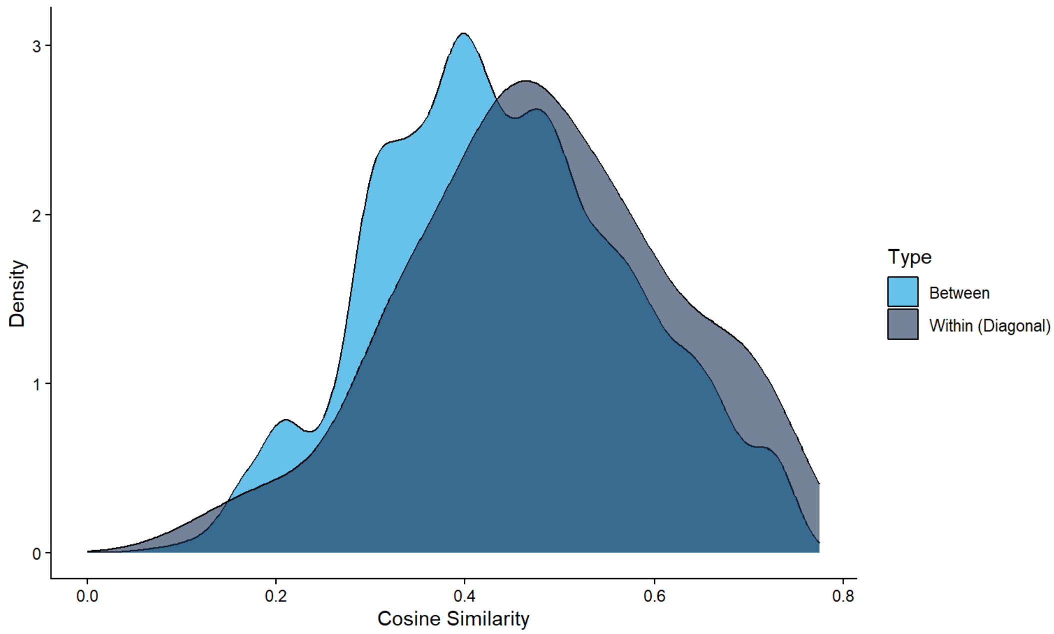

The Figure plots the density of the cosine similarity of supplied skills (measured based on all curricula active in 2025) and demanded skills (grouped across all available years). Both are measured on level 2.

We additionally compare the extracted demanded and supplied skills to those listed in the ESCO occupations pillar as being relevant for a specific occupation in Table 5. To do so, we aggregate the listed relevant skills of all ESCO occupations belonging to a KldB-4 occupation and calculate the share of those skills that we extracted. We calculate the share of extracted ESCO skills in all occupational skills separately for the demand side (job ads) and the supply side (curricula). With our main approach we capture a significant share of the relevant skills in both the curricula and the job ads.

4 Curriculum Updates and Mismatch

Based on our skill mismatch and curriculum data, we analyze how curriculum updates affect skill mismatch, with a particular focus on the role of new technologies. To descriptively assess the potential role of curriculum updates for mismatch, we rely on our aggregated

Table 5: Coverage of occupation-specific ESCO skill lists by extracted skills

	Level 1	Level 1 (thr)	Level 2	Level 2 (thr)
Curriculum-based skill extraction	0.033	0.033	0.352	0.352
Job-ad-based skill extraction	0.306	0.007	0.900	0.119

The table reports the share of extracted ESCO skills in all ESCO skills of an occupation, separately for skills extracted via supply (i.e. curricula) and via demand (i.e. job ads). The statistics are separate for mismatch defined using Level 1 versus Level 2 ESCO skills, and separate for an approach where we use all skills versus where we apply a minimum threshold "(thr)". In the threshold version, we only consider skills that appear in a minimum share of job ads per occupation.

measure of skill demand and regress mismatch (i.e. the share of skills demanded but not supplied) by occupation and year on the time since the last update and occupation fixed effects. Figure 4 plots the results. As curricula get older, mismatch tends to increase – the share of demanded skills actually supplied declines. The coefficients turn downward for the oldest two year groups, but these account for only a small share of observations and may reflect outliers, selection, or reverse causality, rather than a genuine reversal. Overall, the pattern suggests that curriculum updates could potentially reduce skill mismatch.

To estimate the effect of curriculum updates on skill mismatch, we follow Salomons et al. (2025) and create a stacked dataset of curriculum update events to implement a stacked differences-in-differences analysis. Treatment is staggered because curricula receive updates in different years. Whenever a curriculum within a KldB-4 digit occupation is updated, we consider that as an update event. We created a stacked dataset of events as follows: we select all instances where within a KldB-4 digit occupation a curriculum is updated. For that treated occupation, we merge all control occupations that have not seen an update 5 years before or 5 years after the update year of the updated occupation. We stack these events. This implies that the same occupation can serve as a control occupation for different events.

Our key indicator is skill mismatch, defined as the share of demanded skills which are not supplied in the vocational training curricula. We rely on the embedding approach with ESCO level 2 skills due to the reasons outlined, above.[3] We rely on a time-constant definition of skill demand by defining demanded skills using all years from the job ad data. Time variation then solely stems from changes in skill supply via changes in curriculum content over time due to updates. In addition, we know technology exposure of each vocational training curriculum.

[3]We repeat our analyses using the alternative approaches outlined above and find comparable results.

Figure 4: Mismatch by time since last update

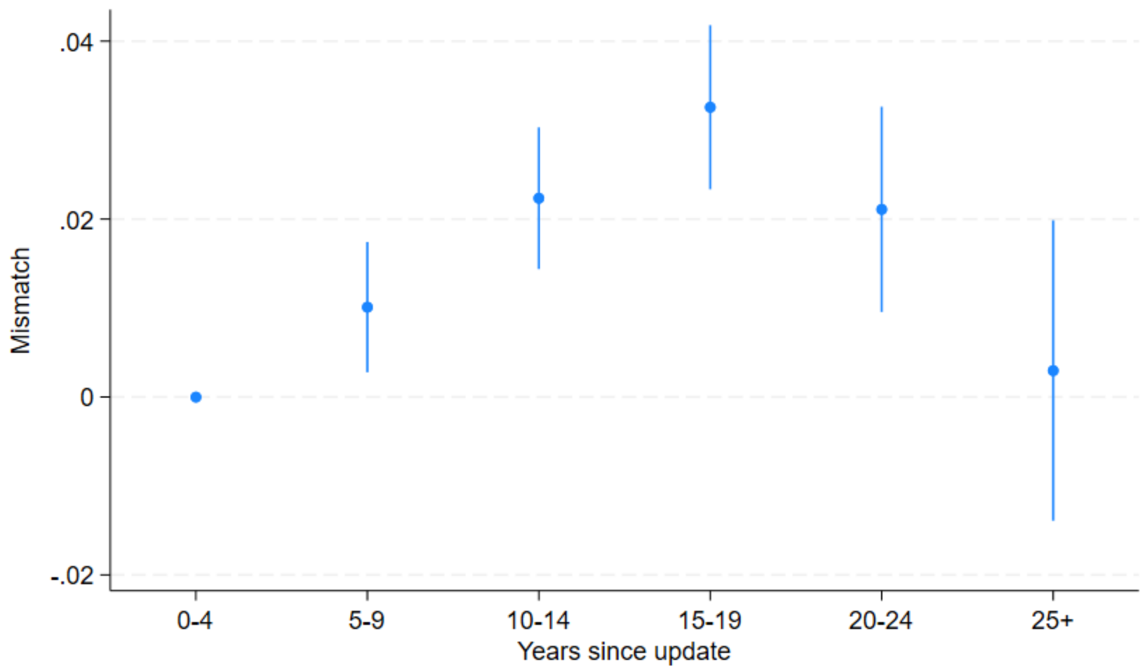

The Figure plots coefficients from a regression of mismatch (share of demanded skills which are not supplied in the vocational training curricula) on years since the last curriculum update (grouped) and broad occupation fixed effects. Bars represent 95 % confidence intervals.

This indicator is depicted from Salomons et al. (2025) and is based on the number of linked digital breakthrough patents from 20 to 25 years ago. We also know the time since the last update of each curriculum.

We implement the following stacked differences-in-differences estimation building on Salomons et al. (2025):

$$\text{skill mismatch}_{it} = \beta \text{Post}_t \times \text{Treated}_i + \psi_i + \phi_t + u_{it}, \tag{2}$$

where post is a treatment dummy (0 pre-update, 1 after an update), and treated is a dummy indicating which occupation is treated (=1) and which is not (=0). We include occupation i and time t (year relative to treatment year) fixed effects. Due to the stacking of the data, i indexes individual occupations by event, and t indexes time by event. Additionally, we control for the time since the last update and the mismatch level prior the update by interacting these controls with the post dummy. Standard errors are clustered at the event by occupation level.

Table 6 provides the key result: curriculum updates reduce skill mismatch by about 4 percentage points (about 5 percent) in treated occupations compared to comparable occupations without an update across all specifications. This highlights that updates of vocational training curricula are successful in bringing skills supplied closer to skills demanded on the labor market.

Table 6: Effect of Curriculum Updates on Skill Mismatch

	(1)	(2)	(3)	(4)
Post × Treated	-0.0389***	-0.0389***	-0.0398***	-0.0403***
	(0.0035)	(0.0034)	(0.0035)	(0.0034)
Pre mismatch		✓		✓
Time since last update			✓	✓
N	36362	36362	36362	36362
Mean of outcome				0.764
SD of outcome				0.134

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table reports stacked differences-in-differences estimation using time-constant definition of skill demand from job ads but time-varying skill supply from curricula. Controlling for occupation and relative time. All controls interacted with event dummies. Standard errors clustered at the event times occupation level.

This effect is comparable between occupations that are more exposed to technological change and those that are less exposed: we split our sample into occupations with an above versus below median exposure to new technologies. Table 7 provides the results. The effects are slightly larger in absolute terms for occupations more exposed to new technologies (a greater reduction in mismatch), though this difference does not appear to be statistically significant. Occupations with a high exposure to new technologies are updated earlier (lower time-since-update when updated) and haver a lower pre-update level of mismatch. Hence, the decline in mismatch is notably larger in relative terms of occupations with a high exposure to new technologies given their lower baseline mismatch.

Salomons et al. (2025) show that curricula are more likely to be updated, and are updated earlier on when they are more exposed to new technologies, resulting in curricula for highly exposed occupations being younger. These curricula also have lower levels of mismatch. This is why we add pre-update-mismatch and time since the last update as control variables. Once we add these controls, effect sizes become more similar between occupations with a high and those with a low exposure to new technologies. The slightly larger update-induced decline in mismatch for occupations exposed to technological change can thus be explained by these occupations receiving updates earlier (due to the new technologies) and having lower pre-update levels of mismatch as they are younger. This implies that new technologies require curricula to be updated more frequently for those curricula to stay well-matched to demand embedded in job ads.

Table 7: Effect of Curriculum Updates on Skill Mismatch by Tech Exposure

	High tech occs				Low tech occs			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × Treated	-0.0404***	-0.0416***	-0.0395***	-0.0408***	-0.0375***	-0.0350***	-0.0411***	-0.0391***
	(0.0051)	(0.0050)	(0.0049)	(0.0048)	(0.0047)	(0.0047)	(0.0051)	(0.0050)
Pre mismatch		✓		✓		✓		✓
Time since update			✓	✓			✓	✓
N	17870	17870	17870	17870	18492	18492	18492	18492
Pre update means:								
Mismatch	0.733				0.799			
SD	(0.119)				(0.106)			
Time since update	12.705				14.482			
SD	(7.220)				(7.156)			

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table reports stacked differences-in-differences estimation using time-constant definition of skill demand from job ads but time-varying skill supply from curricula. Controlling for occupation and relative time. All controls interacted with event dummies. Standard errors clustered at the event times occupation level. Table split by tech exposure (above/below median exposure).

Salomons et al. (2025) further show that curriculum updates raise wages, and they raise wages particularly for tech-exposed curricula. This indicates that the new skills included in the new curricula are particularly valuable when the curriculum is exposed to new technologies. We conclude that curriculum updates bring the skills supplied by training curricula closer to demand — and that this holds both for occupations that are exposed to new technologies, and occupations that are not. The key difference between curricula highly exposed to technology, and those with low exposure, is that the former require more frequent updates to achieve skills that match well with labor market demands. Curriculum updates thus are a key requirement for well-matched curricula, and they are required more often the more exposed a curriculum is to new technologies.

5 Conclusions

This report studies whether updating formal educational content can reduce the mismatch between the skills supplied by vocational training and the skills demanded in the labor market. Using the near-universe of German dual-system vocational curricula linked to job ads for 2017–2025, we construct a multidimensional measure of curriculum–job ad mismatch based on ESCO skills and exploit the timing of curriculum revisions to estimate the effect of updating training content on mismatch. We find that curriculum updates significantly reduce the share of skills demanded in vacancies but not supplied in the corresponding curriculum: an update lowers mismatch by about 4 percentage points relative to comparable occupations that do not experience an update in the same period.

These findings provide direct evidence that institutionalized revisions of vocational curricula are an effective mechanism for aligning the skills taught in formal training with evolving labor-market requirements. Descriptive evidence further shows that mismatch rises with the time elapsed since the last update, consistent with the view that skill bundles embedded in curricula gradually become less well aligned with employer demand as technologies and work organization change. Curriculum updating therefore appears to be an important channel through which education systems can counteract technologically induced skill obsolescence and limit underqualification at the occupational level.

A central question of the report is whether this adjustment mechanism matters especially in occupations more exposed to technological change. The results show that the decline in mismatch following a curriculum update is slightly larger for curricula which are more exposed to digital breakthrough innovations. However, new technologies cause curricula to be updated more frequently, so that highly exposed curricula are younger, and have a lower level of pre-update mismatch. Once controlling for curriculum age and pre-update mismatch, the effects of updates on mismatch are comparable between occupations that have a high versus low exposure to new technologies. Taken together, this implies that occupations exposed to new technologies require more frequent updates to sustain a comparable quality of skill match as occupations that are less exposed. Rapid technological change requires rapid updating of exposed vocational training curricula.

These findings also carry broader policy implications. In economies characterized by rapid technological change, the quality of skill matching depends not only on how firms adjust labor demand, but also on how quickly and effectively education systems revise the content of training. The German dual system provides an example of an institutional setting in which nationally standardized curricula can be updated in response to changing skill needs. The results suggest that the timing of such updates matters: delaying revisions allows mismatch to accumulate, whereas timely reform can bring training content closer to market demand. This points to the value of systematic monitoring of emerging skill requirements, including through job ad data, as an input into curriculum design and revision.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). Ai and jobs: Evidence from online vacancies. Working Paper 28257, NBER.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. American Economic Review, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Discussion Paper 12293, IZA.

Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in u.s. wage inequality. Econometrica, 90(5):1973–2016.

Agrawal, A. K., Kim, D., and Tambe, P. (2025). Skill mismatch over the technology lifecycle.

Andersson, F., Haltiwanger, J. C., Kutzbach, M. J., Pollakowski, H. O., and Weinberg, D. H. (2018). Job displacement and the duration of joblessness: The role of spatial mismatch. The Review of Economics and Statistics, 100(2):203–218.

Atalay, E., Phongthiengtham, P., Sotelo, S., and Tannenbaum, D. (2020). The evolution of work in the united states. American Economic Journal: Applied Economics, 12(2):1–34.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. Quarterly Journal of Economics.

Autor, D. H., Goldin, C., and Katz, L. F. (2020). Extending the race between education and technology. AEA Papers and Proceedings, 110:347–351.

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

Biasi, B. and Ma, S. (2023). The education-innovation gap. Working Paper 29853, NBER.

Boustan, L. P., Choi, J., and Clingingsmith, D. (2022). Automation after the assembly line: Computerized machine tools, employment and productivity in the united states. Working Paper 30400, NBER.

Braxton, J. C., Herkenhoff, K. F., Rothbaum, J. L., and Schmidt, L. (2021). Changing income risk across the us skill distribution: Evidence from a generalized kalman filter. Working Paper 29567, National Bureau of Economic Research.

Conzelmann, J., Hemelt, S. W., Hershbein, B., Martin, S., Simon, A., and Stange, K. (2023). Skills, majors, and jobs: Does higher education respond? Working Paper 31572, NBER.

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. Quarterly Journal of Economics, 135(4):1965–2005.

Di Giacomo, G. and Lerch, B. (2023). Automation and human capital adjustment: The effect of robots on college enrollment. Journal of Human Resources. Forthcoming.

Dillon, E. W., Kahn, L. B., Venator, J., and Dalton, M. (2025). Do workforce development programs bridge the skills gap? Working Paper 34012, NBER.

Fillmore, I. and Hall, J. D. (2021). Technological change and obsolete skills: Evidence from men's professional tennis. Labour Economics, 73:102051.

Fredriksson, P., Hensvik, L., and Skans, O. N. (2018). Mismatch of talent: Evidence on match quality, entry wages, and job mobility. American Economic Review, 108(11):3303–38.

Gathmann, C., Helm, I., and Schönberg, U. (2018). Spillover effects of mass layoffs. Journal of the European Economic Association, 18(1):427–468.

Gnehm, A.-S., Bühlmann, E., and Clematide, S. (2022). Evaluation of transfer learning and domain adaptation for analyzing german-speaking job advertisements. In Proceedings of the 13th Language Resources and Evaluation Conference, Marseille, France. European Language Resources Association.

Goldin, C. and Katz, L. F. (2008). The Race Between Education and Technology. Harvard University Press.

Guo, Y., Langer, C., Mercorio, F., Mezzanzanica, M., Pallucchini, F., Trentini, F., and Wiederhold, S. (2020). Skill demand, skill supply, and the prevalence of skill mismatch in the european union - pillars report. Technical report.

Guvenen, F., Kuruscu, B., Tanaka, S., and Wiczer, D. (2020). Multidimensional skill mismatch. American Economic Journal: Macroeconomics, 12(1):210–44.

Hampole, M., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Working Paper 33509, NBER.

Hanushek, E. A., Janssen, S., Light, J. D., and Simon, L. K. (2025). Adjusters and casualties: The anatomy of labor market displacement. Working Paper 33667, National Bureau of Economic Research.

Janssen, S. and Mohrenweiser, J. (2018). The shelf life of incumbent workers during accelerating technological change. Discussion Paper 11312, IZA.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. Quarterly Journal of Economics, 107(1):35–78.

Kelly, B., Papanikolaou, D., Seru, A., and Taddy, M. (2021). Measuring technological innovation over the long run. American Economic Review: Insights, 3(3):303–320.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Working Paper 31846, NBER.

Light, J. (2024). Student demand and the supply of college courses. Technical report, Working Paper.

MacDonald, G. and Weisbach, M. S. (2004). The economics of has-beens. Journal of Political Economy, 112(S1):S289–S310.

Neffke, F., Nedelkoska, L., and Wiederhold, S. (2024). Skill mismatch and the costs of job displacement. Research Policy, 53(2):104933.

OECD (2025). Bridging the ai skills gap. Technical report, OECD Publishing.

Restrepo, P. (2024). Automation: Theory, evidence, and outlook. Annual Review of Economics.

Salomons, A., vom Baur, C., and Zierahn-Weilage, U. (2025). Expertise at work: New technologies, new skills, and worker impacts. Technical report, SkiLMeeT Project. Deliverable D2.6, link to latest version.

Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands. Journal of Labor Economics, 24(2):235–270.

Storm, E. (2023). Skill mismatch and learning-by-doing: Theory and evidence from time allocation on tasks. Ruhr Economic Papers 1021, Essen.

Tian, Y. (2025). The shattered career: How job displacement channels high performers away from productivity.

Vandeplas, A. and Thum-Thysen, A. (2019). Skills mismatch and productivity in the eu. European Economy Discussion Paper 100, European Commission.

Ziya Sen, Y. (2025). Skill obsolescence and the consequences of job loss.

SkiLMeeT

Follow us:

linkedin.com/company/SkiLMeeT

bsky.app/profile/skilmeet.bsky.social

twitter.com/SkiLMeeT

www.skilmeet.eu

Funded by
the European Union

This project has received funding from the European Union’s Horizon Europe Research and Innovation programme under grant agreement No 101132581.

SkiLMeeT Partners:

TNO innovation for life

rwi

LISER Luxembourg Institute of Socio-Economic Research

Utrecht University

ibs institute for structural research

Fondazione Eni Enrico Mattei

Institute of Philosophy and Sociology

Bulgarian Academy of Sciences

ALDGATE STRATEGY GROUP