

Discussion Paper Series

IZA DP No. 18975

September 2026

Gender Norms and Occupational Segregation: A Spatial Analysis of the Swiss Apprenticeship Market

Andreas Kuhn

Swiss Federal University for Vocational
Education and Training,
University of Bern and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

Gender Norms and Occupational Segregation: A Spatial Analysis of the Swiss Apprenticeship Market*

Abstract

Women have caught up significantly with men in recent decades, but occupational choices remain stubbornly segregated by gender. One particularly popular hypothesis postulates that societal gender norms are the key factor explaining this phenomenon. I attempt to test the validity of this claim as directly as possible in this paper, using a unique combination of different data sources. Data covering the full stock of apprenticeship contracts active in the school year 2021/22 allows me to construct geographically disaggregated measures of occupational segregation, and I use local results from five national-level votes on gender-related policies to construct a measure of local gender norms. The data reveal that there is a statistically significant and practically relevant association between local gender norms and the regional extent of occupational segregation, with more progressive norms associated with less segregation. However, further analyses, carried out by leveraging two distinct identification strategies, point to a small and ultimately insignificant effect of gender norms on occupational segregation once potential confounders are taken into account.

JEL classification

D91, J16, J24

Keywords

occupational choice, occupational segregation, gender norms, apprenticeship, culture, spatial analysis, language border, geographic quasi-experiment

Corresponding author

Andreas Kuhn

andreas.kuhn@ehb.swiss

* I thank Anne Brenoe, Patricia Palffy, Stefan Wolter, Jürg Schweri as well as seminar participants in Wiesbaden, Zollikofen, and Zurich for helpful comments and suggestions. I also thank Laurent Inversin from the Federal Statistical Office for helpful clarifications concerning the main data source and Sally Gschwend-Fisher for proofreading the manuscript.

1 Introduction

There is no doubt that women have made remarkable progress within various measures of both labor-market inclusion and performance over the past couple of decades (e.g. Blau and Kahn, 2017; Blau and Lynch, 2024; Goldin, 2006; Olivetti and Petrongolo, 2016). At the same time, however, it is also well documented that gender segregation in occupations has changed much more slowly and remains at relatively high levels in many countries (e.g. Anderson, 2023; England *et al.*, 2020; Stoet and Geary, 2022). These remaining differences in occupational choices are of both academic and public interest because they are related to differences in various downstream variables, such as labor earnings (e.g. Dahl *et al.*, 2023; Humphries *et al.*, 2024; Sloane *et al.*, 2021). This same pattern is also observed in Switzerland, which will be the focus of the empirical analysis presented in this paper. On the one hand, women have significantly caught up with men along various measures of both educational attainment and labor-market performance in the past couple of decades.¹ Yet, at the same time, occupational segregation remains pronounced even among the most recent cohorts of labor-market entrants; especially among, but not restricted to, those opting for an apprenticeship training at the post-mandatory level (e.g. Finzi, 2007; Kuhn and Wolter, 2022).²

There is no consensus yet within the literature regarding the underlying reasons for the apparent persistence of occupational gender segregation, but it is well understood that the reasons for this phenomenon must most likely be manifold (e.g. Delaney and Devereux, 2021; Stern and Madison, 2022; Stewart-Williams and Halsey, 2021). Nonetheless, it is also true that one specific hypothesis appears to attract a disproportionate amount of attention, both in the academic and public spheres: the hypothesis that societal gender norms describing the proper roles of women and men are one of the primary explanations for the apparent stillstand regarding the de-segregation in occupational choices. *Prima facie*, the hypothesis certainly has some intuitive appeal – exactly because many gender gaps have closed over time. Presumably for that reason, residual differences in wages, for example, are sometimes indirectly attributed

¹For example, the female labor force participation rate equals about 76% among those aged between 15 and 64 years – one of the highest rates across the OECD (BFS, 2020).

²The Federal Statistical Office (FSO) provides information on gender differences in both occupational choice among apprentices as well as among differences in fields of study by tertiary-level students on the following website: <https://www.bfs.admin.ch/bfs/de/home/statistiken/wirtschaftliche-soziale-situation-bevoelkerung/gleichstellung-frau-mann/bildung/berufs-studienfachwahl.html>

to the influence of these norms. However, there are several counter-arguments to such an approach that is arguably rather error-prone. For example, it is worth pointing out that gender norms have demonstrably and substantively changed in the past several decades in Switzerland (Buchmann and Eisner, 1997; Kuhn, 2024b). Moreover, at least in the Swiss context, many employers and occupational associations actively try to encourage both female and male adolescents to enter their professions, such as *Swissmem*, the association for Switzerland’s mechanical and electrical engineering industries and related technology-oriented sectors, or the information-technology sector.³ Consistent with this more anecdotal evidence, recent empirical studies show that there is no evidence that employers use systematic gender discrimination in filling their open apprenticeship positions (Fernandes *et al.*, 2023; Fossati *et al.*, 2020). Finally, and perhaps most importantly, there are plausible alternative explanations for the persistence in occupational gender segregation, including robust gender differences in vocational interests, psychological traits, and economic preferences (e.g. Buser *et al.*, 2026; Falk *et al.*, 2018; Kuhn and Wolter, 2022; Weisberg *et al.*, 2011).⁴ Taken together, I would therefore argue that it is not at all obvious how relevant these societal norms really are in explaining why men and women still tend to work in different occupations – or whether alternative explanations should be more closely analyzed instead. Consequently, we require direct and plausible empirical evidence on the relevance of this hypothesis.

In this paper, I therefore try to empirically assess the relevance of this hypothesis as directly as possible and within the context of occupational choices among Swiss apprentices. The hypothesis under discussion is difficult to assess empirically, mainly for the following two reasons. First, social gender norms are, by their very nature, an intangible and elusive concept from an empirical point of view. I will address the measurement issue by using results from several national-level votes about policies that touched on the issue of further promoting or enforcing legal and practical equality between men and women. For example, the Swiss electorate was asked in the year 1981 to decide over an amendment to the Swiss constitution which explicitly

³For example, *Swissmem*, the employer association of the Swiss technology industry, runs a network for women in the tech industry (“*SwisswoMEMclub*”), and there are several coding workshops specifically for girls (e.g. “*Girls code too*”).

⁴While single features of gender differences tend to be small, with the notable exception of differences in vocational interests along the things-people dimension (Kuhn, 2022), differences across a multitude of dimensions are nonetheless highly predictive with regard to occupational choices (Buser, 2026; Herlitz *et al.*, 2026).

stated that discrimination based on sex is unconstitutional. Second, and beyond the issue of how to best measure gender norms “in the wild”, it is difficult to isolate the partial effect of gender norms using observational data (e.g. Lonati *et al.*, 2024). In the case at hand, when using voting results to measure gender norms, variation in gender norms comes exclusively from variation across different localities. It is therefore important to disentangle the partial effect from gender norms from the influence of other predictors of occupational choices, such as differences in the regional structure of firms, jobs, and thus also training positions (see Kuhn, 2022, 2024a, for direct evidence on spatial variation in both the number and the structure of training positions across local apprenticeship markets). Arguably, these differences tend to be correlated with local gender norms as well, as more rural areas tend to be characterized by more traditional views concerning the role of men and women (e.g. Kuhn and Wolter, 2023).⁵ Thus, in order to empirically isolate the effect of gender norms on occupational segregation, one must try to net out the structural differences across municipalities that also have a direct effect on occupational segregation.

With these challenges in mind, I will provide two different sets of estimates in an effort to further advance this specific discussion. In the first part of the empirical analysis, I rely on the full available regional variation in gender norms across all Swiss municipalities to estimate the association between local gender norms and local occupational segregation. These estimates provide a useful benchmark, in part because they allow me to compare my results with two recent empirical studies that tried to estimate the same parameter, also using data from Switzerland and a similar design. At the same time, however, these estimates are subject to the objection that relevant unobserved confounders might also exist. In the second part of the empirical analysis, and in an effort to move towards estimates which are more plausibly causal, I therefore focus on a geographic quasi-experimental setting – a term borrowed from Galiani *et al.* (2017) – arising from the fact that there is a relatively large and discontinuous gap in gender norms exactly at the language border which separates the German and the French language areas of the country (see also Lonati *et al.*, 2024; Keele and Titiunik, 2016; Keele *et al.*, 2015). Language is, of course, a strong marker and carrier of culture and identity

⁵Several papers have argued that cultural and economic change are contingent on each other, and thus changes in cultural norms, such as gender roles, are likely to be preceded or accompanied by economic and/or demographic changes (e.g. Aguilar-Gomez and Benshaul-Tolonen, 2023; Chabé-Ferret, 2019; Teso, 2019).

(Angerer *et al.*, 2016; Clots-Figueras and Masella, 2013; Davis and Abdurazokzoda, 2016), and thus comparing different language regions within the same country may be helpful to learn about the role of gender norms. The empirical analysis in this part of the paper is inspired by several empirical papers that focus on behavioral discontinuities at this language border (e.g. Cottier, 2018; Eugster *et al.*, 2011, 2017), and the practical implementation of the methodology mirrors in large parts the approach taken by Aepli *et al.* (2021a), who study how local norms in favor of a stronger role of the state in the realm of apprenticeship training impact private employers’ training decisions. The main idea of this strategy is to focus on municipalities from either side of, yet close to the language border, allowing us to implicitly control for at least some unobservable confounders (see also Lonati *et al.*, 2024).

Previous evidence on the effect of gender norms on occupational choice based on a more or less direct measurement of societal gender norms is, for the above-mentioned reasons, strictly limited.⁶ In a first related study, Osikominu *et al.* (2020) focus on the choice of field of study among Swiss university students, finding that male, but not female, students who grew up in more conservative regions are more likely to choose a STEM-related field of study. In another analysis for Switzerland, Kuhn and Wolter (2023) also use municipality-level voting results from several national-level votes and study whether individual apprentices either aspire for or actually choose a gender-stereotypical occupation. While they do find a negative correlation between gender norms and the stereotypicality of individual occupational aspirations, it also turns out that this result is not robust to controlling for parents’ occupation and/or regional controls – suggesting that the association between occupational aspirations and gender norms may be driven by regional differences in the occupational structure. One downside of their analysis is that their sample only covers relatively few municipalities, and only municipalities from the German language part from the canton of Bern – something on which the present paper will substantively improve upon. Finally, in yet another recent study using Swiss data, closely following the setup of Kuhn and Wolter (2023), Palffy *et al.* (2023b) use data from a popular Swiss website that provides them with information on adolescents’ application behavior on the website (i.e. the adolescents can directly apply for open apprenticeship positions on

⁶Empirical studies thus typically use (a combination of) survey questions as measures of gender norms (e.g. Bornatici *et al.*, 2020; Fortin, 2005; Pessin, 2018).

the website); meaning that they observe which occupations the adolescents apply for. They find that individuals living in more conservative regions are more likely to apply for gender-stereotypical occupations, and also that local gender norms appear to influence only male, but not female, adolescents. The comparison across the two later studies, as well as the finding from both Palffy *et al.* (2023b) and Osikominu *et al.* (2020) that it is only men, not women, who drive the association between norms and segregation, raises further questions regarding the validity of a conventional spatial approach.⁷ At the least, it reveals that the set of control variables used in the analysis is relevant in this context, as it likely explains the different findings across studies. In this paper, I thus also make an effort to improve upon the design of previous studies on the subject, by providing estimates based on two different identification strategies.

To the best of my knowledge, there are only few empirical studies from other countries that try to tackle the question at hand, and only in a less direct way. One such study is Aldén and Neuman (2022), who apply the so-called epidemiological approach (e.g. Lonati *et al.*, 2024).⁸ Using Swedish registry data, and further comparing siblings from the same family, they find that the gender gap in STEM-related college major choice is smaller among individuals originating from countries where relatively more women work in STEM occupations. Another study on the subject is by Morales and Marcén (2023), who apply the epidemiological approach to immigrants to the US, finding that the gender gap in occupational choices tends to be less gendered among immigrants from countries with more gender equality, as measured by the gender gap index (WEF, 2026).⁹

The remainder of this paper is organized as follows. In the following section, I present some background information on the Swiss apprenticeship system and the German-French language border, which will play an important role in the design of the second part of the empirical analysis. Section 3 lays out the different data sources and as well as main variables used in the

⁷Somewhat intriguingly, the same authors find in another paper that interventions targeting adolescents' occupational aspirations only work for women, but not for men (Palffy *et al.*, 2023a). This would suggest that women's aspirations are more, not less, malleable than men's.

⁸Pioneered by Fernández and Fogli (2006), this approach compares second-generation immigrants within the same host country, using information about the culture of their parents' originating country along one or several outcomes. See also Lonati *et al.* (2024) on the pros and cons of this specific empirical approach in estimating the causal effect of culture, along with a discussion of alternative approaches.

⁹Also using the epidemiological approach, other papers have tried to estimate the impact of gender norms on the gender gap in both overall academic achievement and in mathematics (e.g. Holmlund *et al.*, 2023; Nollenberger *et al.*, 2016; Rodríguez-Planas and Nollenberger, 2018).

empirical analysis. The empirical analysis itself and the various estimation results are presented in section 4. Section 5 discusses the implications of the results and concludes.

2 Background

2.1 The Swiss apprenticeship system

The empirical analysis in this paper will focus on the occupational choices among adolescents who decided to enter apprenticeship training at the post-mandatory, upper-secondary level (see appendix figure A.1 for a schematic representation of the Swiss education system). Usually, adolescents enter apprenticeship training after completing two years of kindergarten, six years of primary and three years of secondary-schooling, typically in the year they attain age 16 (see SCCRE, 2026; Wettstein *et al.*, 2017, for additional details concerning the apprenticeship system as well as the education system more generally).

Apprenticeship training remains the most frequent educational choice at the post-mandatory level in Switzerland, although general-education tracks have become substantively more popular in the last couple of decades (e.g. Aepli *et al.*, 2021b). In the year 2017, about 71% of the incoming cohort eventually entered some kind of apprenticeship training, in one of the about 250 distinct learnable occupations (SERI, 2021). Examples of apprenticeship occupations include popular occupations such as commercial employee or healthcare assistant, but also rarely chosen occupations such as glass painter or wood sculptor (see also Kuhn, 2022, 2024a). Across all occupations, the vast majority of these apprenticeships are firm-based apprenticeships, i.e. the apprentice (or, rather, his/her parents, in the case the apprentice is underage) signs a contract with an employer and receives practical training, as well as an apprentice wage for providing productive work for the employer, by his/her employer while also attending vocational school for at least one day per workweek.¹⁰ Most apprenticeships last either three or four years, but there are also apprenticeships tailored for academically weaker adolescents that last only two years. The apprenticeship contract ends once the training is successfully completed, i.e. there is no legal obligation on behalf of the employer to retain apprentices once they have completed

¹⁰Specifically, in the year 2017, about 89.6% of all apprenticeships are firm-based, while the remaining 10.4% represent school-based apprenticeships (SERI, 2021). Also note that school-based apprenticeships are only possible in selected occupations (e.g. information technician or commercial employee).

their training (or, for that matter, for the apprentice to stay longer with the training firm). The costs of the practical training provided by the employer are implicitly shared between employers and apprentices (e.g. Gehret *et al.*, 2025; Muehlemann and Wolter, 2014), while vocational schools are publicly financed.

While the focus on apprentices in this study is a decision driven primarily by data availability, note that one may also make a substantive argument in favor of focusing on this subsample of adolescents. Specifically, it is often pointed out that, because apprentices are relatively young when making their occupational choices, they are more receptive for the influence of gender norms (or, alternatively, perhaps they are less influenced by other motives at this stage of their lives, such as actual or expected wage levels or later career opportunities). Another, more practical advantage from the viewpoint of the empirical design is that apprentices, since most of them still live with their parents during their training period, usually do their apprenticeship with an employer located close to their parents' place of residence (see Kuhn, 2022, for direct evidence). This implies that matching their location of residency with the prevailing gender norms in the same location with high reliability is relatively straightforward. This contrasts with studies that focus on university-level students, such as in Osikominu *et al.* (2020), who often have to move relatively far away from the location where they grew up, simply because there are only few universities in the country. From the viewpoint of the design of the study, this feature makes it much easier, and arguably also less error-prone, to link local gender norms with individual-level educational decisions of apprentices rather than university students.

2.2 Switzerland's different language regions and the French-German language border

In the empirical part of the paper, I will also take advantage of the existence of a French-German language border (the so-called "*Röstigraben*" in German, a reminiscence to a popular dish made out of potatoes) that separates the majority of German-speakers residing in the Eastern and Central parts of the country from the minority of French-speaking residents living in the Western part of the country (see Rash, 2002; van Durme, 2002, on the historical evolution of the

language border from a long-run perspective).¹¹ Overall, the French-German language border stretches over a length of approximately 250 kilometers, from the northern to the southern national border (see appendix figure A.2 as well as Aepli *et al.*, 2021a, for additional details). There are three bilingual cantons in Switzerland (Bern, Fribourg, and Valais), meaning that the language border runs through these three cantons; otherwise the language border coincides with cantonal borders (there are five cantons that directly adjoin the French-German language border: Basel-Landschaft, Jura, Neuchâtel, Solothurn, and Vaud). For example, it runs along the borders between the French-speaking canton of Jura and the German-speaking cantons of Solothurn and Basel-Stadt in the northern part of the country (again, see appendix figure A.2; as evident from that figure, the longer part of the French-German language border runs through the three bilingual cantons).

It is important to point out that the educational context and legal regulation is, at least in principle, the same on both sides of the language border because the legal framework is set at either the federal (in the case of vocational education and training) or primarily at the cantonal level (in case of secondary and upper-secondary schooling).¹² Thus, the fact that the French-German language border runs through some of the cantons allows us to compare adolescents from different cultural contexts, but who are subject to essentially the same institutional and legal framework regulating the educational system (remember that the vocational system is regulated at the national level in any case).

The language border thus separates two distinct cultural regions with broadly identical, or at least very similar institutional and legal contexts. Moreover, the language border remains relatively sharp to this day, i.e. there is a large and discontinuous change in the fraction of German- and French-speaking individuals, respectively, directly at the language border (see Aepli *et al.*, 2021a; Eugster *et al.*, 2011). Several empirical papers have documented that there are different behavioral differences that closely align with the cultural-language border (e.g.

¹¹There are four official languages in Switzerland: German, French, Italian, and Romansh. Both Italian and Romansh are spoken by small minorities of the overall population (in the year 2021, only 8.3% and 0.5%, respectively, stated that this is their main language spoken at home). See the following website for additional details: <https://www.bfs.admin.ch/bfs/de/home/statistiken/bevoelkerung/sprachen-religionen/sprachen.html>

¹²Municipalities are the smallest administrative and political units with some autonomy in both taxation and spending. On average, municipalities are small both in terms of geographic expansion and population size; at the same time, there are relatively large differences in size across municipalities. These differences will also show up in the empirical analysis below.

Aeppli *et al.*, 2021a; Cottier, 2018; Eugster *et al.*, 2011, 2017; Steinhauer, 2018). This literature has documented not only differences in mothers’ labor-supply behavior (Steinhauer, 2018), but also in the likelihood of employers providing apprenticeship positions (Aeppli *et al.*, 2021a). This first result shows that the difference in gender norms has an impact on individual behavior, while the second implies that there are differences in local apprenticeship markets that should also be taken into account in the analysis.

3 Data and key variables

I next describe the main data sources and the key variables, focusing mainly on the measurement of occupational segregation and gender norms at the municipality level.

3.1 Gender-stereotypical occupational choices and occupational segregation

My primary data source is a one-time extract from the statistics of basic vocational education and training (“*Statistik der beruflichen Grundbildung*”, or SBG for short), which covers all individual-level apprenticeship contracts involving apprentices residing within Switzerland as well as employers also located within Switzerland. The data cover the population of all current dual contracts for the 2021/22 apprenticeship year (i.e. the apprenticeship year starting in the late summer of the year 2021). This includes all apprenticeships that newly started in that year, but also apprentices who have advanced in but not yet finished their training. These data are collected mainly for administrative purposes, as the cantonal administrations are responsible, among other things, for the approval of each contract and control of both employers’ training activity and of the individual apprenticeship contracts (for example, the administration is involved in cases where an apprenticeship contract is terminated prematurely, for whatever reason). The cantonal data are then compiled together and prepared for further usage by the Federal Statistical Office (FSO). The FSO then uses these data to document occupational choices among adolescents, and in other areas as well.

The data contain the exact apprenticeship occupation (e.g. retail assistant, healthcare assistant, or commercial employee) as well as the gender and the place of residence of each

individual apprentice, among other things. This information allows me to construct a reliable measure of occupational gender segregation at the municipality level, even for relatively sparsely populated municipalities, as discussed in more detail below.¹³ Moreover, the fact that there is detailed and very accurate information on apprentices' occupation makes it unlikely that there is any noticeable bias in the measurement of segregation caused by the (mis)measurement and/or the aggregation of occupation categories (e.g. Martin-Caughey, 2021; Speer, 2016).

Data cleaning and main estimation sample

In a preliminary step, I consolidate the occupational descriptions in the raw data. Most importantly, I remove the information about different specializations within the same apprenticeship occupation for most of the analyses.¹⁴ This brings the number of unique occupational descriptions down from 460 to 248. Moreover, I eliminate observations with missing information on one or more of the key variables, which yields a maximum set of 185'736 individual-level apprenticeship contracts (in most of these cases, it is the municipality of residence that is missing).

After aggregating the data to the municipality level and after merging all the controls with the key variables (as detailed in appendix table A.1), the main estimation sample consists of 2,107 municipalities, implicitly covering 181'487 individual apprenticeship contracts. In those parts of the empirical analysis where I focus on the whole of Switzerland, this is the underlying sample data.

In contrast, however, in section 4.3, where I focus on municipalities close to the French-German language border, I will only use a relatively small subset of municipalities. Moreover, I do use the full sample of contracts when estimating the proportion of female apprentices within occupations and, in a few of the robustness checks later on, the estimation sample will be somewhat smaller than that because I add several further sets of control variables (these are listed in appendix table A.2).

¹³See <https://www.becc.admin.ch/becc/public/bvz/beruf/grundbildungen> for the current list of learnable occupations in the Swiss apprenticeship system.

¹⁴More popular occupations tend to have more such specializations, and thus aggregating the data by occupation will likely acerbate the skewness in the number of contracts across occupations. Nonetheless, table 2 below will show that this decision has no substantial impact on the estimate of the main parameter of interest.

Measuring occupational segregation at the regional level

A look at the proportion of female and male apprentices within the twenty most popular occupations reveals that a clear majority of either female or male adolescents chooses a given occupation and that, consequently, gender-balanced occupations are the exception rather than the rule. In fact, out of the twenty most popular occupations, only four have about a relatively equal proportion of both female and male apprentices (e.g. commercial employee or retail professional).

Figure 1

To quantify the local extent of occupational gender segregation I first determine, for each apprenticeship occupation o , the fraction of female apprentices choosing that specific occupation in the apprenticeship year 2021/22, denoted by

$$\pi_o^f = (N_o^f / N_o), \quad (1)$$

and with N_o (N_o^f) denoting the total number of contracts (the number of contracts with female apprentices) in occupation o . It is notable that there are large differences in the number of contracts per occupation, with N_o varying between only one contract in several occupations (e.g. in the case of a wood sculptor) up to a maximum of 26,397 contracts for the occupation of a commercial employee (see also Kuhn, 2022, 2024a).¹⁵

For each individual apprentice i , I can then determine the gender-stereotypicity of his or her specific occupational choice as follows:

$$\pi_i = \pi_o^f \cdot \mathbf{1}(i : \text{female}) + (1 - \pi_o^f) \cdot \mathbf{1}(i : \text{male}), \quad (2)$$

with $\mathbf{1}(i : \text{female})$ and $\mathbf{1}(i : \text{male})$ indicating whether individual i is female or male, respectively. Thus π_i simply reflects the proportion of same-sex apprentices within the occupation chosen by individual i . Further note that π_i may vary within the unit interval, excluding the value of zero (i.e. it is not possible for individual i to choose an occupation with no one from the same

¹⁵Table 2 below provides, among other things, a few robustness checks with regard to estimating π_o^f from the available data (such as using a somewhat more extensive data base).

gender as individual i), across individuals and that values closer to 0 (closer to 1) indicate that individual i has chosen a gender-atypical (gender-typical) occupation.¹⁶

In a final step, I aggregate π_i by apprentices' municipality of residence to get a continuous measure of the extent of gender segregation in occupational choices prevailing in a given municipality:

$$\pi_j = \frac{1}{N_j} \sum_{i=1}^N \pi_i \cdot \mathbf{1}(i : \text{resides in municipality } j), \quad (3)$$

with N_j denoting the number of apprenticeship contracts signed with adolescents residing in municipality j . Thus π_j describes whether apprentices in a given municipality, on average, tend to choose rather gender-stereotypical or -atypical occupations, with high (low) values of π_j indicating a high (low) degree of occupational gender segregation in municipality j .¹⁷ This will be the variable of main interest in the empirical analysis below. Note that $\pi_j \in [0, 1)$ and that, as occupational segregation decreases, π_j is expected to converge towards 0.5 (if there is an equal number of female and male apprentices in the population).

Another notable feature is that, there are also large differences in the number of contracts observed across the different municipalities: N_j varies between a low of only one contract in several municipalities to a maximum of 6,265 in the municipality/city of Zurich (Kuhn, 2022, 2024a). There are also several municipalities with zero observed contracts in the year 2021/22. Note that the large differences in the number of contracts by municipality necessitate access to population-level data; otherwise, many municipalities would most likely not be covered by the data at all with only sample data at hand.

One practical advantage of using π_j to measure regional levels of occupational segregation over other segregation measures, such as the popular Duncan index (Duncan and Duncan, 1955), is that this measure also yields meaningful values for municipalities with only few apprenticeship contracts. In contrast, the Duncan index very often takes on its maximum value in municipalities with only few contracts.¹⁸ This has an impact on the distribution of the measure

¹⁶Kuhn and Wolter (2022, 2023) present specific examples of typical and atypical occupational choices among both female and male apprentices.

¹⁷In the available individual-level data, the actual share of female apprentices amounts to about 41%, reflecting the fact that more girls than boys continue with general education at the upper-secondary level, implying that $\pi_j = 0.5162$ would indicate a fully integrated occupational structure in a given municipality.

¹⁸This is due to the fact that there is often no overlap at all in the chosen occupations at all between female

of local occupational segregation, and it may also have influence on the estimated association between gender norms and occupational segregation. Moreover, π_j can also be constructed for female and male apprentices separately, something that is not possible with the Duncan index. This allows me to check whether I can reproduce the finding from both Osikominu *et al.* (2020) and Palffy *et al.* (2023b) that the association between gender-typical occupational choice and gender norms may be driven primarily by male adolescents' choices. Of course, however, I also explicitly show that the choice of π_j over the Duncan index does not drive the results (cf. appendix table A.4).

Figure 2

These measures consistently show that apprentices' occupational choices remain highly segregated by gender, at different aggregation levels of the data. First, panel (a) of figure 2 plots the distribution of π_o^f across the 248 distinct occupations, showing that most apprenticeship occupations still tend to be strongly segregated by gender – consistent with previous evidence (Kuhn and Wolter, 2022, 2023).¹⁹ In addition, panel (b) of figure 2 shows the frequency distribution of π_i for the full population of individual-level apprenticeship contracts. Evidently, the distribution is heavily skewed towards larger values of π_i , with an average value of about 0.73 across all contracts, indicating that a majority of adolescents still prefers to choose a gender-stereotypical occupation.

Figure 3

Figure 3, in turn, shows the distribution of π_j across all municipalities with at least one contract signed with an adolescent living in that particular municipality. Consistent with figure 2 above, this figure shows that most municipalities are characterized by a relatively high degree of gender segregation across all apprentices and occupations, which is evident from the fact that the distribution of π_j is centered at around 0.78. The smallest and largest values are presumably driven by municipalities with only few contracts, and thus panel (b) of figure 3 further shows the

and male apprentices in these regions; in which case the Duncan index takes on its maximum value (cf. appendix figure A.3).

¹⁹As it turns out, some of the occupations with a more equal balance between female and male apprentices are among the most popular occupations overall, such as commercial employee or retail assistant (Kuhn and Wolter, 2022).

distribution of π_j across those municipalities with at least ten active apprenticeship contracts. Among these municipalities, π_j varies between a minimum of 0.47 (in the municipality of *Vufflens-le-Château*, located in the French-speaking canton of Vaud) and a maximum of about 0.91 (in the small municipality of *Epsach*, located in the German-speaking part of the canton of Bern). Thus, while there is notable variation in the regional level of occupational segregation, it is at the same time also true that there is a clear prevalence of gender-typical occupational choices in a vast majority of municipalities.

Occupational aspirations versus occupational choices

At this point, one may object that the apprenticeship contract data is already subject to the potential influence of discriminatory behavior on the part of the training firms, given that the data reflect not only adolescents', but also employers' preferences and, ultimately, hiring decisions. Several different pieces of evidence counter this objection. First, previous studies using a very similar design and the same measure of gender norms (Kuhn and Wolter, 2023; Palffy *et al.*, 2023b) report very similar estimates of the raw association between gender norms on the one hand and occupational aspirations and search queries for open apprenticeship positions, respectively, on the other hand (see also appendix table A.3 below). Second, recent experimental field studies, including a correspondence study, for the Swiss labor market show that there is no evidence of employers systematically discriminating on the basis of gender (Fernandes *et al.*, 2023; Fossati *et al.*, 2020), which also suggests that the differences between individuals are primarily driven by their own aspirations and choices. Third, as mentioned in the introduction, many employers and professional associations actively seek to increase the number of female (or male) apprentices, which also runs counter to the argument that employers somehow enforce segregation in adolescents' occupational choices. In addition, one may add the fact that many occupations are currently struggling to find apprentices to fill their open training positions, which further undermines that argument of gendered gatekeeping by employers.

3.2 Local gender norms

Measuring local gender norms is considerably trickier, but as mentioned in the introduction, we are able to use results from national-level votes to approximate these norms. More specifically, I use municipality-level voting results from five different national-level votes on issues and policies related to the role of women in society and the labor market (see appendix table 1 for some additional information). For example, in the year 1981, the Swiss electorate was asked to vote on an amendment to the constitution explicitly stating that discrimination by sex is unconstitutional, which later led to the Gender Equality Act (“*Gleichstellungsgesetz*”). In general, note that these votes tend to be consequential, which makes them a-priori more plausible than hypothetical survey measures (e.g. Mullainathan *et al.*, 2001).

As the regressor of main interest, I will simply use the mean share of supporting votes across these five votes, with a larger share of supporting votes rather directly reflecting a more progressive norm towards (more) gender equality in each of the five votes:²⁰

$$N_j = \frac{1}{5} \cdot (y_j^{306} + y_j^{336} + y_j^{458} + y_j^{461} + y_j^{513}), \quad (4)$$

where y_j^v denotes the share of supporting votes for vote number v in municipality j (note that this measure aggregates information from various years and that it has a natural scaling). One may object that even the most recent vote is “too old” to be relevant for the present analysis. However, note that the influence of local gender norms is usually understood to flow, to a significant though certainly not to an exclusive extent, from the older to the younger generation, i.e. from parents, teachers and older siblings to children (e.g. Legros and Cislighi, 2020; Young, 2015). From this perspective, one may argue that these votes suitably reflect the prevalent gender norms until about the time the adolescents covered by the data were born; and thus arguably reflect the norms that their parents and teachers grew up with.

Table 1

Using voting results as a measure of regional gender norms has already been done in previous

²⁰In most specifications, I will use the same parameterization of gender norms as Kuhn and Wolter (2023), who provide additional robustness checks related to the use of this measure (they show, among other things, that the results from the different votes are highly correlated with each other and that these norms are associated with various items taken from survey data that explicitly ask individuals about specific topics related to gender).

studies; namely in research on the association between gender norms and subjective well-being (Lalive and Stutzer, 2010), on the potential impact of these norms on the gender pay gap within firms (Janssen *et al.*, 2016), or to explore the relation between regional norms towards a stronger role of the state and employers’ willingness to provide apprenticeship positions (Aepli *et al.*, 2021a; Kuhn *et al.*, 2022). In a closely related paper, Kuhn and Wolter (2023) use the same aggregation of voting results to study the relation between municipality-level gender norms and individual-level occupational aspirations among a sample of adolescents from the German-speaking part of the canton of Bern in 8th grade, as well as occupational choices for the same sample of adolescents later on in 9th grade.²¹

Importantly, note that these voting results do not only quantify the (variation in the) prevailing attitudes on specific gender-related issues, they are also public information – in the sense that the voting results become known immediately after the respective vote has taken place, as the results for national-level votes are discussed and interpreted in both national and local media outlets shortly after the votes have taken place. This makes these voting results almost ideally suited as a measure of societal norms (see also Brennan *et al.*, 2013, for a discussion of this specific aspect of social norms).

Figure 4

As shown in figure 4, there is significant variation in the voting outcomes at the municipality level, and thus arguably also in the underlying latent gender norms (see also appendix figure A.5, which shows the corresponding histogram). Across all Swiss municipalities with at least ten apprenticeship contracts, the mean vote share in support of more equality between men and women varies from a minimum of about 11,6% (in the municipality of “*Alpthal*”, a rural municipality located in the canton of Schwyz) to a maximum of about 72.9% (“*Chêne-Bourg*”, another small municipality, located in the French-speaking canton of Geneva).

Two additional features that are evident from figure 4 are noteworthy. First, there appears

²¹One potentially relevant problem with using these votes as a measure of regional gender norms, especially when focusing on differences across language regions, is that they may also reflect differences in attitudes towards the role of the state (cf. Aepli *et al.*, 2021a; Kuhn *et al.*, 2022). For example, the introduction of paid maternity leave is often understood as a policy to foster equality between men and women in the labor market. At the same time, it is also a state-mandated policy financed by mandatory levies. Thus, variation in the proportion of supporting votes may be driven both by non-financial and financial motives. I will come back to this potentially relevant issue in the empirical part of the paper (cf. section 4.3).

to be a rather sharp discontinuity in attitudes towards more gender equality between the different language-cultural regions; specifically, both the French-German and the German-Italian language border are immediately visible from the figure (this also shows up in the bimodal frequency distribution of appendix figure A.5). Second, there also appears to be a marked difference between urban and rural municipalities, at least within the German-speaking regions of the country.²² These two features also point towards the two different identification strategies that are possible in the present context. A first strategy may thus simply focus on the overall spatial variation in local gender norms, and this will be the first empirical approach that I will take in the analysis below. Another, second strategy is to focus on the apparent discontinuity in gender norms along the French-German language border. In the empirical analysis below, I will implement both of these strategies, with the idea that this increases the plausibility of the estimates – if the two strategies yield consistent results.

3.3 Additional variables

Because the main analysis is done at the municipality level, it is easy to add further control variables, e.g. from published statistics from the FSO. A list of these variables, along with their source and basic descriptives, is given in appendix table A.1.

A first set of controls are variables that describe geographical and topographical features of a municipality, such as the median elevation or its coordinates. A second set of controls refers municipalities' demographics, such as the size of the local resident population. The third set of controls is a full set of cantonal dummies. The fourth set of controls describes the local income distribution within a municipality. The fifth and final set of control variables is a set of dummies representing an official typology of municipalities (e.g. whether a municipality is an agricultural or a tourist region).

²²Based on this map alone, one can easily spot the larger agglomerations of the country. For example, the larger darker-shaded area in the north-eastern part of the country corresponds to the wider agglomeration area around the city of Zurich, stretching from the cities of Baden and Winterthur to Uster, Wetzikon, and Rapperswil-Jona to the city of Zug.

4 Empirical analysis

I now move on to the empirical analysis, which proceeds in three consecutive steps.

4.1 The spatial association between gender norms and occupational segregation

I start with a graphical description of the spatial correlation between gender norms and the extent of occupational segregation at the municipality level. Figure 5 thus plots π_j on the y-axis against N_j on the x-axis, with smaller (larger) values on the x-axis corresponding to more traditional (progressive) gender norms. For better visibility, municipalities with less than ten apprenticeship contracts are excluded from drawing the figure (as shown in table 2 below, this has virtually no impact on the estimated size of the association between the two variables). Moreover, and also for a more informative visualization, I plot observations weighted by the number of apprenticeship contracts – even though most of the estimates are based on unweighted data (on whether to use weights or not in such a setup, see the discussion by Solon *et al.*, 2015). On top, panel (b) of figure 5 further highlights a few specific municipalities within the scatterplot from panel (a), such as the larger cities of Lausanne or Zurich. Not surprisingly, the figure suggests that the larger cities tend to cluster on the right-hand side of the scatterplot.

Figure 5

Evidently, there is a clear and unambiguous negative association between more progressive gender norms and the extent of occupational segregation. More specifically, and as expected, occupational choices tend to be less gender-stereotypical in those municipalities characterized by a more pronounced local norm towards equality between men and women.

Simple regression estimates

Table 2 quantifies the strength of the association evident from figure 2, by running a simple linear regression of local occupational segregation π_j on local gender norms N_j :

$$\pi_j = \alpha + \beta N_j + \epsilon_j, \tag{5}$$

with π_j and N_j as defined in equation (3) and (4), respectively. This yields a statistically significant point estimate of $\hat{\beta} = -0.110$, with a robust standard error of about 0.014, as shown in column (1) of table 2. To put the size of this coefficient into perspective, note that the observed range of N_j is about 0.7 (i.e. N_j varies from a minimum of about 0.1 to a maximum of about 0.8). Therefore, the above point estimate translates into an effective maximum effect of about $-0.077 (= 0.7 \cdot -0.110)$. In relative terms, the point estimate thus implies a hypothetical relative effect of about 10 percent, given that the predicted value among the most conservative municipalities equals about 0.767 (i.e. $-0.077/0.767 \simeq -0.1$).

Table 2

The remaining columns of table 2 show that the size of the correlation is virtually independent of some basic ancillary parameters already chosen at this point. Notably, the point estimate of β is very similar when π_o^f is calculated in a slightly alternative way (columns (2) to (4)): Either using the more detailed level of occupational descriptions, allowing for differences across language regions when computing π_o^f , or using an alternative data base to compute π_o^f all yield very similar estimates ($\hat{\beta} = -0.108$, $\hat{\beta} = -0.108$, and $\hat{\beta} = -0.126$, respectively). The estimate of β is also robust to variations in the measurement of local gender norms (columns (5) to (7)). Using only the earlier two votes from table 1 or, alternatively, only the later three votes, yields estimates that are comparable to the baseline specification ($\hat{\beta} = -0.089$ and $\hat{\beta} = -0.103$, respectively). Moreover, using the vote share in the 2019 election of the national parliament in favor of either the social-democratic party or the green party, as an alternative measure of gender norms, also yields a similar point estimate of parameter β ($\hat{\beta} = -0.101$).²³ Finally, restricting the sample to municipalities with at least ten apprenticeship contracts or weighting the observations by the absolute number of contracts again yield similar point estimates of β ($\hat{\beta} = -0.108$ and $\hat{\beta} = -0.136$, respectively), as shown in columns (8) and (9). Moreover, note that all estimates from table 2 are statistically significant at any conventional level.²⁴

²³Appendix figure A.4 shows that these election results closely align with the measure of gender norms at the municipality level. Local gender norms thus statistically predict left-wing election results later on, underlining the validity of using the voting results as a measure of local gender norms.

²⁴In addition, appendix table A.3 compares the estimate based on actual occupational choices with comparable estimates from Kuhn and Wolter (2023), who use survey-data on occupational aspirations (as well as occupational choices) from a sample of adolescents residing in the German-speaking part of the canton of Bern. This comparison shows that the resulting estimates are of comparable size, in turn suggesting that the focus on occupational choices, rather than aspirations, is most likely not a crucial issue.

At the same time, however, figure 5 also shows that occupational choices remain strongly segregated even in those municipalities characterized by the most progressive societal gender norms. Thus, for example, even in those municipalities characterized by strongly progressive gender norms (such as in the city of Geneva), female and male apprentices still tend to choose very different occupations on average. In fact, among the municipalities with the most progressive views, the predicted value of occupational segregation still amounts to $\hat{\pi}_j \simeq 0.689$ – far away from an integrated situation. Another feature evident from figure 2 is that there is considerable variation in occupational segregation for a given level of local gender norms, suggesting (potentially many) other factors must be relevant in explaining the cross-municipality variation in occupational gender segregation (in this regard, also note the relatively low R-squared in table 2).²⁵

Differences across subsamples and subregions

Before moving on to estimates that control for potential confounders, I show some additional estimates for specific subsamples in table 3 (the first column replicates the baseline estimate from column (1) of table 2).

Table 3

To start with, columns (2) and (3) show, respectively, that the spatial association between occupational segregation and gender norms exists both among the subset of German and French language municipalities ($\hat{\beta} = -0.135$ and $\hat{\beta} = -0.130$). In fact, the two resulting estimates turn out to be very similar, almost identical to, and not statistically different from each other. Moreover, one might argue that, given that the voting results only reflect the norms among Swiss citizens, resident adolescents without Swiss nationality may be more prone to be influenced by norms measured by voting results. Column (4) therefore shows the association between local gender norms and π_j when only focusing on resident apprentices. The resulting point estimate of $\hat{\beta} = -0.121$ is, however, similar to the baseline estimate from the first column. Similarly, one may expect that societal gender norms are more likely to influence younger

²⁵However, because the larger municipalities tend to be located closer to the regression line, a feature that is also evident from figure 5, the R-squared is considerably larger when the observations are weighted by the number of contracts (cf. column (9) of table 2).

apprentices than older ones, and I thus restrict the sample to those apprentices not older than 18 years when starting their apprenticeship in column (5). However, the resulting point estimate equals $\hat{\beta} = -0.073$, and thus turns out to be somewhat smaller, not larger, than the baseline estimate (though not significantly so). Finally, the remaining two columns of table 3 show the association between gender norms and the extent of occupational segregation for female and male apprentices separately. As mentioned in the introduction, both Osikominu *et al.* (2020) and Palffy *et al.* (2023b) have found that it is men, not women, whose decisions appear to be aligned with local gender norms. I can reproduce this finding with the data at hand, in the sense that I find a large and significant correlation between segregation and gender norms among male ($\hat{\beta} = -0.225$, with a robust standard error of about 0.015), but not among female ($\hat{\beta} = 0.025$, with a robust standard error of about 0.018) apprentices. While this could indeed imply that male and female adolescents differ in their reactivity to gender norms, it is certainly not the only explanation for this finding. For example, an alternative explanation could be that male adolescents are more constrained by the local availability of apprenticeship positions (i.e. male apprentices may prefer to work in a rather things-oriented occupation, but there are regional differences in the relative supply of such positions; acting as a constraint on their occupational choices). As the following results will show, the results of my analysis point to the latter explanation, consistent with the findings from Kuhn and Wolter (2023) for the German-speaking part of the canton of Bern.

4.2 A conventional spatial-regression approach

More generally, it remains to be seen if, and/or to what extent, the spatial association between gender norms and occupational segregation can really be understood as a causal relation between the two involved variables. For example, note that figure 5 points to the fact that larger and more urban municipalities tend to have more progressive gender norms than smaller and more rural municipalities. This could reflect differences in the economic and job structure across different regions, which of course influences the supply of apprenticeship positions by employers (Kuhn, 2022).

In a next step, I thus estimate a series of linear regressions that take the following form:

$$\pi_j = \alpha + \beta N_j + \gamma X_j + \epsilon_j, \quad (6)$$

with X_j denoting the inclusion of additional control variables. Parameter β remains of main interest, as it now quantifies the conditional association between gender norms and occupational gender segregation. As above, I show heteroscedasticity-robust standard errors throughout the analysis.

Table 4

The first specification replicates, for the ease of comparison, the estimate from column (1) of table 2 above. The remaining columns of table 4 add additional control variables to the regression, in an effort to take potentially confounding variables into account (see also appendix table A.1). I use several sets of controls, starting with the variables that are, arguably, the most exogenous.

More specifically, the second column adds several variables describing both geographical and topographical features of municipalities, including the median elevation above sea level, among other things. Together, these pick up a substantial part of the raw correlation between gender norms and occupational segregation, yielding a smaller point estimate of $\hat{\beta} = -0.062$ (with a robust standard error of about 0.023, the estimate remains significant). In the third column, I further add a couple of basic demographic controls, such as the log total resident population in a municipality. This leads to a further, yet small decrease in the relevant point estimate ($\hat{\beta} = -0.056$). In contrast, adding a full set of cantonal dummies yields a small increase in the point estimate ($\hat{\beta} = -0.075$), compared to the previous specification. Further taking controls that describe that the local income distribution into account, as in column (5), again leads to a somewhat smaller point estimate ($\hat{\beta} = -0.056$). Finally, adding dummy variables that represent a qualitative typology of municipalities, e.g. whether a municipality is a tourist municipality, yields a point estimate of $\hat{\beta} = -0.034$, which is only about 30 percent the size of the raw coefficient from column (1). Moreover, this final estimate is also statistically insignificant (while, at the same time, the different confidence bands largely overlap across the different specifications).

There is a more or less subtle caveat at this stage of the analysis. Note that, as the point estimate of β becomes smaller, the standard error associated with $\hat{\beta}$ increases quite substantively as additional control variables are included; suggesting that collinearity among the variables might be an issue. In other words, it turns out to be difficult to precisely disentangle the potential effect from gender norms from the effect stemming from the economic or demographic factors. The overall pattern appears to be quite clear nonetheless. Controlling for potential confounders picks up a substantial amount of the raw correlation between gender norms and the extent of occupational segregation, in turn suggesting that both the extent of occupational segregation and the prevailing gender norms are driven to a significant extent by structural differences across the various regions within Switzerland. This is in line with previous evidence on the subject (Kuhn and Wolter, 2023).

4.3 Focusing on municipalities close to the German-French language border

Moreover, one may still argue that other, unobservable confounders may mask an effect of gender norms on occupational segregation. For example, it is difficult to net out all of the relevant factors that influence employers' location decisions, which in turn affect the occupational structure within a region, in a plausible way. In the next and final part of the empirical analysis, I will therefore use a different design approach that will focus on a comparison of municipalities located close to the French-German language border. Arguably, these estimate are closer to a causal estimate of gender norms on occupational gender segregation because at least some otherwise unobservable factors, such as their location relative to specific employers or, more generally, the availability of training positions in different occupations within a given distance, are implicitly taken into account in this kind of setup (Lonati *et al.*, 2024; Keele and Titiunik, 2016; Keele *et al.*, 2015). In the practical implementation of the design, I closely follow the approach taken by Aepli *et al.* (2021a).²⁶

²⁶See also appendix B for additional information concerning the setup in this part of the analysis.

Estimating the local language border contrast in gender norms

In an ancillary step, I first document that a large, unambiguous and robust local language-border contrast (LLBC) in gender norms indeed exists at the French-German language border.

Figure 7

The top panel of figure 7 thus plots local gender norms against the minimum travelling distance to the German-French language border (see also appendix figure 6 for additional plots). More specifically, the figure plots N_j on the y-axis against the minimum travelling time by public transportation, in minutes, to the language border from the corresponding municipality on the x-axis (note that, in the figure, negative values are associated with municipalities on the French-speaking part of the border, and positive values with municipalities with a German-speaking majority). The figure suggests that there is a sizeable difference in gender norms directly at the language border and with gender norms being less progressive on the German-speaking side on average, as expected.

Table 5

The regression estimates presented in table 5 confirm that the LLBC in gender norms is significant in both substantive and statistical terms, as well as robust to the inclusion of additional control variables. The estimates come from regressions that take on the following basic form:

$$N_j = \alpha + \beta G_j + \gamma X_j + \varepsilon_j \quad \forall j \in \{j : d_j \leq bw\}, \quad (7)$$

with N_j again denoting municipality-level gender norms, but now serving as the dependent variable. G_j is a dummy variable indicating whether municipality j is in a German-speaking region, in which case $G_j = 1$, and d_j corresponds to the minimum travelling distance from municipality j to the nearest municipality on the other side of the language border, as discussed in section 6 above. Note that equation (7) is estimated using only observations that are relatively close to the language border, which necessitates the choice of a specific bandwidth bw ; the choice of this ancillary parameter is important because it defines the specific set of municipalities that actually enter the analysis sample (see appendix B for details).

Additional control variables are again denoted by X_j , and they include the same full set of variables as in the preceding step of the analysis. Parameter β remains of primary interest here as it quantifies the LLBC in gender norms. The left-hand panel of table 5 shows the resulting estimates.

In the specification shown in column (1) of table 5, the estimated LLBC in gender norms amounts to about -0.165 percentage points (this estimate corresponds to the difference in gender norms between the two regions highlighted in the bottom panel of figure 7). This difference is highly significant, with a robust standard error of 0.010. This corresponds to a large relative difference of about 29% ($= 100\% \cdot (-0.161/0.553)$) in gender norms between the two language regions. Column (2) of table 5 shows the point estimate of β when also including the full set of control variables. Evidently, the controls do pick up some of the difference in gender norms between the two language regions, but the LLBC in gender norms remains large and statistically significant ($\hat{\beta} = -0.132$, with a robust standard error of about 0.016). Thus even after taking structural differences along various dimensions into account, a substantively and statistically large difference in gender norms remains between the two language regions. It now remains to be seen if there is a corresponding difference in the extent of occupational segregation.

Estimating the local language border contrast in occupational segregation

In a next and analogous step, figure 8 thus illustrates the LLBC in occupational segregation, by plotting municipality-level occupational segregation on the y-axis against its distance from the language border on the x-axis. In contrast to the local contrast in gender norms, figure 8 does not suggest that there is local contrast in occupational segregation, at least not a large one.

Figure 8

However, the impression from the graphical illustration may be misleading, and I thus next turn to estimating the reduced-form contrast in occupational segregation between French- and German-speaking municipalities in close vicinity to the language border:

$$\pi_j = \alpha + \beta G_j + \gamma X_j + \epsilon_j \quad \forall \quad j \in \{j : d_j \leq bw\}, \quad (8)$$

which equals the specification from equation (7), but where the dependent variable is now the measure of local occupational segregation, π_j . As before, β is the parameter of main interest, which will now quantify the estimated difference in the mean level of occupational segregation between the two cultural regions, in close vicinity to the language border.

The first estimate with regard to π_j (column (3) of table 5) simply quantifies the unconditional local language border contrast in occupational segregation. The resulting point estimate equals $\hat{\beta} = 0.013$, which is not statistically significant (with a robust standard error of about 0.009).²⁷ Moreover, and as shown in the final column of table 5, adding the full set of controls also yields a small and statistically insignificant point estimate of $\hat{\beta} = 0.014$, with a robust standard error of about 0.016.

Table 6

As figure 8 suggests a couple of potential outliers on both sides of the language border, I re-estimate the four specifications from table 5 using quantile regression (e.g. Wooldridge, 2002). Specifically, the estimates from table 6 show the LLBC in the (un)conditional median, rather than the (un)conditional mean, of the dependent variables. This apparently has no effect on the qualitative pattern of the results, however, as a large and significant difference in gender norms and a small and statistically insignificant difference in occupational segregation between the two language regions remains.

Robustness

In an additional step, table 7 presents a series of additional estimates that probe the robustness of the finding above that there is a null LLBC in segregation occupational segregation (again, for the ease of comparison, the first column replicates the estimate from column 4 of table 5 above, using the full set of controls).²⁸

Table 7

²⁷To compare this estimate to the one from column (1) of table 2, we may simply multiply it with about 4.24 (corresponding to the ratio of 0.165 to 0.7, i.e. the ratio of the LLBC in gender norms to the total range in local gender norms observed in the full data). This yields $4.24 \cdot 0.013 \simeq 0.055$, which is only about half the size of the corresponding estimate from table 2.

²⁸For the sake of completeness, appendix table A.5 provides the complementary estimates for N_j , showing that the difference in gender norms is persistent and robust across all specifications.

A first check is to see whether the estimate of the LLBC is robust to the choice of the bandwidth. As detailed in appendix B, the baseline bandwidth of a travelling time by public transport (i.e. train and/or bus) to the nearest municipality on the other side of the language border of 40 minutes or less is chosen on a set of ancillary results. Based on these results, other bandwidths also appear plausible a-priori, and thus columns (2) to (6) show estimates using a different bandwidth, ranging from a low of 30 to a high of 60 minutes. In each of these specifications, the resulting point estimate remains small and statistically insignificant ($\hat{\beta} = -0.004$, $\hat{\beta} = -0.002$, and $\hat{\beta} = -0.005$). In a related step, columns (5) and (6) show estimates based on a specification where both a maximum and a minimum distance to the language border is applied to select the estimation sample (i.e. a so-called donut specification). These two specifications also yield a small and statistically insignificant estimate of $\hat{\beta} = 0.000$ and $\hat{\beta} = -0.029$. A third check consists of choosing an alternative distance metric or an alternative travelling mode, respectively; notably in using travelling distances (instead of travelling times) by public transport and travelling times by private rather than by public transportation. Because of differences in scaling, I adjust the bandwidth in these cases such that the number of observations remains about the same. In both cases, the resulting estimates of β remains small and statistically insignificant, as evident in columns (7) and (8), with $\hat{\beta} = 0.013$ and $\hat{\beta} = 0.003$, respectively. Another potentially interesting check consists of focusing on the other notable language border within Switzerland: that separating the German-speaking majority from the Italian-speaking minority (again, see figure 4 as well as appendix figure A.2). However, and in contrast to the German-French border, this language border runs along the Alps, rendering this contrast much less convincing than the German-French language border.²⁹ In practical terms, this implies that the municipalities from the two language regions do in fact not as directly border each other as in the case of the French-German language border; I therefore have to increase the bandwidth significantly to get a sufficient number of municipalities. With that caveat in mind, the resulting point estimate turns out to have the “wrong” sign, but also to be statistically insignificant once again. In column (10), I further restrict the estimation sample to municipalities located in one of the three bilingual cantons (Bern, Fribourg, and Valais). Again, this yields

²⁹This becomes, for example, apparent in appendix figure A.6, which shows the frequency distribution of municipalities along their distance to the corresponding language border for both language borders.

a small and statistically insignificant estimate of parameter β ($\hat{\beta} = 0.008$). As a final check, I again weight the observations by the absolute number of apprenticeship contracts signed with adolescents residing within a given municipality, while at the same time using standard errors that are clustered at the municipality level. As in the preceding checks, the point estimate remains small and statistically insignificant ($\hat{\beta} = -0.006$ with a cluster-robust standard error of 0.016).

Across the full set of robustness checks, the point estimate of β thus remains consistently small and statistically insignificant, corroborating the main finding from table 5 that there is no substantive LLBC with regard to occupational segregation.

Additional controls

In the final step of my analysis, I show a couple of estimates that include some additional, but at least in part potentially endogenous control variables (such as direct controls for the local occupational structure among apprenticeships; appendix table A.2 lists the different sets of additional controls).³⁰

Table 8

The first three columns of table 8 again focus first on the local contrast in gender norms. The first column replicates the estimate from column (2) of table 5, while the specification in the second column adds the mean voter turnout across the five votes used to construct the measure of local gender norms, leaving the point estimate of β virtually unchanged. In the third column I further add an empirical measure of the prevailing local norm towards the role of the state (in the realm of vocational education and training). This variable mirrors the measure used by Kuhn *et al.* (2022) in their study on the determinants of employer participation in apprenticeship training, and it is also based on local voting results. As mentioned earlier, it appears likely that considerations about the role of the state should also play a role in some of the votes from table 1. And indeed, this specification yields a substantially smaller point estimate of the language-border contrast in gender norms, $\hat{\beta} = -0.102$, implying that gender

³⁰For the sake of completeness, appendix table A.6 shows the corresponding estimates using the conventional spatial-regression approach, without further discussion. This table documents that the LLBC in gender norms remains robust across all additional specifications.

norms and attitudes towards the state are indeed correlated with each other.³¹ At the same time, however, a large and statistically significant difference in gender norms of close to ten percentage points remains.

The remaining five columns of table 8 present additional estimates with π_j as the dependent variable. First, column (4) replicates the baseline estimate of column (4) from table 5 above. In the next column, column (5), I add a few additional controls, also constructed from the main data source, that describe the apprentices' demographics, such as the proportion of female apprentices. The resulting point estimate remains small and statistically insignificant ($\hat{\beta} = 0.003$). Column (6) further adds three variables that are meant to capture differences in the selectivity of the two main educational pathways at the upper-secondary level (e.g. the distance to the nearest baccalaureate school, which is meant to influence the decision between choosing an apprenticeship or choosing further general training). Adding these three variables has, again, only a marginal impact on the point estimate ($\hat{\beta} = 0.007$). On top of that, I add a few additional variables in column (7) that reflect important features of the prevailing local apprenticeship market; specifically, the total number of apprenticeship contracts as well as the total number of training firms. Again, this specification yields a small and insignificant point estimate of about $\hat{\beta} = -0.010$, with a robust standard error of about 0.015. In the final column, I also include a restricted set of occupational dummies that directly control for differences in the occupational structure across municipalities.³² This, again, yields a statistically insignificant point estimate of $\hat{\beta} = -0.010$, with a robust standard error of 0.014.

If anything, then, adding these additional control variables pushes the point estimate of β somewhat into the “wrong” direction. Obviously, though, given the relatively large standard errors, a significant overlap in the confidence bands associated with the various point estimates remains, and also with the baseline estimate from the first column.

³¹Consistent with this, there is a local-language border contrast in this measure as well, as shown in appendix figure A.7.

³²As mentioned in section 3 above, there are 248 distinct apprenticeship occupations observed in the year 2021. Thus, to fully describe the occupational distribution in principle, the same number of dummy variables would be necessary; at the same time, note that there are only 232 municipalities in the sample when using the default bandwidth. I therefore decided to only use a restricted set of occupational dummies; specifically, I constructed dummies for only the 30 largest apprenticeship occupations. Because contracts are highly concentrated among the most popular occupations (cf. Kuhn, 2022, 2024a), however, these 30 occupations account for almost 75% of all apprenticeship contracts.

5 Conclusions

Using a unique combination of different data sources, I set out to test the popular hypothesis postulating that the evident persistence in occupational segregation by sex is driven in large parts, or even exclusively, by the persistence of traditional gender norms. Specifically, I use local voting results from five national-level votes about legislation or policies that focus on gender equality at the municipality level to measure local gender norms. Moreover, having access to population-level data on apprenticeship contracts for the apprenticeship year 2021/22 allows me to precisely estimate the extent of occupational segregation across all Swiss municipalities, even for the less densely populated municipalities.

In the empirical analysis itself, I leverage two different identification strategies; the first one relying on the full spatial variation in the key variables, the second one focusing exclusively on municipalities located on either side, but close to the French-German language border. Notably, the second strategy relies on a small subset of all available municipalities, covering only about 11% of the full sample, and thus the two strategies could well lead to different results. Moreover, the second strategy yields estimates that more plausibly reflect a causal estimate because it also allows for a partial control for unobservable confounders.

I find that there is an unambiguous spatial association between local gender norms and the regional level of occupational segregation in the expected direction. That is, municipalities characterized by more progressive gender norms tend to exhibit substantively and significantly lower levels of occupational segregation. The raw correlation alone shows that gender norms cannot be the only remaining predictor of occupational segregation. At the same time, the same data also clearly show that occupational choices remain strongly segregated even in the most progressive areas. Moreover, the estimates that take potential confounders into account show that the raw correlation between local gender norms and occupational segregation is, to a large extent, driven by structural differences between the municipalities. More specifically, it turns out that most of the estimates from either identification strategy are in their majority statistically insignificant, but I would not stress this feature too much, given that these estimates are somewhat imprecise. Nonetheless, and taken together, the key take-away from the analysis is that the data provide only very limited evidence for gender norms having a substantive, if

any, impact on the regional extent of occupational segregation.

The finding that there is no, or at most only a small, effect of gender norms on occupational gender segregation also squares well with the findings from other recent studies on the Swiss apprenticeship market. Specifically, recent experimental evidence on the relevance of gender discrimination in the apprenticeship market by employers also appears to be consistent with this finding (Fernandes *et al.*, 2023; Fossati *et al.*, 2020). In addition to that, field experiments that tried to directly influence Swiss adolescents' occupational choices also tend to yield small or null effects (e.g. Brenoe *et al.*, 2026; Palffy *et al.*, 2023a). In other words, it turns out to be very difficult to directly influence individuals' occupational aspirations and choices in practice – a finding that also appears to be consistent with only a small or even null effect of gender norms on adolescents' occupational choices.

References

- Aepli, M., Kuhn, A., and Schweri, J. (2021a). Culture, norms, and the provision of training by employers: Evidence from the Swiss language border. *Labour Economics*, **73**, 102057.
- Aepli, M., Kuhn, A., and Schweri, J. (2021b). Der Wert von Ausbildungen auf dem Schweizer Arbeitsmark. Bern: Staatssekretariat für Wirtschaft SECO.
- Aguilar-Gomez, S. and Benshaul-Tolonen, A. (2023). The evolution and persistence of women’s roles: Evidence from the Gold Rush. *Journal of Economic Behavior & Organization*, **215**, 364–381.
- Aldén, L. and Neuman, E. (2022). Culture and the gender gap in choice of major: An analysis using sibling comparisons. *Journal of Economic Behavior & Organization*, **201**, 346–373.
- Anderson, G. (2023). Has Canada’s 21st-Century Grand Gender Convergence Stalled? Male and Female Income and Human Resource Stock Distributions Viewed Through an Equal Opportunity Lens. *Review of Income and Wealth*, **69**(4), 907–930.
- Angerer, S., Glätzle-Rützler, D., Lergetporer, P., and Sutter, M. (2016). Cooperation and discrimination within and across language borders: Evidence from children in a bilingual city. *European Economic Review*, **90**, 254–264.
- BFS (2020). Schweizerische Arbeitskräfteerhebung (SAKE). Erwerbsbeteiligung der Frauen 2010–2019. Neuchâtel: Bundesamt für Statistik (BFS).
- Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, **55**(3), 789–865.
- Blau, F. D. and Lynch, L. M. (2024). Fifty years of breakthroughs and barriers: women in economics, policy, and leadership. *The ANNALS of the American Academy of Political and Social Science*, **711**(1), 225–244.
- Bornatici, C., Gauthier, J.-A., and Le Goff, J.-M. (2020). Changing attitudes towards gender equality in switzerland (2000–2017): period, cohort and life-course effects. *Swiss Journal of Sociology*, **46**(3), 559–585.
- Brennan, G., Eriksson, L., Goodin, R. E., and Southwood, N. (2013). *Explaining norms*. Oxford University Press.
- Brenoe, A. A., Carlana, M., Zöllner, T., and Wolter, S. C. (2026). Backlash from gender equality? Working Paper, University of Bern.
- Buchmann, M. and Eisner, M. (1997). The transition from the utilitarian to the expressive self: 1900–1992. *Poetics*, **25**(2-3), 157–175.
- Buser, T. (2026). Too competitive to care? the overall explanatory power of personality for occupational gender segregation. IZA@LISER Working Paper No. 18721.
- Buser, T., Niederle, M., and Oosterbeek, H. (2026). Can competitiveness predict education and labor market outcomes? evidence from incentivized choice and survey measures. *Review of Economics and Statistics*, **108**(3), 755–773.
- Chabé-Ferret, B. (2019). Adherence to cultural norms and economic incentives: evidence from fertility timing decisions. *Journal of Economic Behavior & Organization*, **162**, 24–48.


- Clots-Figueras, I. and Masella, P. (2013). Education, language and identity. *Economic Journal*, **123**(570), F332–F357.
- Cottier, L. (2018). Culture, financial constraints, and retirement decision. *Labour Economics*, **53**, 128–145.
- Dahl, G. B., Rooth, D.-O., and Stenberg, A. (2023). High school majors and future earnings. *American Economic Journal: Applied Economics*, **15**(1), 351–382.
- Davis, L. S. and Abdurazokzoda, F. (2016). Language, culture and institutions: Evidence from a new linguistic dataset. *Journal of Comparative Economics*, **44**(3), 541–561.
- Delaney, J. M. and Devereux, P. J. (2021). The Economics of Gender and Educational Achievement: Stylized Facts and Causal Evidence. In A. Banerjee, editor, *Oxford Research Encyclopedia of Economics and Finance*. Oxford Academic.
- Duncan, O. D. and Duncan, B. (1955). A methodological analysis of segregation indexes. *American Sociological Review*, **20**(2), 210–217.
- England, P., Levine, A., and Mishel, E. (2020). Progress toward gender equality in the United States has slowed or stalled. *Proceedings of the National Academy of Sciences*, **117**(13), 6990–6997.
- Eugster, B., Lalive, R., Steinhauer, A., and Zweimüller, J. (2011). The demand for social insurance: does culture matter? *Economic Journal*, **121**(556), F413–F448.
- Eugster, B., Lalive, R., Steinhauer, A., and Zweimüller, J. (2017). Culture, work attitudes, and job search: Evidence from the Swiss language border. *Journal of the European Economic Association*, **15**(5), 1056–1100.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2018). Global evidence on economic preferences. *Quarterly Journal of Economics*, **133**(4), 1645–1692.
- Fernandes, A., Huber, M., and Plaza, C. (2023). When does gender occupational segregation start? An experimental evaluation of the effects of gender and parental occupation in the apprenticeship labor market. *Economics of Education Review*, **95**, 102399.
- Fernández, R. and Fogli, A. (2006). Fertility: The role of culture and family experience. *Journal of the European Economic Association*, **4**(2-3), 552–561.
- Finzi, I. (2007). *Occupational gender segregation and gender wage gap in Switzerland*. Ph.D. thesis, Università della Svizzera italiana.
- Fortin, N. M. (2005). Gender role attitudes and the labour-market outcomes of women across OECD countries. *Oxford Review of Economic Policy*, **21**(3), 416–438.
- Fossati, F., Wilson, A., and Bonoli, G. (2020). What signals do employers use when hiring? evidence from a survey experiment in the apprenticeship market. *European Sociological Review*, **36**(5), 760–779.
- Galiani, S., McEwan, P. J., and Quistorff, B. (2017). External and internal validity of a geographic quasi-experiment embedded in a cluster-randomized experiment. In M. D. Cattaneo and J. C. Escanciano, editors, *Regression Discontinuity Designs. Theory and Applications*, volume 38 of *Advances in Econometrics*. Emerald Publishing Limited.

- Gehret, A., Kuhn, A., and Schweri, J. (2025). Is training apprentices worthwhile? company perspectives on costs, benefits and training quality 2025. Zollikofen: Swiss Federal University for Vocational Education and Training.
- Goldin, C. (2006). The quiet revolution that transformed women’s employment, education, and family. *American Economic Review*, **96**(2), 1–21.
- Herlitz, A., Frostegård, J. K. E., Asperholm, M., Bränström, R., Guest, E., Martinsen, J., Sonnegård, H., Sorjonen, K., Thorell, L. B., and Persson, B. N. (2026). The power of many small sex differences in cognition, personality, and interests. *Scientific Reports*, **16**(1), 16572.
- Holmlund, H., Rainer, H., and Reich, P. (2023). All geared towards success? Cultural origins of gender gaps in student achievement. *Journal of Economic Behavior & Organization*, **209**, 222–242.
- Humphries, J. E., Joensen, J. S., and Veramendi, G. F. (2024). The gender wage gap: Skills, sorting, and returns. In *AEA Papers and Proceedings*, volume 114, pages 259–264. American Economic Association.
- Janssen, S., Tuor Sartore, S., and Backes-Gellner, U. (2016). Discriminatory social attitudes and varying gender pay gaps within firms. *Industrial & Labor Relations Review*, **69**(1), 253–279.
- Keele, L. and Titiunik, R. (2016). Natural experiments based on geography. *Political Science Research and Methods*, **4**(1), 65.
- Keele, L. J., Titiunik, R., *et al.* (2015). Geographic boundaries as regression discontinuities. *Political Analysis*, **23**(1), 127–155.
- Kuhn, A. (2022). The Geography of Occupational Choice. Empirical Evidence from the Swiss Apprenticeship Market. IZA Discussion Paper No. 15679.
- Kuhn, A. (2024a). Eine Flotte aus Jollen und Barken, Kuttern und Tankern. Transfer. Berufsbildung in Forschung und Praxis 9(1).
- Kuhn, A. (2024b). The Times Have Changed: Tracking the Evolution of Gender Norms over Time. In B. Elsner and S. Polachek, editors, *Big Data Applications in Labor Economics, Part B*, volume 52 of *Research in Labor Economics*, pages 131–170. Emerald Publishing Limited.
- Kuhn, A. and Wolter, S. C. (2022). Things versus people: Gender differences in vocational interests and in occupational preferences. *Journal of Economic Behavior & Organization*, **203**, 210–234.
- Kuhn, A. and Wolter, S. C. (2023). The strength of gender norms and gender-stereotypical occupational aspirations among adolescents. *Kyklos*, **76**(1), 101–124.
- Kuhn, A., Schweri, J., and Wolter, S. C. (2022). Local Norms Describing the Role of the State and the Private Provision of Training. *European Journal of Political Economy*, **75**, 102226.
- Lalive, R. and Stutzer, A. (2010). Approval of equal rights and gender differences in well-being. *Journal of Population Economics*, **23**(3), 933–962.
- Legros, S. and Cislighi, B. (2020). Mapping the social-norms literature: An overview of reviews. *Perspectives on Psychological Science*, **15**(1), 62–80.

- Lonati, S., Lalive, R., and Efferson, C. (2024). Identifying culture as cause: Challenges and opportunities. *Evolutionary Human Sciences*, pages 1–66.
- Martin-Caughey, A. (2021). What’s in an occupation? Investigating within-occupation variation and gender segregation using job titles and task descriptions. *American Sociological Review*, **86**(5), 960–999.
- Morales, M. and Marcén, M. (2023). The effect of gender norms on gender-based sorting across occupations. *Review of Economics of the Household*, pages 1–34.
- Muehlemann, S. and Wolter, S. C. (2014). Return on investment of apprenticeship systems for enterprises: Evidence from cost-benefit analyses. *IZA Journal of Labor Policy*, **3**(1), 1–22.
- Mullainathan, S., Bertrand, M., *et al.* (2001). Do people mean what they say? implications for subjective survey data. *American Economic Review*, **91**(2), 67–72.
- Nollenberger, N., Rodríguez-Planas, N., and Sevilla, A. (2016). The math gender gap: The role of culture. *American Economic Review*, **106**(5), 257–261.
- Olivetti, C. and Petrongolo, B. (2016). The evolution of gender gaps in industrialized countries. *Annual Review of Economics*, **8**, 405–434.
- Osikominu, A., Grossmann, V., and Osterfeld, M. (2020). Sociocultural background and choice of STEM majors at university. *Oxford Economic Papers*, **72**(2), 347–369.
- Palfy, P., Lehnert, P., and Backes-Gellner, U. (2023a). Countering gender-typicality in occupational choices: An information intervention targeted at adolescents. Working Paper. University of Zurich, Department of Business Administration (IBW).
- Palfy, P., Lehnert, P., and Backes-Gellner, U. (2023b). Social norms and gendered occupational choices of men and women: Time to turn the tide? *Industrial Relations: A Journal of Economy and Society*, **62**(4), 380–410.
- Pessin, L. (2018). Changing gender norms and marriage dynamics in the United States. *Journal of Marriage and Family*, **80**(1), 25–41.
- Rash, F. (2002). The German-Romance Language Borders in Switzerland. In J. Treffers-Daller and R. Willemyns, editors, *Language Contact at the Romance–Germanic Language Border*, pages 112–136. Multilingual Matters LTD.
- Rodríguez-Planas, N. and Nollenberger, N. (2018). Let the girls learn! it is not only about math... it’s about gender social norms. *Economics of Education Review*, **62**, 230–253.
- SCCRE (2026). *Swiss Education Report 2026*. Aarau: Swiss Coordination Centre for Research in Education (SCCRE).
- SERI (2021). Vocational and professional education and training in Switzerland – Facts and figures 2021. State Secretariat for Education, Research and Innovation SERI.
- Sloane, C. M., Hurst, E. G., and Black, D. A. (2021). College majors, occupations, and the gender wage gap. *Journal of Economic Perspectives*, **35**(4), 223–248.
- Solon, G., Haider, S. J., and Wooldridge, J. M. (2015). What are we weighting for? *Journal of Human Resources*, **50**(2), 301–316.


- Speer, J. D. (2016). How bad is occupational coding error? A task-based approach. *Economics Letters*, **141**, 166–168.
- Steinhauer, A. (2018). Working moms, childlessness, and female identity. CEPR Discussion Paper No. DP12929.
- Stern, C. and Madison, G. (2022). Sex differences and occupational choice: Theorizing for policy informed by behavioral science. *Journal of Economic Behavior & Organization*, **202**, 694–702.
- Stewart-Williams, S. and Halsey, L. G. (2021). Men, women and STEM: Why the differences and what should be done? *European Journal of Personality*, **35**(1), 3–39.
- Stoet, G. and Geary, D. C. (2022). Sex differences in adolescents’ occupational aspirations: Variations across time and place. *PLOS One*, **17**(1), e0261438.
- Teso, E. (2019). The long-term effect of demographic shocks on the evolution of gender roles: Evidence from the transatlantic slave trade. *Journal of the European Economic Association*, **17**(2), 497–534.
- van Durme, L. (2002). Genesis and Evolution of the Romance-Germanic Language Border in Europe. In J. Treffers-Daller and R. Willems, editors, *Language Contact at the Romance–Germanic Language Border*, pages 9–21. Multilingual Matters LTD.
- WEF (2026). Global Gender Gap Report 2023. Geneva: World Economic Forum (WEF).
- Weisberg, Y. J., DeYoung, C. G., and Hirsh, J. B. (2011). Gender differences in personality across the ten aspects of the big five. *Frontiers in psychology*, **2**, 11757.
- Wettstein, E., Schmid, E., and Gonon, P. (2017). *Swiss vocational and professional education and training (VPET): forms, systems, stakeholders*. hep.
- Wooldridge, J. M. (2002). *Econometric analysis of cross section and panel data*. MIT Press.
- Young, H. P. (2015). The evolution of social norms. *Annual Review of Economics*, **7**(1), 359–387.

Figure 1: Gender differences in apprentices' occupational choices




Notes: The figure shows the percentage proportion of female and male apprentices among the twenty most popular occupations in the apprenticeship year 2021/22. Occupations are ranked in descending order according to the total (i.e. female plus male) number of active apprenticeship contracts. These twenty occupations account for almost 64% of all active contracts in that specific year. Where available, I use the official english description for the occupations.

Figure 2: Measuring occupational segregation




(a) Occupational-level share of female apprentices, π_o^f




(b) Gender-stereotypicity of individual-level occupational choices, π_i

Notes: Panel (a) shows the frequency distribution of the occupational-level share of female apprentices across the 248 distinct occupations, π_o^f . Panel (b) shows the frequency distribution of π_i across the 185'736 individual-level apprenticeship contracts.

Figure 3: Occupational segregation at the local level




(a) The frequency distribution of π_j across all municipalities with at least one apprenticeship contract



(b) The frequency distribution of π_j across municipalities with at least ten apprenticeship contracts


Notes: Panel (a) shows the frequency distribution of π_j across all municipalities with a positive number of apprenticeship contracts (2,107 municipalities). Panel (b) shows the frequency distribution of π_j across those municipalities with at least ten active apprenticeship contracts (1,831 municipalities).

Figure 4: Spatial variation in local gender norms




Notes: The figure shows the spatial variation in gender norms across all Swiss municipalities, with the measure of gender norms based on the results from five national-level votes from table 1. Municipalities shaded in darker (lighter) blue are characterised by a higher (lower) mean share of supporting votes across the five different votes from table 1, indicating stronger (weaker) support for additional policies and legislation pushing gender equality; areas shaded in white show lakes or areas with missing information.

Figure 5: Occupational gender segregation and gender norms



(a) All municipalities with at least 10 contracts



(b) All municipalities with at least 10 contracts, several municipalities highlighted

Notes: The figure plots π_j , describing the regional extent of occupational segregation, on the y-axis against N_j , the measure of local gender norms, on the x-axis. Each circle represents a municipality, with the size being proportional to the number of apprenticeship contracts signed with adolescents living in that municipality. For better visualization, municipalities with less than ten contracts are not included in this figure (there are a few heavy outliers in the data; see also panel (a) of figure 3). Panel (b) replicates panel (a), while further highlighting the position of several specific municipalities within the scatterplot.


Figure 6: Estimating local language border contrasts along the French-German language border



(a) The French- and German-language regions and the French- (b) Municipalities reachable by public transportation from the other side of the language region within 40 minutes or less


Notes: The figure illustrates the basic idea of the geographic quasi-experimental design, focusing on the cantons along the French-German language border in the western part of the country. Panel (a) shows the French-German language border, based on whether there is a relative majority of either French- or German-speaking individuals. Panel (b) highlights the municipalities included in the baseline analysis sample (i.e. those municipalities not located farther away than 40 minutes from the nearest municipality on the other side of the language border). The two figures zoom in on the cantons along the French-German language border (i.e. Basel-Land, Basel-Stadt, Bern, Fribourg, Jura, Neuchâtel, Solothurn, Wallis, and Vaud).

Figure 7: The local contrast in gender norms at the French-German language border



Notes: The figure plots the local gender norm N_j in a given municipality on the y-axis against the minimum distance of the municipality to the language border, measured by the travelling time by public transport in minutes, on the x-axis. To draw the figure, travelling times for observations from the French-speaking part are multiplied with -1 (and are thus located on the left-hand side of the scatterplot). The two dashed horizontal lines indicate the corresponding (unweighted) mean of local gender norms on either side of the language border and within the baseline bandwidth of 40 kilometers.

Figure 8: The local contrast in occupational segregation at the French-German language border



Notes: The figure plots π_j , describing the the gender-stereotypicity of occupational choices in municipality j , on the y-axis against the minimum distance of a municipality to the language border, measured by the travelling time by public transport in minutes, on the x-axis. To draw the figure, travelling times for observations from the French-speaking part are multiplied with -1 (and are thus located on the left-hand side of the scatterplot). The two dashed horizontal lines indicate the corresponding (unweighted) mean gender-stereotypicity of occupational choices on either side of the language border and within the bandwidth of 40 kilometers.

Table 1: Selected national-level votes on gender issues in Switzerland

Nr.	Date	Key subject of the vote	Result	Approval
306	14.06.1981	“Equal rights for men and women”	Accepted	60.30%
336	22.09.1985	Revision of the civil code (marital law)	Accepted	54.70%
458	13.06.1999	Introduction of paid maternity leave	Rejected	39.00%
461	12.03.2000	Introduction of gender quota within the Federal Administration	Rejected	18.00%
513	26.09.2004	Introduction of paid maternity leave	Accepted	55.50%

Notes: The table lists the five votes used to measure local gender norms. The vote number refers to the official numbering used by the Federal Statistical Office. Additional information on the different votes can be found on the website of the Swiss Federal Administration (<https://www.admin.ch/gov/en/start/documentation/votes.html>). The approval rate equals the percentage proportion of supporting votes among all valid votes.

Table 2: The spatial association between gender norms and occupational segregation

	π_j									
N_j	-0.110*** (0.014)	-0.108*** (0.013)	-0.108*** (0.014)	-0.126*** (0.014)	-0.089*** (0.013)	-0.103*** (0.014)	-0.101*** (0.017)	-0.108*** (0.010)	-0.136*** (0.008)	
Computation of π_o^f :										
#Occupations	248	459	248	248	248	248	248	248	248	248
By language region	No	No	Yes	No	No	No	No	No	No	No
Data base	SBG	SBG	SBG	FSO	SBG	SBG	SBG	SBG	SBG	SBG
#Votes	5	5	5	5	2	3	-	5	5	5
Left-wing vote share	No	No	No	No	No	No	Yes	No	No	No
#Contracts $\geq$	1	1	1	10	1	1	1	10	1	1
Weights	No	No	No	No	No	No	No	No	No	Yes
N	2107	2107	2107	2107	2061	2107	2044	1831	2107	
R-Squared	0.041	0.039	0.038	0.047	0.029	0.040	0.021	0.074	0.211	

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses (the standard error in the final column is clustered by municipality). #Occupations denotes the number of distinct occupational categories; #Contracts the minimum number of individual contracts per municipality; and #Votes the number of votes used to construct the measure of local gender norms, N_j . SBG refers to the use of the main data source (i.e. the stock of apprenticeships in the year 2021/22) in constructing occupation-level shares of female apprentices, while FSO refers to the use of inflow data over the years 2013–2017, taken from the FSO’s website (to be clear, both data sets are provided by the FSO). Left-wing vote share corresponds the vote share in support of either the social-democratic or the green party in the 2019 elections for the national parliament.

Table 3: Estimates for different subregions and subpopulations

N_j	π_j						
	-0.110*** (0.014)	-0.135*** (0.023)	-0.130*** (0.047)	-0.121*** (0.015)	-0.073*** (0.016)	0.025 (0.018)	-0.225*** (0.015)
Language region	All	German	French	All	All	All	All
Subpopulation	Full	Full	Full	Natives	Adolescents	Women	Men
N	2107	1380	599	2105	2100	2065	2095
R-Squared	0.041	0.033	0.015	0.042	0.013	0.001	0.114

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses. The native subpopulation only considers apprentices with Swiss citizenship; adolescents focuses on apprentices no older than 18 years when starting their apprenticeship.

Table 4: Spatial regression estimates

	π_j					
N_j	-0.110*** (0.014)	-0.062*** (0.023)	-0.056** (0.023)	-0.075** (0.032)	-0.056* (0.034)	-0.034 (0.039)
Controls:						
Geo-/Topography	No	Yes	Yes	Yes	Yes	Yes
Demographics	No	No	Yes	Yes	Yes	Yes
Cantonal dummies	No	No	No	Yes	Yes	Yes
Income distribution	No	No	No	No	Yes	Yes
Municipal typology	No	No	No	No	No	Yes
N	2107	2107	2107	2107	2107	2107
R-Squared	0.041	0.074	0.088	0.105	0.108	0.116

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses. See appendix table A.1 for additional details concerning the different sets of control variables.

Table 5: Estimating the local-language border contrast in gender norms and in occupational segregation

	N_j		π_j	
G_j	−0.165*** (0.010)	−0.132*** (0.016)	0.013 (0.009)	0.014 (0.016)
Full set of controls	No	Yes	No	Yes
N	232	232	232	232
R-Squared	0.546	0.836	0.010	0.315

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses.

Table 6: Re-estimating the local-language border contrast in gender norms and in occupational segregation using quantile (median) regression

	N_j		π_j	
G_j	−0.175*** (0.013)	−0.143*** (0.012)	0.006 (0.009)	−0.002 (0.019)
Full set of controls	No	Yes	No	Yes
N	232	232	232	232
Pseudo R-Squared	0.349	0.663	0.002	0.199

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses. The target parameter in this table is the (un)conditional median of the corresponding dependant variable.

Table 7: Local-language border contrast in occupational segregation, robustness

π_j												
G_j	0.014 (0.016)	-0.004 (0.022)	-0.002 (0.014)	-0.005 (0.012)	0.000 (0.019)	-0.029 (0.023)	0.013 (0.015)	0.003 (0.015)	0.022 (0.025)	0.008 (0.020)	-0.006 (0.016)	
Full controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Bandwidth:												
Maximum	40	30	50	60	40	40	19	20	120	40	40	
Minimum	0	0	0	0	10	20	0	0	0	0	0	
Travelling mode	Public	Public	Public	Public	Public	Public	Public	Private	Public	Public	Public	
Distance metric	Time	Time	Time	Time	Time	Time	Distance	Time	Time	Time	Time	
Contrast	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-IT	DE-FR	DE-FR	
Bilingual cantons only	No	No	No	No	No	No	No	No	No	Yes	No	
Weights	No	No	No	No	No	No	No	No	No	No	Yes	
N	232	144	339	451	187	148	235	255	324	158	232	
R-Squared	0.315	0.441	0.213	0.171	0.430	0.503	0.322	0.309	0.205	0.374	0.345	

Notes: ***, **, * and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses (the standard error in the final column is clustered by municipality). Each specification includes the full set of baseline controls. The distance metric is the minimum travelling time necessary to reach the closest municipality on the other side of the language border, either by public or private transport. The language-border contrast is either between the German- and the French-speaking regions (DE-FR) or between the German- and the Italian-speaking regions (DE-IT). The estimate in the final column is weighted by the absolute number of apprenticeship contracts, and its standard error is clustered by municipality. See also appendix table A.5 for the corresponding estimates with regard to local gender norms, N_j .


Table 8: Local-language border contrast, additional control variables

G_j	N_j			π_j		
	-0.132^{***} (0.016)	-0.131^{***} (0.015)	-0.102^{***} (0.015)	0.014 (0.016)	0.003 (0.013)	0.007 (0.014)
Baseline controls	Yes	Yes	Yes	Yes	Yes	Yes
Voter turnout	No	Yes	Yes	-	-	-
Attitudes towards role of state	No	No	Yes	-	-	-
Apprentices' demographics	-	-	-	No	Yes	Yes
Selection into apprenticeship	-	-	-	No	No	Yes
Apprenticeship market	-	-	-	No	No	Yes
Occupational structure	-	-	-	No	No	Yes
N	232	232	232	232	232	214
R-Squared	0.836	0.842	0.866	0.315	0.522	0.568
						0.684

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses. See also appendix table A.2 for details and basic descriptives regarding the different sets of additional control variables.

A Additional figures and tables


Figure A.1: The Swiss education system



Source: State Secretariat for Education, Research and Innovation (SERI): <https://www.sbf.admin.ch/en>.


Notes: Vocational education (i.e. apprenticeship training) is positioned on the left-hand side of the illustration. Note that there are various pathways into tertiary education (in blue) from basic vocational education at the upper-secondary level ("Federal Diploma of Vocational Education and Training").

Figure A.2: The four language regions of Switzerland




Notes: The figure illustrates the four language regions of Switzerland (German, French, Italian, and Romansh). The black lines delineate cantonal borders. The French-German language border, the so-called "*Röstigraben*", is located in the western part of the country and runs from the northern to the southern country border.

Figure A.3: Comparison between two alternative segregation measures




Notes: Panels (a) and (b) show the frequency distribution of the local Duncan index; panels (c) and (d) show the association between the local Duncan index and π_j ; panels (a) and (c) include all municipalities, while panels (b) and (d) only include those municipalities with 30 or more contracts.

Figure A.4: Spatial association between left-wing vote shares and gender norms




(a) Gender norms and left-wing vote share



(b) Gender norms and left-wing vote share, with some municipalities highlighted


Notes: The figures plot the spatial association between the left-wing vote share (i.e. vote share in support of either the social-democratic or the green party) in the 2019 elections for the national parliament (on the x-axis) and gender norms (on the y-axis). Figure (b) replicates figure (a), but again highlights the position of several specific municipalities (i.e. the same as in figure 5 in the main text). The correlation between the two variables equals $\hat{\rho} = 0.696$ for the unweighted data (and $\hat{\rho} = 0.762$ for the weighted data).

Figure A.5: The distribution of local gender norms




Notes: The figure shows the frequency distribution of municipality-level gender norms. The evident bimodal shape of the distribution is driven by differences across language regions (i.e. German- versus non-German speaking regions; see also figure 4).

Figure A.6: The distribution of municipalities along their distance to the language border




(a) Distance to the French-German language border (*"Röstigraben"*)



(b) Distance to the Italian-German language border (*"Polentagraben"*)

Notes: The figure plots the number of the Swiss municipalities ordered along their distance to either the French-German (panel (a)) or the Italian-German (panel (b)) language border. To draw the figure, travelling times for observations from the French- or the Italian-speaking part are multiplied with -1 (and are thus located on the left-hand side of the scatterplot). Note the different pattern between the two language region (which is broadly due to geography).

Figure A.7: Support for additional intervention by the state in the realm of vocational education and training



Notes: The figure shows the local-language border contrast in local norms towards the role of the state, in the realm of vocational education and training. This variable is identical to the measure used in Kuhn *et al.* (2022), and it is also based on local results from national-level votes. To draw the figure, travelling times for observations from the French-speaking part are multiplied with -1 (and are thus located on the left-hand side of the scatterplot).

Table A.1: Descriptive statistics for the full sample, key variables and main control variables

Variable	Mean	SD	Source
<i>(a) Key variables</i>			
Occupational segregation, π_j	0.734	0.069	SBG
Gender norms, N_j	0.420	0.125	FSO
<i>(b) Control variables: geo-/topography</i>			
Minimum elevation (in meters)	510.731	197.049	FOT
Maximum elevation (in meters)	1137.073	857.787	FOT
Median elevation (in meters)	755.516	478.168	FOT
X-coordinate of the municipality centroid	631726.009	70133.091	FOT
Y-coordinate of the municipality centroid	203409.160	50391.020	FOT
Perimeter of the municipality (in meters)	18650.740	13354.886	FOT
Total area of the municipality (in square kilometers)	17.819	31.982	FOT
Urban area (proportion)	15.121	14.894	FOT
Agricultural land (proportion)	45.963	19.131	FOT
Forest area (proportion)	32.449	16.024	FOT
Minimum distance to the country border (in meters)	23319.328	17882.398	ARE/FOT
Mean accessibility by public transportation (in minutes)	191.114	42.519	ARE/FOT
Distance to the nearest city (in meters)	21701.683	14622.120	ARE/FOT
<i>(c) Control variables: demographics</i>			
# residents	3995.638	13042.647	FSO
Population density	445.903	802.188	FSO
Percentage proportion non-native residents	17.030	9.740	FSO
Marriage rate (percent of the resident population)	4.168	2.477	FSO
Divorce rate (percent of the resident population)	1.905	1.563	FSO
Birth rate (percent of the resident population)	9.243	3.813	FSO
Death rate (percent of the resident population)	7.619	3.727	FSO
Vacancy rate (rental homes, percent)	1.972	1.698	FSO
<i>(d) Control variables: cantonal dummies^(*)</i>			
Canton of Bern (yes = 1)	0.154	0.361	FSO
Canton of Fribourg (yes = 1)	0.057	0.232	FSO
<i>(e) Control variables: taxable income</i>			
# taxpayers	2294.919	7826.591	FTA
Median income (in CHF)	52117.181	9778.047	FTA
Gini coefficient	0.442	0.066	FTA
<i>(f) Control variables: municipal typology^(*)</i>			
Tourist municipality (yes = 1)	0.023	0.151	FSO
Industrial municipality (yes = 1)	0.038	0.192	FSO

Notes: The table shows descriptives for the key variables (panel (a)) and the control variables (panels (b) to (g)) in the full estimation sample; all variables are at the municipality level. SBG = statistics of basic vocational education and training (“*Statistik der beruflichen Grundbildung*”); FSO = Federal Statistical Office (“*Bundesamt für Statistik*”); FFA = Federal Finance Administration (“*Eidgenössische Finanzverwaltung*”); FOT = Federal Office of Topography (“*Bundesamt für Landestopografie*”); ARE = Federal Office for Spatial Development (“*Bundesamt für Raumentwicklung*”). ^(*) denotes categorical variables for which only a few selected categories are shown.

Table A.2: Descriptive statistics for the LLBC sample, additional controls

Variable	Mean	SD	Source
(a) Key variables:			
Occupational segregation	0.728	0.066	SBG
Gender norms	0.469	0.111	FSO
Distance to the language border	3.544	26.145	ARE
(b) <i>Voter turnout and attitudes towards the role of the state</i>			
Voter turnout (mean across votes)	0.405	0.077	FSO
Attitudes towards the role of the state	25.416	7.142	FSO
(c) <i>Apprentices' demographics</i>			
# apprentices	75.897	118.934	SBG
Age at entry into apprenticeship (among apprentices)	18.094	1.014	SBG
Proportion female (among apprentices)	0.399	0.127	SBG
Proportion Swiss (among apprentices)	0.852	0.133	SBG
(d) <i>Selection into apprenticeship</i>			
(# contracts / # adolescents (15-21))	0.342	0.110	SBG/FSO
(# adolescents (15-21) / # individuals)	0.072	0.016	FSO
Distance to the nearest Gymnasium	15.571	11.642	ARE/FSO
(e) <i>Apprenticeship market</i>			
# trainees	91.005	203.267	SBG
# training firms	32.033	65.862	SBG
# distinct training occupations	20.547	22.402	SBG
(f) <i>Occupational structure</i>			
Proportion of contracts within a given occupation ^(*) :			
Commercial employee (yes = 1)	0.082	0.107	SBG
Health care assistant (yes = 1)	0.057	0.143	SBG
Information technician (yes = 1)	0.007	0.016	SBG

Notes: The table shows descriptives for the additional control variables used in table 8 in the main text; all variables are at the municipality level, and descriptives are for the LLBC sample only. SBG = statistics of basic vocational education and training (“*Statistik der beruflichen Grundbildung*”); FSO = Federal Statistical Office (“*Bundesamt für Statistik*”); ARE = Federal Office for Spatial Development (“*Bundesamt für Raumentwicklung*”). ^(*) denotes categorial variables for which only a few selected categories are shown.

Table A.3: Occupational aspirations versus occupational choices, weighted by the number of individuals

Data source	π_i				
	Kuhn and Wolter (2023)		SBG 2021/22		
N_j	−0.207** (0.099)	−0.185** (0.062)	−0.136*** (0.008)	−0.169*** (0.025)	−0.212*** (0.035)
Outcome	Aspiration	Choice	Choice	Choice	Choice
Controls	No	No	No	No	No
Region	Sample	Sample	Switzerland	Bern	Sample
N	24	24	2,107	325	24
R-Squared	0.004	0.004	0.211	0.165	0.540

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. All regressions in this table are weighted by the number of individuals. Robust standard errors, clustered by municipality, are given in parentheses. The study by Kuhn and Wolter (2023) covers a sample of individuals residing 24 different municipalities located in the German-speaking part of the canton of Bern. Column (3) is the same estimate as in the final column of table 6; column (4) restricts the sample to municipalities from the canton of Bern; column (5) restricts the sample further to those 24 municipalities covered in the analysis by Kuhn and Wolter (2023).

Table A.4: Comparison across different segregation measures

	π_j		f_j		D_j	
N_j	-0.110*** (0.014)	-0.034 (0.039)	-0.264*** (0.025)	-0.047 (0.068)	-0.034 (0.022)	0.009 (0.038)
Controls	No	Yes	No	Yes	No	Yes
N	2107	2107	2107	2107	2053	2053
R-Squared	0.041	0.116	0.063	0.179	0.001	0.531

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses. f_j corresponds to the fraction of contracts signed in an occupation with at least 70% of the apprentices having the same sex, i.e. $f_j = \frac{1}{N_j} \sum \mathbf{1}(\pi_i > 0.7) \cdot \mathbf{1}(i : \text{municipality} = j)$. D_j denotes the Duncan segregation index. These two alternative measures are often used in empirical studies on occupational segregation.

Table A.5: Local-language border contrast in gender norms, robustness

G_j	N_j											
	-0.132*** (0.016)	-0.121*** (0.018)	-0.134*** (0.014)	-0.133*** (0.012)	-0.114*** (0.021)	-0.105*** (0.025)	-0.140*** (0.017)	-0.136*** (0.016)	-0.046** (0.021)	-0.122*** (0.018)	-0.108*** (0.012)	
Full controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Bandwidth:												
Maximum	40	30	50	60	40	40	19	20	120	40	40	
Minimum	0	0	0	0	10	20	0	0	0	0	0	
Travelling mode	Public	Public	Public	Public	Public	Public	Public	Private	Public	Public	Public	
Distance metric	Time	Time	Time	Time	Time	Time	Distance	Time	Time	Time	Time	
Contrast	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-FR	DE-IT	DE-FR	DE-FR	
Bilingual cantons only	No	No	No	No	No	No	No	No	No	Yes	No	
Weights	No	No	No	No	No	No	No	No	No	No	Yes	
N	232	144	339	451	187	148	235	255	324	158	232	
R-Squared	0.836	0.858	0.817	0.804	0.861	0.879	0.822	0.828	0.867	0.801	0.911	

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses (the standard error in the final column is clustered by municipality). The distance metric is the minimum travelling time necessary to reach the closest municipality on the other side of the language border, either by public or private transport. The language-border contrast is either between the German- and the French-speaking regions (DE-FR) or between the German- and the Italian-speaking regions (DE-IT). The estimate in the final column is weighted by the absolute number of apprenticeship contracts, and its standard error is clustered by municipality.

Table A.6: Spatial regression estimates, additional control variables

	π_j						
N_j	-0.034 (0.039)	-0.033 (0.039)	-0.049 (0.037)	-0.041 (0.034)	-0.032 (0.036)	0.000 (0.033)	0.005 (0.033)
Controls:							
Full baseline controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Voter turnout	No	Yes	Yes	-	-	-	-
Attitudes towards the state	No	No	Yes	-	-	-	-
Apprentices' demographics	No	-	-	Yes	Yes	Yes	Yes
Selection into apprenticeship	No	-	-	No	Yes	Yes	Yes
Apprenticeship market	No	-	-	No	No	Yes	Yes
Occupational structure	No	-	-	No	No	No	Yes
N	2107	2107	2107	2107	2085	1964	1964
R-Squared	0.116	0.117	0.117	0.217	0.221	0.239	0.276

Notes: ***, **, and * denote statistical significance on the 1%, 5%, and 10% level, respectively. Robust standard errors are given in parentheses. See appendix tables A.1 and A.2 for additional details concerning the different sets of control variables.

B Setting up the local language-border contrast

In this part of the appendix, I discuss how I decided to set up the local language-border contrast (presented in section 4.3 in the main text).

B.1 Measuring the distance to the language border

First, I have to approximate the distance one needs to travel to cross the French-German language border, for which I use data from the Federal Office for Spatial Development (“*Bundesamt für Raumentwicklung*”). These data cover approximate average effective travelling distances and travelling times from one municipality to the other, for all possible pairs of municipalities, and for both private and public transportation.

To construct the travelling distance from a given municipality to the language border, I first split the sample according to whether they have a French- or German-speaking majority. Then, for each municipality located on the German-speaking part of the country, I look for the nearest municipality on the French-speaking side of the language border; and vice versa for the municipalities from the French-speaking part of the country.¹ Figure B.1 maps the spatial variation in this variable, showing that the variable closely matches the official classification of the municipalities according to their majority language (as shown in appendix figure A.2).²

Figure B.1

Moreover, and in contrast to Aepli *et al.* (2021a), note that I use travelling times by public transportation as the main distance metric, instead of distance by private car or motorcycle. I decided to use travelling times by public transport because most apprentices commute by public transport to their training firms (BFS, 2021), and it is therefore the more natural distance metric in the context of this study. At the same time, given the different distance metrics are all highly correlated with each other, and the choice of the distance metric is therefore unlikely to be crucial (see also table 7 in the main text for an empirical verification of this claim).

B.2 Choosing a baseline bandwidth

The main auxiliary parameter in the geographic quasi-experimental design is the bandwidth, which in turn determines the subset of municipalities used in that part of the analysis. In defining the baseline bandwidth around the language border, I essentially follow the procedure implemented by Aepli *et al.* (2021a):

¹This implies that the data do in fact not directly measure the distance to the language border per se, but the distance to the nearest municipality located on the other side of the language border. This creates some inherent fuzziness in the distance metric. While this would certainly be an issue in a conventional regression-discontinuity design, it is unlikely to be a relevant in the conditional language-border contrasts presented in section 4.3.

²See also appendix B in Aepli *et al.* (2021a) for additional details and a discussion of potential problems with this distance measure.

- In a first step, I determine a minimum bandwidth to ensure that municipalities along the (more or less) full length of the language border enter into the analysis sample.

Figure B.2

This step is graphically illustrated in figure B.2, and it suggests a minimum bandwidth of at least 20, or perhaps even 30 minutes. Note that quite a discontinuous group of municipalities results when choosing a smaller bandwidth.

- In a next step, I run a series of ancillary regressions to determine a maximum bandwidth, such that gender norms are approximately constant on either side of the language border. More specifically, I estimate the following regression equation:

$$N_j = \alpha_0 + \alpha_1 G_j + \alpha_2 d_j + \alpha_3 (G_j \cdot d_j) + \varepsilon_j \quad \forall j \in \{j : d_j \leq bw\}, \quad (\text{B.1})$$

for a bandwidth bw running from 160 minutes down to 20 minutes, in steps of 1 minute.³ For each of these regressions, I then test whether $\alpha_2 = \alpha_3 = 0$. I find that gender norms become approximately constant on either side of the language border once the bandwidth is chosen to be about 40 kilometers or less (see panel (a) figure B.3 below).

Figure B.3

Importantly, however, and as discussed in section 4.3 in the main text, the two language regions do differ significantly in gender norms on average (i.e. $\alpha_1 \neq 0$, even for small bandwidths); this is shown in panel (b) of figure B.3, which plots the estimated parameter $\hat{\alpha}_1$ over the different bandwidths (along with the 95% confidence band).

- In a third step, I check whether the additional covariates become better balanced when choosing a narrower window around the language border. To this end, I run an additional series of ancillary regressions, this time of the following form:

$$G_j = \alpha_0 + \alpha_1 X_j + \varepsilon_j, \quad \forall j \in \{j : d_j \leq bw\}, \quad (\text{B.2})$$

with the dependent variable now being the dummy variable indicating whether a municipality is located on the German- or the French-speaking side of the language border (i.e. $G_j = 1$ again denotes that municipality j is located on the German-speaking side of the language border). For each regression, I then test whether $\alpha_1 = 0$ (note that α_1 denotes a parameter vector here). See panel (c) of figure B.3 for the resulting pattern in the associated F-statistic.

As expected, this exercise shows that the smaller the chosen bandwidth, the smaller the

³A bandwidth larger than 160 minutes would only add regions on the German-speaking side of the language border as there is no French-speaking region farther away than that.

resulting F-statistic. However, because choosing a narrower bandwidth also shrinks the sample size, there is an obvious trade-off between precision and bias involved here.⁴

- As a by-product, moving the focus closer to the language border also implies that the number of municipalities from the two regions becomes more balanced (however, the design does not require that there be an equal number of observations from either side). This is shown in panel (d) of figure B.3, which plots the ratio of observations from the German-speaking regions to the number of observations from the French-speaking side of the language border (using the left-hand y-scale).

The figure also shows the number of different cantons in which the sample municipalities are located (right-hand y-scale). Moving closer to the language border also implies that the sample is drawn from a smaller set of cantons, which eliminates some unobserved heterogeneity (due to, for example, institutional differences in educational policies between the different cantons).

Based on the above results, I decided to choose a travelling time by public transportation of 40 minutes or less as the baseline bandwidth in section 4.3. However, I also show in that section that the main estimate is robust with respect to the exact choice of bandwidth (within the boundaries suggested by the results from this appendix; cf. tables 7 and A.5).


Moreover, note that these additional results also show that explicitly controlling for observable confounders (see also Lonati *et al.*, 2024) remains necessary. While the municipalities become more similar if the bandwidth becomes smaller, differences between the municipalities from the two sides of the language regions with regards to observable characteristics remain. This reflects the simple fact that even those municipalities close to the language border are not simply randomly assigned to either side of the border (as in a conventional regression-discontinuity design).

References

- Aeppli, M., Kuhn, A., and Schweri, J. (2021a). Culture, norms, and the provision of training by employers: Evidence from the Swiss language border. *Labour Economics*, **73**, 102057.
- BFS (2021). Pendlermobilität in der Schweiz 2019. Mit einer Vertiefung zur Länge der Arbeitswege. Neuchâtel: Bundesamt für Statistik (BFS).
- Lonati, S., Lalive, R., and Efferson, C. (2024). Identifying culture as cause: Challenges and opportunities. *Evolutionary Human Sciences*, pages 1–66.


⁴This would also become evident at the lower, truncated end of the x-axis in panel (c) of figure 6, where there is in fact a sudden and large increase in the robust F-statistic.

Figure B.1: Minimum travelling time to the French-German language border




Notes: The figure maps the minimum travelling time, by public transportation, to the nearest municipality located on the other side of the French-German language border. Regions shaded in light blue (dark blue) are located closer to (farther away from) the language border. In drawing this map, note that all municipalities without a French-speaking majority are coded as German-speaking (or, rather, as non French-speaking).

Figure B.2: Choosing the bandwidth around the language border



Notes: The figure shows the set of municipalities that make up the analysis sample on either side of the French-German language border for a bandwidth varying between 60 and 10 minutes. Note that the maps focus on the subset of cantons located along the French-German language border. The figures zoom in on the cantons along the French-German language border (i.e. Basel-Land, Basel-Stadt, Bern, Fribourg, Jura, Neuchâtel, Solothurn, Valais, and Vaud).

Figure B.3: Defining the baseline bandwidth for the language-border contrast



Notes: The figure shows various statistics used to determine the baseline bandwidth used in section 4.3 in the main part of the paper. Panel (a) plots the F-statistic associated with the null hypothesis of constant gender norms on both sides of the language border; panel (b) shows estimated difference in gender norms between the two language regions, along with the 95% confidence band; panel (c) shows the F-statistic associated with the null hypothesis of no overall difference in the observable control variables between the two regions; panel (d) shows the ratio of the number of municipalities from the German-speaking part to the number of municipalities from the French-speaking part along with the number of different cantons in which the sample observations are located.