

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Things versus People: Gender Differences in Vocational Interests and in Occupational Preferences

Andreas Kuhn^{a,b,c,1,*}, Stefan C. Wolter^{b,c,d,e}

^a Swiss Federal University for Vocational Education and Training, Zollikofen, Switzerland

^b Centre for Research in Economics of Education, University of Bern, Switzerland

^c Institute for the Study of Labor (IZA), Bonn, Germany

^d Swiss Coordination Centre for Research in Education, Aarau, Switzerland

^e Center for Economic Studies (CESifo), Munich, Germany

ARTICLE INFO

Article history:

Received 21 November 2021

Revised 24 June 2022

Accepted 2 September 2022

JEL classification:

J16

J24

D91

Keywords:

Occupational choice

Occupational segregation

Things versus people

Preferences

Gender differences

Job content

ABSTRACT

Occupational choices remain strongly segregated by gender, for reasons not yet fully understood. In this paper, we use detailed information on the cognitive requirements in 130 distinct learnable occupations in the Swiss apprenticeship system to describe the broad job content in these occupations along the things-versus-people dimension. We first show that our occupational classification along this dimension closely aligns with actual job tasks, taken from an independent data source on employers job advertisements. We then document that female apprentices tend to choose occupations that are oriented towards working with people, while male apprentices tend to favor occupations that involve working with things. In fact, our analysis suggests that this variable is by any statistical measure among the most important proximate predictors of occupational gender segregation. In a further step, we replicate this finding using individual-level data on both occupational aspirations and actual occupational choices for a sample of adolescents at the start of 8th grade and the end of 9th grade, respectively. Using these additional data, we finally show that the gender difference in occupational preferences is largely independent of a large number of individual, parental, and regional controls.

© 2022 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Introduction

Many countries have seen profound and often surprisingly rapid changes in various measures of women's participation and performance in both the educational system and in the labor market, paralleled by corresponding changes in individuals'

* Corresponding author.

E-mail addresses: andreas.kuhn@ehb.swiss (A. Kuhn), stefan.wolter@unibe.ch (S.C. Wolter).

¹ We thank Manuel Aepli, Anne Brenøe and Jürg Schweri for comments on an early draft of the paper as well as seminar participants in Zurich and two anonymous reviewers for helpful comments and suggestions. We also thank Walter Götze for providing us with the raw data describing the cognitive requirements of the various occupations, Christoph Düby from the cantonal department of education of the canton of Bern for providing us with the data on the population of apprenticeship contracts in the canton of Bern, Miriam Grønning for helping us with the occupational crosswalks used to merge the different data sets, and Sally Gschwend-Fisher for proofreading the manuscript. Corresponding author: Andreas Kuhn, Swiss Federal University for Vocational Education and Training, Kirchlindachstrasse 79, 3052 Zollikofen, Switzerland

attitudes towards the equality between women and men. Nonetheless, there remains a persistently high degree of occupational gender segregation, even in the most progressive countries in this regard (see, for example, Charles and Grusky, 2004) – and often despite explicit and considerable public effort to decrease the extent of occupational segregation.² These remaining differences in occupational choice show up in women's underrepresentation in STEM jobs (e.g. Kahn and Ginther, 2017), for example, and they are related to the remaining gender gap in wages (e.g. Blau and Kahn, 2017; Olivetti and Petrongolo, 2016). Moreover, both the legitimacy and the potential effectiveness of measures directed towards establishing a more equal gender balance across occupations are directly related to the underlying causes of the observed difference in occupational choices. Therefore, understanding the causes of the persistence in occupational gender segregation is of both academic and public interest (Cortes and Pan, 2018; Bertrand, 2011; Williams and Halsey, 2021).

A first potential explanation starts from the observation that men and women differ in their economic preferences and psychological traits that may influence their occupational choices. Indeed, a voluminous and growing number of empirical studies documents gender differences in various economic preferences (Croson and Gneezy, 2009; Bertrand, 2011) and across many countries (Falk et al., 2018).³ Moreover, a few studies have explicitly studied whether these gender differences in preferences are associated with differences in occupational choices. For example, Bonin et al. (2007) showed that more risk-tolerant individuals select into occupations with a higher earnings risk. In a closely related study, Fouarge et al. (2014) showed that both risk preferences and time preferences are related to occupational choices, with more patient individuals choosing occupations with a steeper earnings profile later on. In the context of occupational choice, however, competitiveness has presumably received the most attention from economists. In one prominent study on the subject, Buser et al. (2014) found that the gender difference in competitiveness can explain some of the gender gap in choice of study subject in the Netherlands. Other studies have found similar results, such as Kleinjans (2009) for Denmark or Buser et al. (2017), who estimate the effect of competitiveness on study choices among Swiss college students (see also Buser et al., 2022, who extend the analysis to non-college students). Differences along personality dimensions are also well documented empirically, such as in agreeableness or neuroticism (e.g. Weisberg et al., 2011; Giudice et al., 2012), and they become large when the different dimensions are considered simultaneously (Del Giudice, 2019); see also Archer (2019), who provides an overview across a large number of traits and measures. The differences in personality between men and women also hold up across many different countries (e.g. Schmitt et al., 2008). Similar to economic preferences, some personality traits have been shown to be associated with individuals' occupational choices (e.g. Lee et al., 2015; John and Thomsen, 2014; Coenen et al., 2021; Taber et al., 2011), and there also appears to exist overlap across different concepts and measures (see, for example, Houston et al., 2015).

A closely related and partially overlapping, and in some way even more obvious hypothesis is that men and women may differ in their occupational choices because they have different vocational interests, i.e. that they differ in their preferences over the general task content within a given occupation. Indeed, there is a considerable amount of empirical evidence, mainly from psychologists, consistently documenting large and persistent gender differences in vocational interests, especially along the things-versus-people dimension, an idea that goes back at least to the work of Holland (1959).⁴ In this and other publications, Holland (1959) has set out a multidimensional conceptual typology of vocational personalities, such as investigative or entrepreneurial personalities. Interest in working with inanimate things and working with other people, respectively, are the two dimensions of interest to us, and it has been suggested that they may describe interests along the same bipolar dimension, as first argued by Prediger (1982).⁵ A representative study along these lines is Morris (2016), who describes vocational interests for about 1.28 million individuals aged 14 to 63 from the US, surveyed between 2005 and 2014. He finds a large gender difference in vocational interests between men and women, mainly along (but not restricted to) the things-people dimension, but only small to very small associations with either age, year or ethnicity. Recent empirical studies by Lordan and Pischke (2022) and by Gelblum (2020) confirm this regularity, showing that women and men have different preferences regarding job content. Moreover, empirically, it appears that this specific dimension of vocational interests can best discriminate between tasks men and women prefer, respectively (e.g. Lippa, 1998). Interestingly, this specific dimension of vocational interests may also explain a substantive part of the gender segregation within relatively narrow groups of occupations, such as within STEM occupations (e.g. Su and Rounds, 2015; Thelwall et al., 2019). Moreover, Lippa et al. (2014) show that the explanatory power of an occupation's task content along the things-people dimension has

² In fact, several cross-country studies have consistently documented that gender differences in economic preferences, as well as in personality traits, tend to increase, rather than decrease, in countries with more progressive views towards equality between men and women (Falk and Hermle, 2018; Schmitt et al., 2008; Stoet and Geary, 2022). The same holds true for occupational gender segregation (e.g. Charles, 2017; Charles and Bradley, 2009). In contrast, gender differences in academic achievement do not appear to be consistently related to measures of gender equality (e.g. Stoet and Geary, 2015).

³ More specifically, empirical research has documented gender differences in time preferences (e.g. Ditttrich and Leipold, 2014), risk preferences (e.g. Charness and Gneezy, 2012), social preferences (e.g. Kamas and Preston, 2015), and competitiveness (e.g. Niederle and Vesterlund, 2011), to name but the most prominent examples.

⁴ As laid out in Zickar and Min (2019), efforts to measure vocational interests dates back to the late 19th century. See also Nauta (2010) on a discussion of Holland's influential conceptualization of vocational interests. Empirically, it turns out that observed gender differences in vocational interests along the things-people dimension are among the largest differences between men and women measured by any psychometric instrument. For example, in their meta-study on the subject, Su et al. (2009) document an average effect size of 0.93. Similarly, Lippa (2010) reports an average effect size of 1.18. As we will show in Section 4.4 below, we find a gender difference in vocational interests along the things-versus-people dimension of a comparable size.

⁵ More precisely, Prediger (1982) argued that the six different vocational types or personalities (i.e. scales) proposed by Holland (1959) can be mapped onto four dimensions, which in turn can be represented by the two bipolar dimensions of things versus people and data versus ideas (see also Morris, 2016).

become more, not less, important over time in the US labor market, suggesting that the gender difference in occupational preferences has been rather stable over time.

Moreover, the psychological literature suggests that the gender difference in interest towards things and towards people may run much deeper than just reflecting differences in vocational interests. Specifically, it has been argued that one's orientation towards things or towards other people is a much more encompassing trait, which manifests itself also in everyday behavior, such as in book preferences or even in the objects individuals are looking at (Graziano et al., 2011; Intyre and Graziano, 2019; 2016). And then there is another, also closely related literature on the distinction between individuals' systemizing skills, i.e. their propensity towards seeking patterns and being interested in understanding both natural and technical systems, and empathizing, i.e. the degree to which people are able to navigate the social world without much effort, which also emphasizes that individuals may be differently predisposed to and apt in dealing with these two different domains (e.g. Groen et al., 2018; Cohen, 2004). Both these concepts, which also overlap to some extent, have been shown to be associated with occupational choices (Svedholm-Häkkinen and Lindeman, 2016; Yang and Barth, 2015; Greenberg et al., 2018; Wright et al., 2015). Again, see also the overview by Archer (2019), which suggests that gender differences in the different measures associated with how strongly individuals are oriented or interested towards things versus people are among the largest documented differences between men and women across all documented traits and measures.

In this paper, we add further empirical evidence to this important and fascinating debate, using a unique combination of different data sources to describe gender differences in occupational preferences among Swiss adolescents. Specifically, we use detailed information on the cognitive requirements in 130 apprenticeship occupations to describe occupations' broad task content along the things-versus-people dimension. Importantly, note that we directly describe occupations, not individuals' subjective interests, deviating from the original, mainly psychological literature in this regard (more on this in Section 3.1 below). In a first step, however, we show that occupational aspirations are indeed highly segregated, consistent with analogous findings for the country as a whole (e.g. Aepli et al., 2019). We then show that the proportion of men and women, respectively, in an occupation is strongly correlated with its position on the things-versus-people dimension. In fact, we find that the broad task content of an occupation is an extremely powerful predictor of whether males or females predominantly choose the occupation. In the second part of our empirical analysis, we replicate this finding using individual-level data for a sample of adolescents from the canton of Bern. The majority of these individuals were surveyed twice, at the start of 8th grade as well as at the end of 9th grade. As we will explain below, this implies that the survey contains information on participants' occupational aspirations at the time they started the vocational selection process as well as about their actual occupational choices later on. This allows us not only to see whether we can replicate the analysis based on the occupational level data, but also to show that the association between gender and the task content of an occupation shows up already in adolescents' occupational aspirations (and thereby mitigating the concern that the difference in actual occupational choices is driven by external factors, such as employers' discriminatory behavior). Moreover, using these additional data, we are also able to check whether the gender difference in occupational preferences is robust to the inclusion of additional controls at the individual, parental and regional level.

The remainder of this paper proceeds as follows. In the following section, we shortly describe some of the key features of the Swiss educational system, mainly focusing on the apprenticeship system. Section 3 discusses the different data sources used in the empirical analysis. Our empirical analysis is presented and discussed in several consecutive steps in Section 4. In the first part of the analysis (Sections 4.1 to 4.3), we analyze data at the occupational level. In the second part of the analysis (Section 4.4), we use individual-level data to replicate and expand on the main finding from the analysis at the occupational level. Section 5 concludes.

2. The Swiss educational system

Compulsory schooling in Switzerland lasts eleven years, of which two years are spent in kindergarten, six in primary school, and three in secondary school (see SCCRE, 2018, for a detailed description of the Swiss educational system). Children usually enter primary school in the year they turn seven years old, and thus most of them are round 15 years old when they finish the mandatory part of education. Consequently, they usually enter post-mandatory, upper-secondary schooling/training in the year they reach the age of 16.

At the upper-secondary level, there are two main options (see appendix Fig. A.1 for a schematic illustration). First, there is the possibility to enter further general education via a baccalaureate school ("Gymnasium" in German, about equivalent to high school), which prepares for and gives access to further studies at the tertiary (usually university) level. The other, far more popular route at this stage is to enter the apprenticeship system, which also gives opportunities to enter further education and training at the tertiary level later on.

2.1. The apprenticeship system

At the national level, a large majority of about 62% of adolescents eventually enters the apprenticeship system, typically through a dual-track apprenticeship which combines practical training within a firm with vocational school (e.g. SERI, 2022).⁶ In the canton of Bern, from where our sample of adolescents is drawn (see Section 3 below for details), the

⁶ See Wettstein et al. (2017) for a detailed discussion of the Swiss apprenticeship system.

proportion of adolescents entering the apprenticeship system is actually higher than the national average, with almost 70% of recent cohorts entering the apprenticeship system (SCCRE, 2018).

Most apprenticeships last three or four years, and one day per week (in some cases two days per week) are spent in vocational school (SERI, 2022). There are also two-year apprenticeships for adolescents with lower academic standards (see also Section 4.2). Overall, there are about 240 different learnable occupations within the Swiss apprenticeship system. It is also worth noting that the apprenticeship system is regulated at the national level, in contrast to the rest of the educational system (where both the cantons and the municipalities play the lead role; that is also one of the main reasons why there are large regional differences in fraction of adolescents entering vocational training versus further general education at the upper-secondary level).

Moreover, there is basically a market for apprenticeship positions, with adolescents looking for apprenticeship positions and with employers simultaneously advertising vacant positions (furthermore, apprentice wages may differ across employers for the same occupation). In case of a mutual match, the host company and the apprentice both sign an apprenticeship contract, which lasts until the completion of the apprenticeship. This implies that an apprenticeship does not automatically lead to an employment contract with the training firm, even though many training firms retain their apprentices after they have successfully finished their apprenticeships. During their training period, apprentices are paid an apprentice wage, which is substantively lower than that of a fully trained worker in the same occupation.

3. Data

We next describe the different data sources used in our empirical analysis.

3.1. Describing the task content of occupation

We use two different and independent data sources to describe the task content of occupations.

Cognitive requirements by occupation

First, we use data on the cognitive requirements in the different learnable occupations within the Swiss apprenticeship system. These data come from a project initiated, and partially funded, by the Swiss State Secretariat for Education, Research and Innovation (SERI), and with the primary goal of providing adolescents and their parents, as well as people working within the VET system (such as teachers in vocational schools or vocational advisers) with comparable information on the cognitive requirements in the different learnable occupations (labelled “Anforderungsprofile” in German).⁷ For this purpose, the data contain quite detailed requirements in native and foreign languages, mathematics, and natural sciences for the various learnable occupations within the Swiss apprenticeship system. Within each of the four subjects, there is a more detailed breakdown by subtopics.⁸ However, we will almost exclusively focus on the aggregated scores by main topic in our own analysis (we do use some of the subcategories for a validation exercise; see Section 4.1 below). This is our main data source for describing the task content of occupations, as explained in more detail in Section 4.1 below.

That being said, note that our approach is fundamentally from, but at the same time complementary to the approach typically chosen in the (mainly) psychological literature, which focuses on the psychometric measurement of individuals’ (stated) vocational interests (for a short introduction, see Chernyshenko et al., 2019). Practically, this is done using a large and tested battery of survey items, such as the “Strong Interest Inventory” (e.g. Donnay and Borgen, 1996; Donnay et al., 2005). In contrast, our approach focuses on variables that describe the cognitive requirements judged necessary to enter a given learnable occupation, i.e. variables that essentially describe occupations’ task content and overall difficulty. In other words, our approach focuses on adolescents’ vocational interests as revealed by their occupational aspirations and actual occupational choices. Both approaches are complementary in that they both assume that individuals self-select into the various occupations at least in part based on the match between their vocational interests and the actual task content of an occupation (see, for example, Nye et al., 2019, on this issue). Of course, one potential issue with such an approach is that the evidence with regard to the influence of vocational interests remains indirect, and we will come back to this issue in Section 4.1 below.

Task descriptions from actual job postings

We complement this information with data from actual job advertisements, taken from an additional and independent source of data collected for the main purpose of monitoring the demand side of the Swiss job market over time (the Swiss Job Market Monitor, “Stellenmarktmonitor”).⁹ These data contain samples of actual job advertisements by both private and

⁷ More information about these data is available online (however, only in German, French, and Italian) at: www.anforderungsprofile.ch. Appendix Fig. A.2 shows an example comparison between two popular occupations (healthcare assistant versus mechanical engineer) as available directly from the website.

⁸ For example, in mathematics, there are the following five subtopics: algebra (“Zahl und Variable”), geometry (“Form und Raum”), units of measurement (“Größen und Masse”), calculus (“Funktionale Zusammenhänge”) and statistics (“Daten und Zufall”). See again appendix Fig. A.2 for a concrete example.

⁹ Additional details are available online from the project website: www.stellenmarktmonitor.uzh.ch.

Table 1
The occupations ranked highest/lowest along the things–people dimension.

Rank	Occupation	C_j^{things}
<i>(a) The ten most things-oriented occupations</i>		
1.	Plant and apparatus engineer (“Anlagen- und Apparatebauer”)	2.260
2.	Design engineer (“Konstrukteur”)	2.084
3.	Mechanical engineer (“Polymechaniker”)	2.052
4.	Automatic technician (“Automatikmonteur”)	2.027
5.	Mechanical technician (“Produktionsmechaniker”)	2.027
6.	Gunsmith (“Bchsenmacher”)	1.974
7.	Insulation contractor (“Isolierspengler”)	1.933
8.	Automation engineer (“Automatiker”)	1.884
9.	Carpenter (“Schreiner”)	1.538
10.	Industrial ceramist (“Industriekeramiker”)	1.517
<i>(b) The ten most people-oriented occupations</i>		
1.	Podiatrist (“Podologin”)	-3.049
2.	Healthcare assistant (“Fachfrau Gesundheit”)	-2.715
3.	Information and documentation specialist (“Fachfrau Information und Dokumentation”)	-2.634
4.	Certified social care worker (“Fachfrau Betreuung”)	-2.590
5.	Druggist (“Drogistin”)	-2.538
6.	Pharmacy assistant (“Pharma-Assistentin”)	-2.234
7.	Hairdresser (“Coiffeuse”)	-2.233
8.	Bookseller (“Buchhndlerin”)	-2.218
9.	Hairdresser (“Coiffeuse”)	-2.136
10.	Customer dialogue specialist (“Fachmann Kundendialog”)	-2.072

Notes: The table shows the ten occupations with the highest (lowest) scores on C_j^{things} . The official German description is given in parentheses, along with the English translation suggested by the State Secretariat for Education, Research and Innovation (SERI), where available (otherwise we chose our own translation). We show the female (male) description in German if an occupation is chosen by a majority of female (male) apprentices.

public employers from the years 1950 until 2018 (currently, the data collection is still going) and sampled from newspapers, company websites, as well as online job portals (Buchmann et al., 2020).

For our purpose, we focus on the more recent advertisements from the years 2010 to 2015, with a total of 24,368 individual job postings from all over Switzerland. Among other things, the data record the main activity of each job posting.¹⁰ We use this variable to validate our approach of categorizing the different occupations along the interest in things versus people dimension based on the data describing the cognitive requirements (see Section 4.1 below for details).

3.2. Describing the extent of occupational gender segregation

Moreover, we have also access to a dataset that contains the full population of individual-level apprenticeship contracts involving either apprentices and/or employers from the canton of Bern as of August 2014, with more than 45,000 individual apprenticeship contracts covered. We use these data to compute the proportion of females by occupation, which in turn allows us to describe the extent of occupational gender segregation. We can do this with reasonable precision for most occupations because the data cover so many individual-level contracts. Nonetheless, some occupations are still only rarely chosen, and we exclude these occupations from most of the analysis, as discussed in more detail below.¹¹

3.3. Individual-level survey among adolescents

Our final data source is a computer-assisted individual-level survey among 1,514 adolescents at the start of 8th grade (with an average age of about 14 years at the time the main survey was administered) from the German language area of the canton of Bern (see Buser et al., 2017, for additional details). The adolescents are from 28 different schools spread across 24 different municipalities in the German language area of the canton of Bern.¹² The main survey was administered in the autumn of 2013, and a majority of adolescents (about 96%) was successfully contacted a second time in late spring of 2015, at the end of 9th grade (i.e. at the end of compulsory education).

¹⁰ The variable containing the main activity has 21 different values designating broad groups of tasks, such as “planning, engineering, designing/drawing” or “educating/teaching, advising”, and it is available from 1995 onwards only (see Table 2 below for the full list of activities coded). Moreover, the data also contain several occupational codes, allowing us to merge the two datasets.

¹¹ It is worth noting that the distribution of individuals across occupations is highly skewed, with a large proportion of all apprenticeship contracts concentrated in relatively few occupations only. This is clearly evident from appendix Table A.1, which shows the ten most popular occupations among girls and boys, respectively. Among boys (girls), the ten most popular occupations account for 42.3% (65.3%) of all apprenticeship contracts.

¹² These data have been used before to study whether competitiveness has an influence on adolescents’ occupational aspirations (Buser et al., 2017; 2022). Jaik and Wolter (2019) look at the correlation between occupational aspirations and occupational choices. Finally, Kuhn and Wolter (2022) use the same data to study the impact of societal gender norms on gender-stereotypical occupational aspirations.

Table 2

Validating our classification of occupations using information from actual job postings.

	C_j^{things}	
Agricultural tasks	-0.460*	-0.399
Manufacturing	0.125	0.243
Installation, assembly, construction	0.875**	1.045***
Set up, operate	0.669	0.902
Repair, restore	0.823***	0.843***
Store and transport	0.352***	0.381***
Purchasing/sales, cashier, customer service	-2.101***	-2.341***
Writing, correspondence, administration	-4.958***	-5.395***
Accounting and finance	7.292	9.431
IT, programming	-0.742	-0.967
Hospitality services	-2.977***	-3.012***
Ironing, cleaning, waste management	-0.116	-0.017
Guarding	-4.339	-1.952
Analysis/research, controller	-0.481	-1.082
Planning, engineering, designing/drawing	1.918***	1.530***
Supervising, hiring	9.892	17.444
Disposing, organizing, leading	3.282	2.355
Educating/teaching, advising	-4.056*	-4.781**
Administration of justice	-1.569	3.343
Medical and cosmetical care	-1.944***	-2.053***
Publishing, creative work	-0.881	-1.027
C_j^{demand}		0.152***
Number of observations	130	130
R-Squared	0.673	0.698
Adjusted R-Squared	0.610	0.636

Notes: *, **, *** denotes statistical significance on the 10%, 5%, and 1% level, respectively. The table shows point estimates from a regression suppressing the constant term (robust standard errors are not shown to keep the table compact). The labels for the activities are carried over from the “Stellenmarktmonitor” data (Buchmann et al., 2020).

Table 3

Discriminating between two-year and regular (i.e. three- or four-year) apprenticeships.

	Two-year apprenticeship (yes = 1)	
C_j^{demand}	-0.182*** (0.019)	-0.183*** (0.018)
C_j^{things}		-0.096*** (0.019)
Number of observations	130	130
R-Squared	0.492	0.573
Adjusted R-Squared	0.488	0.567

Notes: *, **, *** denotes statistical significance on the 10%, 5%, and 1% level, respectively. Robust standard errors are given in parentheses. The dependent variable is a dummy, taking on the value of 1 if an occupation is learnable through a two-year apprenticeship (and 0 otherwise), and thus explicitly tailored towards academically weaker adolescents.

The first and main round of the survey was administered in class and during school hours, and it took about one school lesson to complete the survey in class. A small team of research assistants and students went to each class who participated, equipped with a set of laptops, and the adolescents could thus fill out the survey electronically. This also allowed to implement experimental and incentivized measures of both competitiveness and risk preferences. The second, significantly less extensive part of the survey, which mainly focused on individuals' actual choices already taken by the end of 9th grade, was administered as a short paper questionnaire of two pages only. It was also administered in class and during school hours, which is why attrition between the two rounds of the survey was so low.

Occupational aspirations and occupational choices among adolescents

In the first and main round of the survey, adolescents were asked about their occupational aspirations, i.e. adolescents were simply asked in which occupation they would like to work (What apprenticeship would you most like to complete?).

This information is in raw-text form, but we assigned an occupational code to each occupation that is learnable through an apprenticeship occupation. The adolescents were further asked about their actual occupational choices in a second, later round of the survey. At this stage, most adolescents tend to have chosen their occupation and already have signed an apprenticeship contract with their prospective employer and instructor (or have decided instead that they want to go on with further general education, in which case there is no occupational choice to be made yet).

In the empirical analysis reported below, we will report results for both actual choices and aspirations. The data on individual-level occupational choices obviously allow us to replicate the analysis based on the occupational-level data. At the same time, occupational aspirations are also of interest because they are arguably not (or less) influenced by external restrictions on occupational choice, such as by the availability of apprenticeship positions in the desired occupation, or by their prospective employers and/or their parents. On the other hand, however, it is often argued that people are more prone to be influenced by societal norms when they are younger.¹³ This allows us to show that the results using actual choices, both at the individual and occupational levels, are not simply driven by such external forces.

4. Empirical analysis

The primary aim of our analysis is to see whether there are gender differences in occupational preferences that align with differences in the task content of the occupations. In a first step, however, we need to classify the different occupations according to their task content, along the things-versus-people dimension.

4.1. Describing the task content of occupations

We therefore first describe how the learnable occupations differ in their broad task content.

The task content of occupations: things versus people orientation

We first run a principal-components analysis to describe the task content of occupations (e.g. James et al., 2013; Gorsuch, 2014). More specifically, we use the four variables describing the cognitive requirements in mathematics, natural sciences, as well as native and foreign languages as input variables into the analysis. The results show that the first three principal components can reproduce almost 98% of the overall variation in the four original variables (cf. appendix Table A.2). Also note that, by construction, the resulting principal components are uncorrelated with each other at the occupational level.

Not surprisingly, the first principal component (PC) loads on all four input variables, which suggests that this PC can be interpreted as the overall level of cognitive demands in a given occupation, and we will thus use the shorthand C_j^{demand} for this variable subsequently. Indeed, comparing the occupations with the highest and lowest values on this PC supports this interpretation (cf. appendix Table A.3). The ten occupations with the lowest overall cognitive requirements, for example, are exclusively apprenticeships that only last for two years (i.e. these are occupations which were intentionally set up for academically underachieving youths), such as timber worker or tire work assistant. At the upper end of the distribution of this variable, on the other hand, we find occupations such as optometrist or geomatics expert. A more formal validation check concerning this variable is presented in the empirical part (see Section 4.2 below).

In what follows, however, we will mainly focus on the second and the third PCs, which together describe the task content of the occupations along the interest in things versus interest in people dimension, as we will now argue. The second PC loads positively on mathematics and natural sciences as well as negatively on both native and foreign languages, while the third PC loads positively on mathematics and foreign languages, but negatively on natural sciences and native languages. Thus the second PC classifies occupations according to whether they are oriented towards the importance of mathematics and sciences rather than towards the usage of languages, and we use the shorthand C_j^{math} for this PC in what follows. The third PC is more ambiguous than the previous two, but we think that one may interpret this PC as indicating whether an occupation deals with things or people in an either more abstract or a more direct and practical way, and we thus denote this PC as C_j^{abstract} subsequently, but we acknowledge that the interpretation of this PC is less clear-cut than for the first two.¹⁴ Occupations oriented towards things are presumably tilted towards both mathematics/sciences as well as abstract content, and we thus combine the second and the third PCs into a common variable, simply by summing them up.¹⁵ In what follows, we use the shorthand C_j^{things} (i.e. $C_j^{\text{things}} = C_j^{\text{math}} + C_j^{\text{abstract}}$) for this variable, with larger values

¹³ See also Jaik and Wolter (2019), who study how occupational aspirations in 8th grade deviate from actual occupational choices later on using the same data. Kuhn and Wolter (2022) find a null effect of societal gender norms on both occupational aspirations and actual occupational choices. Together, the two studies suggest that adolescents' aspirations may often differ from their choices, without influencing the gender-stereotypicality of their preferences, however. Our results are also in line with this conclusion.

¹⁴ The fourth principal component has no obvious interpretation and is not used in the empirical analysis at all (also note that this PC explains only a negligible fraction of the variation in the input variables; cf. appendix Table A.2).

¹⁵ Remember that the two PCs are independent of each other at the occupational level by construction. This implies that C_j^{things} still has a mean of zero, and a variance that equals the sum of the variances of the two variables that are added up.

indicating that the broad task content in an occupation is tilted towards working with inanimate things, rather than with other people.¹⁶

A look at some specific occupations suggests that this classification is both plausible and meaningful. Table 1 therefore lists the ten occupations with the largest and with the smallest values on C_j^{things} , respectively. Not surprisingly, technical occupations such as mechanical engineer or carpenter are found on the things-intense tail of the distribution of this variable, while occupations such as healthcare assistant or hairdresser are found on the opposite end of the distribution (see Appendix B for a more formal evaluation).¹⁷

Further note that the statistical separation between occupations' overall demand level and their task content along the things-versus-people dimension implies another, more subtle difference between our approach and many of the preceding empirical studies on the subject. Specifically, by separating the vertical (i.e. demand level) from the horizontal (i.e. task content) dimension of occupational choice we can more clearly focus on the horizontal dimension of occupational choice. This is interesting for at least two reasons. First, empirically, the horizontal dimension turns out to be much more relevant than the vertical dimension in predicting the proportion of male and female apprentices choosing a given occupation (as shown in Sections 4.3 and 4.4 below). Second, at the conceptual level, the distinction between a horizontal and a vertical dimension of occupations could also help in better understanding why and how certain characteristics, such as vocational interests, have an influence on individuals' occupational choices.

4.2. Validating our classification of the task content of occupations

We next provide two validation checks to support our occupational classification.

Classification along the things-versus-people dimension

We first use the data on employers' actual job postings to validate our classification of the occupations along the main dimension of interest (i.e. things-versus-people) because one might raise the objection that the data on the cognitive requirements may not have much to do with what people working in these occupations actually do. Using the additional data from employers' job postings allows us to check empirically whether our classification of the occupations based on the cognitive requirements alone roughly corresponds to what employers really expect people actually working in these occupations to do.

Table 2 describes the full correlational pattern by regressing our measure of the task-intensity along the things-people dimension C_j^{things} on a set of 21 variables, each essentially measuring the relative frequency with which a given main activity was explicitly mentioned in a real job posting (i.e. each of these variables may take on values within the unit interval). The first column of Table 2 shows estimates without controlling for the overall demand level in an occupation, while the second column adds C_j^{demand} as a regressor (because the task data from the SJMM may also reflect differences in the demand level of an occupation). Both columns show that the value on the things-versus-people variable of an occupation closely aligns with the actual activity pattern in an occupation as made salient through job postings by employers looking for individuals to work in these occupations. Correspondingly, the resulting R-squared is very high, 0.673 without and 0.698 with the inclusion of C_j^{demand} as a regressor. Given that we use two fully independent sources of data, we view these results as quite a strong confirmation that our occupational classification based on the cognitive requirements is meaningful and that it picks up real differences between the different occupations.¹⁸

The overall level of cognitive requirements

A complementary validation check focuses on the overall demand level of an occupation (i.e. C_j^{demand}) rather than on its task content. To implement this test, we focus on the institutional feature that there exist two-year apprenticeships that are specifically designed for adolescents with lower academic abilities ("regular" apprenticeships last either three or four years). The successful completion of these apprenticeships also leads to a nationally recognized diploma, and successful completion of such an apprenticeship opens up further educational possibilities.¹⁹ Therefore, in this context, we expect that C_j^{demand}

¹⁶ In Appendix B, we provide a data-based rationale for combining the two PCs into a single task measure, and we also deal with potential objections to this specific parameterization. Moreover, we also provide some additional robustness checks concerning this variable. Importantly, we show that our main result is not driven by the decision to combine the two PCs into a single task measure.

¹⁷ See also appendix Table A.4, which shows the breakdown between C_j^{math} and C_j^{abstract} for the occupations from Table 1. The table suggests that, while all of the most math-intensive occupations are also high on abstraction, there is notable variation on C_j^{abstract} among the most language-intensive occupations. That is, among the language-intensive occupations, there are occupations that score either low (e.g. healthcare assistant) or high (e.g. customer dialogue specialist) in terms of abstract job content. Quite evidently, this mainly reflects the fact that these occupations differ in the extent to which they have direct (e.g. physical contact with patients) or only indirect (e.g. contact with customers by phone or email) contact with other people. Appendix B presents additional evidence and results on this specific issue.

¹⁸ An additional validation exercise based on the subcategories for the cognitive requirements in both mathematics and in native language is discussed in Appendix B.2.

¹⁹ These two-year apprenticeships make up less than 10% of all apprenticeships contracts and they are more often chosen by male than by female apprentices (SERI, 2022). See also appendix Fig. A.1.

Table 4
Occupational-level regressions.

	F_j							
C_j^{things}	-0.202*** (0.014)	-0.202*** (0.015)	-0.232*** (0.021)	-0.238*** (0.022)	-0.188*** (0.018)	-0.188*** (0.019)	-0.228*** (0.013)	-0.220*** (0.016)
C_j^{demand}		0.050*** (0.014)		0.057*** (0.017)		0.048*** (0.014)		0.045** (0.019)
Estimation Method	OLS	OLS	Frac. Probit	Frac. Probit	OLS	OLS	WLS	WLS
Number of observations	130	130	130	130	153	153	130	130
R-Squared	0.477	0.524	0.414	0.462	0.386	0.430	0.771	0.799
Adjusted R-Squared	0.473	0.517	-	-	0.382	0.423	0.769	0.795
Pseudo R-Squared	-	-	0.216	0.239	-	-	-	-

Notes: *, **, *** denotes statistical significance on the 10%, 5%, and 1% level, respectively. Robust standard errors are given in parentheses. For easier comparison we report marginal effects, evaluated at mean values, instead of parameter coefficients in columns 3 and 4. Weights are equal to the number of apprenticeship contracts in the canton of Bern as of August 2014 in a given occupation.

should be predictive of whether a given occupation is set up as a two-year rather than as a regular apprenticeship, and C_j^{things} should have a lower, if any, predictive value.

The estimates from Table 3 show, as expected, that the overall level of cognitive requirements is indeed highly predictive of an occupation being targeted towards academically weaker adolescents (the R-squared associated with the regression from column 1 equals 0.492). In contrast, also including C_j^{things} as regressor adds comparatively little fit to the regression, even though the point estimate is also statistically significant, as shown in column 2 of Table 3 (with an associated R-Squared of 0.573).

4.3. Analysis of the occupational-level data

In the second part of our analysis, we show that men and women chose occupations that differ strongly in their broad task content along the things-versus-people dimension.

Describing occupational gender segregation

In a preliminary step, however, we first show that there indeed is strong occupational gender segregation within the Swiss apprenticeship system (cf. Aepli et al., 2019; Kuhn and Wolter, 2022). As we mentioned above, access to data on the population of apprenticeship contracts in the canton of Bern allows us to calculate the proportion of female apprentices within each of the 214 different occupations covered in the data with reasonable accuracy. We do, however, exclude occupations with very low overall frequencies, i.e. occupations with less than ten contracts as of August 2014.²⁰ This restriction leads to the exclusion of 44 (of a total of 214) occupations. However, because these occupations are rarely chosen, they represent only 0.32% of the total of all apprenticeship contracts. Another restriction is due to the fact that we do not have information on the cognitive requirements for all of the occupations; due to this restriction, we lose another 40 occupations. We thus end up with a total of 130 different occupations for which we have at least ten apprenticeship contracts as well as information on the cognitive requirements in that occupation (unless stated otherwise, all results in this section are based on this set of occupations). These 130 occupations cover about 90% of the total of apprenticeship contracts in the canton of Bern as of August 2014.

Fig. 1 shows a histogram of the proportion of females by occupation. The most obvious feature is the bimodality of the distribution, with many occupations clustering at the two extreme values (specifically, there are many occupations that are primarily chosen by men; obviously, the least frequent occupations are those with a more or less balanced sex ratio).

This is more explicitly evident from Fig. 2, which plots the number of individual apprenticeship contracts by males (females) in an occupation characterized by a given overall proportion of females choosing an occupation. Obviously, and as expected, both male and female adolescents tend to cluster in occupations mainly chosen by individuals of the same gender (cf. Kuhn and Wolter, 2022). On average, the typical occupation chosen by female adolescents is 71.4% female, while the typical occupation of male adolescents is 77.5% male. Thus, without any doubt, the data show that there is strong occupational gender segregation in the Swiss apprenticeship system (Aepli et al., 2019, show similar results for the country as a whole).

Men and women choose occupations with different task content

In a next step, we merge the variables describing the task content of occupations with the information on the proportion of females within each of these occupations. This allows us to see whether occupations females predominantly choose differ in their task content from those occupations males primarily select. Based on the existing empirical literature on the subject,

²⁰ Architectural model builder (“Architekturmodellbauer”) or glass painter (“Glasmaler”) are just two examples of rarely chosen occupations. The occupational-level results are robust to the inclusion of the smaller occupations, however, as shown in columns 5 and 6 of Table 4.






Fig. 1. Proportion of females, by occupation. Notes: The figure shows the absolute number of learnable occupations with a given proportion of females choosing that occupation. Only occupations with ten or more apprenticeship contracts are included in the figure ($J = 130$; see Sections 3 and 4.3 in the main text for details). There are 12 (3) occupations with a female share of exactly 0 (1).



(a) Male apprentices



(b) Female apprentices

Fig. 2. Gender segregation in occupational choices. Notes: The figure shows the distribution of individuals across occupations characterized by varying proportions of females (shown on the x-axis of each panel). Panel (a) shows the distribution of male apprentices; panel (b) the distribution of female apprentices. The two figures are based on the population of apprenticeship contracts in the canton of Bern as of 2015. See also Sections 3 and 4.3 in the main text for details.




Fig. 3. Proportion of females versus broad task-content, by occupation. Notes: The figure plots, by occupation, the proportion of females against the task content along the people-versus-things dimension (positive values indicate that an occupation is more oriented towards things, negative values indicate that an occupation is oriented towards people). Only occupations with at least ten apprenticeship contracts are considered ($J = 130$). The size of the markers is proportional to the number of apprenticeship contracts in an occupation in the canton of Bern as of August 2014. See main text for additional details.

we expect to find a negative association between the degree to which an occupation is heavy on things-related content and the proportion of women in that occupation.

Fig. 3 thus shows a scatterplot of the fraction of females on the x-axis against C_j^{things} , the principal component describing the task content of an occupation j along the people-versus-things dimension on the y-axis (the dotted horizontal line corresponds to the overall proportion of women in the population of all apprenticeship contracts). The relation between the two variables could hardly be more evident. As expected, occupations mainly chosen by males are oriented towards things, while occupations predominantly chosen by females are tilted towards working with people (see also appendix Fig. A.3).²¹

In a next step, we run several linear regression models of the following form:

$$F_j = \pi_0 + \pi_1 C_j^{\text{things}} + \pi_2 C_j^{\text{demand}} + \epsilon_j, \tag{1}$$

where the dependent variable, F_j , denotes the proportion of women choosing occupation j , and thus $F_j \in [0, 1]$. The regressor of main interest is C_j^{things} , which describes whether occupation j is oriented towards working with things (as discussed in Section 4.1 above), and we include the overall demand level C_j^{demand} as an additional regressor in some of the specifications. We report heteroscedasticity-robust standard errors throughout.

Table 4 shows the resulting parameter estimates. In the first column, we include C_j^{things} as the only regressor. This specification yields an estimate of $\hat{\pi}_1 = -0.202$ with a robust standard error of about 0.014. The point estimate of π_1 implies that the predicted proportion of women in an occupation essentially shifts from zero to one as we move from the lowest to the highest values on C_j^{things} , consistent with the pattern shown in Fig. 3. Also note the large R-squared of 0.477 associated with this simple regression – also suggesting that the variable describing the task content along the things-versus-people dimension is likely one of the more important predictors of whether an occupation is predominantly chosen by men or women. In the second column, we add the overall demand level C_j^{demand} as regressor. By construction, this does not change the point estimate of π_1 – remember that the two variables are statistically independent from each other at the occupational level by construction – but it shows that the predictive value of the overall level of cognitive requirements is far lower than that of things-versus-people dimension. In substantive terms, it shows that women tend to favor somewhat more demanding occupations than men.²² In the next two columns, we estimate the same two specifications using a fractional probit regression (e.g. Wooldridge, 2010), to take into account that the dependent variable can only take on values within the unit interval (accordingly, we report marginal effects, evaluated at mean values, in these two columns). This yields comparable findings, both in terms of marginal effects and model fit. Next, we also include the smaller occupations in columns 5 and 6 (i.e. those occupations with only ten or less observations in the data on the individual-level apprenticeship contracts; as discussed in Section 3.2 above), again yielding a very similar point estimate ($\hat{\pi}_1 = 0.188$ in both specifications). As expected, however, the R-squared associated with these specifications is lower than in the baseline specifications because the fraction of female apprentices tends to be more variable in rarely chosen occupations. Finally, we re-estimate the two

²¹ A further notable feature of Fig. 3 is the fact that two of the most popular occupations among both males and females are located somewhere in the middle along the things-versus-people axis: “commercial employee” and “retail professional” (cf. appendix Table A.1).

²² This is well in line with existing empirical evidence on gender differences in academic achievement, showing that girls tend to (slightly) outperform boys academically (e.g. Voyer and Voyer, 2014).

specifications using weighted least squares in the last two columns, with weights that are proportional to the number of apprenticeship contracts in a given occupation. The resulting point estimates ($\hat{\pi}_1 = -0.228$ and $\hat{\pi}_1 = -0.220$, respectively) are again very similar to our baseline estimates. The fact that the R-Squared is considerably higher than in the unweighted regressions (i.e. 0.771 and 0.799 versus 0.477 and 0.524) is presumably due to the fact that the more popular occupations tend to be close to the estimated regression function (cf. Fig. 3 above).

Overall, the estimates reported in Table 4 are not only in line with our expectations (and also with common sense, we would add), they also turn out to be surprisingly large, by any statistical measure. This is even more remarkable considering that our main task measure represents only a fraction (about 36%) of the total variation in the underlying four variables (cf. appendix Table A.2). This in turn suggests that vocational interests along the things–people dimension presumably ranks amongst the most important proximate predictors of occupational gender segregation, and we will provide some more direct comparisons in Section 4.4 below.

4.4. Analysis of individual-level data

In the final part of our empirical analysis, we will replicate this result in our individual-level data and then show that the gender difference in occupational preferences is robust to the inclusion of a variety of control variables. We therefore combine the information describing occupations' task content with the individual-level survey among the adolescents from the German language area of the canton of Bern. This will allow us to go beyond the analysis at the occupational level, checking whether the finding from Section 4.3 is robust to the inclusion of additional predictors of occupational aspirations or choices, such as school grades or parental characteristics.

Gender differences in occupational aspirations

We therefore next estimate linear regressions that take the following form:

$$c_i^{\text{things}} = \alpha + \beta f_i + \gamma_x x_i + \gamma_p p_i + \delta_p (f_i \times p_i) + \gamma_r r_i + \delta_r (f_i \times r_i) + \epsilon_i, \quad (2)$$

where the dependent variable c_i^{things} characterizes the job content of the occupation for which individual i aspires while in 8th grade (i.e. c_i^{things} corresponds to the value of C_j^{things} associated with the occupation that individual i aspires for).²³

The main regressor in this case is a dummy variable indicating whether individual i is a female, in which case $f_i = 1$. Step by step, we will then also control for various individual-level controls x_i , parental-level controls p_i , as well as regional-level controls r_i . Note that controls at both the parental and the regional level are expected to be uncorrelated with an adolescent's gender because parents have no control over the sex of their children, and simply including these variables as additional controls will therefore not have any noticeable impact on the estimate of β . For that reason, we also include the interaction terms between the female dummy and these controls and then check whether or not we can reject the null hypothesis that the parameters associated with the interaction terms are simultaneously equal to zero (we will show robust F-statistics and associated p-values). Throughout, we are primarily interested in the estimated size of gender differences in occupational preferences, which is given by parameter β in equation (2), as well as in the coefficients associated with the interaction terms, i.e. δ_p and δ_r . We report robust standard errors and test statistics for the full set of estimates.

At the individual level, we include an adolescent's exact age, an indicator for being a single child, the number of siblings, school grades (in mathematics as well as in German, French, and English), an indicator of the educational track at the upper-secondary level, as well as survey measures of competitiveness, risk preferences, and locus of control as additional regressors (see Buser et al., 2017; Jaik and Wolter, 2019, for additional details). Parental-level controls include the highest educational attainment of both mother and father, as well as their occupations (major, one-digit ISCO group). Finally, in the full-blown specification, we include a full set of local labor markets dummies to further control for regional differences in the availability of apprenticeship positions in different occupations as well as for spatial differences in gender norms.

Table 5 presents the resulting OLS estimates. In the first column, we simply regress the things-orientation of individual's occupational aspirations on the female dummy. This yields an estimate of $\hat{\beta} = -1.775$, which implies that male and female adolescents in the sample differ widely in terms of the task content of their occupational aspirations. More specifically, this estimate implies that the mean difference in occupational preferences along the things-versus-people dimension between male and female adolescents equals more than one standard deviation, a result which is consistent with previous empirical evidence (e.g. Lippa, 2010; Su et al., 2009).²⁴ With a robust standard error of about 0.062, the estimate is also statistically

²³ More precisely, in the survey, individuals stated their occupational aspirations (with zero, one, or more than one possible occupations given). In a first step, and where possible, we assigned the same occupational coding to the raw text as those used for the occupational-level data. We then merged the corresponding value on C_j^{things} for each occupation with a valid code. In cases where the same individual named more than one occupation, we simply use the mean value of C_j^{things} across all occupations that individual in question named. We lose some observations because of missing answers ("don't know") or because we are not able to assign a learnable occupation to the answer (e.g. in cases where the adolescent stated an occupation which requires studying at the university, such as a medical doctor for example).

²⁴ In the occupational level data, the standard deviation of C_j^{things} equals about 1.234. In the individual-level survey data, the standard deviation of $c_{j|i}^{\text{things}}$ amounts to about 1.387. Both Su et al. (2009) and Lippa (2010) report similar-sized gender differences along the things-versus-people dimension of job content.

Table 5
Individual-level regressions, occupational aspirations at the start of 8th grade.

Female _i (yes = 1)	$c_{j i}^{\text{things}}$			
	-1.775*** (0.062)	-1.721*** (0.065)	-1.556*** (0.242)	-1.792*** (0.267)
<i>Individual-level controls:</i>				
Demographics	No	Yes	Yes	Yes
School grades	No	Yes	Yes	Yes
Preferences	No	Yes	Yes	Yes
<i>Parental-level controls:</i>				
Education	No	No	Yes	Yes
Occupation	No	No	Yes	Yes
<i>Regional-level controls:</i>				
Local labor market	No	No	No	Yes
<i>Main effects:</i>				
Parental education			0.477 (0.929)	0.492 (0.921)
Parental occupation			2.737 (0.000)	2.280 (0.003)
Local labor market				1.325 (0.198)
<i>Interaction effects:</i>				
Female × parental education			1.236 (0.253)	1.350 (0.184)
Female × parental occupation			1.305 (0.186)	1.310 (0.183)
Female × local labor market				0.900 (0.546)
Number of observations	1,191	1,191	1,191	1,191
R-Squared	0.409	0.423	0.455	0.473
Adjusted R-Squared	0.409	0.416	0.421	0.428

Notes: *, **, *** denotes statistical significance on the 10%, 5%, and 1% level, respectively. Robust standard errors are given in parentheses. The table also reports robust F-statistics (with associated p-values in parentheses below) from testing the null hypothesis that the corresponding interaction terms are equal to zero. Full regression results are available upon request. The full list of controls is as follows: “demographics” includes the age at the time of the survey, the number of siblings, and a dummy for being a single child; “school grades” includes grades in mathematics, German, French and English; “preferences” includes survey measures of competitiveness, risk preferences, and locus of control; parental education includes a full set of dummies for the highest attained education of both mother and father; parental occupation includes a full set of dummies for the occupation of both mother and father (major ISCO group); local labor market includes a full set of local-labor markets dummies.

highly significant. Moreover, and similar to the occupational-level regressions, the resulting R-squared of 0.409 is unusually large. Overall, this first specification based on individual-level data is thus consistent with our finding based on the occupational-level data from Section 4.3 above.

We next check whether this difference still holds when we add various control variables. We first add individual-level controls in the second column, which results in an estimate of $\hat{\beta} = -1.721$ (with a robust standard error of 0.065, the estimate also remains highly statistically significant). Obviously, adding these controls has hardly any impact on the size of $\hat{\beta}$, either because there are no or only small gender differences in these variables and/or because these variables do not predict occupational preferences along the things-versus-people dimension. In the next column, we further add parental-level controls as well as the interaction terms between these and the female dummy. Here we are mainly interested in the coefficients associated with the interaction terms.²⁵ The F-test associated with the null hypothesis that the gender difference in occupational preferences does not vary with parental education is small and statistically insignificant ($F = 1.236$, with a p -value of 0.253); similarly for parental occupation ($F = 1.350$, $p = 0.184$). In column 4, we further add dummies representing local labor markets and their interactions with the female dummy. This also yields a small and insignificant test statistic ($F = 0.900$, $p = 0.546$).

Taken together, and consistent with the results from our occupational-level analysis above, the estimates from Table 5 show that there is a large and statistically significant gender difference over the task content along the things-versus-people dimension of individual occupational aspirations. Moreover, we also find that this difference is not driven by any of the individual- or parental-level control variables considered in the regression analysis.

²⁵ As expected, however, and consistent with the results of Kuhn and Wolter (2022), the F-tests associated with the main effects of the parental controls tend to be statistically significant. Thus, children from different parental backgrounds do differ in their occupational preferences along the things-people dimension, but there appears to be no differential pattern between girls and boys in this regard.

Table 6

Individual-level regressions, overall demand-level and average recommended apprentice wage (associated with occupational aspirations at the start of 8th grade).

	$c_{j it}^{\text{demand}}$		$\ln(\bar{w}_{j it})$	
Female _{<i>i</i>} (yes = 1)	0.310*** (0.063)	-0.114 (0.258)	-0.042*** (0.010)	-0.059 (0.045)
<i>Individual-level controls:</i>				
Demographics	No	Yes	No	Yes
School grades	No	Yes	No	Yes
Preferences	No	Yes	No	Yes
<i>Parental-level controls:</i>				
Education	No	Yes	No	Yes
Occupation	No	Yes	No	Yes
<i>Regional-level controls:</i>				
Local labor market	No	Yes	No	Yes
<i>Main effects:</i>				
Parental education		0.658 (0.792)		0.520 (0.903)
Parental occupation		1.654 (0.050)		1.645 (0.052)
Local labor market		2.334 (0.006)		2.010 (0.021)
<i>Interaction effects:</i>				
Female × parental education		0.881 (0.566)		0.494 (0.919)
Female × parental occupation		1.118 (0.332)		3.024 (0.000)
Female × local labor market		1.436 (0.143)		1.613 (0.082)
Number of observations	1,191	1,191	1,129	1,129
R-Squared	0.020	0.178	0.015	0.117
Adjusted R-Squared	0.019	0.108	0.014	0.037

Notes: *, **, *** denotes statistical significance on the 10%, 5%, and 1% level, respectively. Robust standard errors are given in parentheses. The table also reports robust F-statistics (with associated p-values in parentheses below) from testing the null hypothesis that the corresponding interaction terms are equal to zero. Full regression results are available upon request. The full list of controls is as follows: “demographics” includes the age at the time of the survey, the number of siblings, and a dummy for being a single child; “school grades” includes grades in mathematics, German, French and English; “preferences” includes survey measures of competitiveness, risk preferences, and locus of control; parental education includes a full set of dummies for the highest attained education of both mother and father; parental occupation includes a full set of dummies for the occupation of both mother and father (major ISCO group); local labor market includes a full set of local-labor markets dummies.

Gender differences in the overall demand level and apprentice wages

In an additional step, we present analogous estimates for two other key characteristics of learnable occupations, namely their overall cognitive demand level as well as the apprentice wage which goes along with a given occupational choice (to save space, we only show estimates without any additional controls and with the full set of controls, respectively).

The first two columns of [Table 6](#) show estimates using the overall cognitive demand level associated with individuals' occupational aspirations at the start of 8th grade as the dependent variable. Consistent with [Table 4](#) above, we find a much smaller gender difference in the demand level of individuals' occupational aspirations. Thus even though the demand level represents a much larger proportion of the variation in the underlying (cf. appendix [Table A.2](#)), it is occupations' task content that is much closer associated with adolescents' gender. Moreover, the difference becomes statistically insignificant once we include the (full set of) control variables, but note that this is at least partially driven by a large increase in the associated standard error. In the remaining two columns, the dependent variable is the natural logarithm of the average apprentice wage that is associated with adolescents' occupational aspirations.²⁶ As shown in the third column, there is again a significant, but comparatively small gender difference of about 4 percentage points. This shows that a gender gap in wages already appears at this early stage of individuals' labor market career which is associated with differences in occupational aspirations (obviously, recommended wages do not vary by gender, which would be unconstitutional; moreover, appren-

²⁶ We use wage recommendations for the corresponding occupation provided by the responsible professional association, averaged across the different years of training. The data are taken from an online platform providing information about open apprenticeship positions in Switzerland (www.yousty.ch). While employers are not required to comply with the wage recommendations, actual apprentice wages tend to be close to the corresponding recommendation. However, actual apprentice wages within the same occupation often vary across regions because of differences in the overall wage level.

Table 7
Individual-level regressions, occupational choices at the end of 9th grade.

	$c_{j l}^{things}$			
Female _i (yes = 1)	-1.691*** (0.068)	-1.664*** (0.070)	-2.182*** (0.480)	-2.581*** (0.624)
<i>Individual-level controls:</i>				
Demographics	No	Yes	Yes	Yes
School grades	No	Yes	Yes	Yes
Preferences	No	Yes	Yes	Yes
<i>Parental-level controls:</i>				
Education	No	No	Yes	Yes
Occupation	No	No	Yes	Yes
<i>Regional-level controls:</i>				
Local labor market	No	No	No	Yes
<i>Main effects:</i>				
Parental education			1.974 (0.024)	2.102 (0.015)
Parental occupation			2.921 (0.000)	3.366 (0.000)
Local labor market				1.681 (0.066)
<i>Interaction effects:</i>				
Female × parental education			2.437 (0.004)	2.494 (0.003)
Female × parental occupation			4.693 (0.000)	4.232 (0.000)
Female × local labor market				0.789 (0.662)
Number of observations	953	953	953	953
R-Squared	0.397	0.411	0.465	0.479
Adjusted R-Squared	0.396	0.403	0.423	0.423

Notes: *, **, *** denotes statistical significance on the 10%, 5%, and 1% level, respectively. Robust standard errors are given in parentheses. The table also reports robust F-statistics (with associated p-values in parentheses below) from testing the null hypothesis that the corresponding interaction terms are equal to zero. Full regression results are available upon request. The full list of controls is as follows: “demographics” includes the age at the time of the survey, the number of siblings, and a dummy for being a single child; “school grades” includes grades in mathematics, German, French and English; “preferences” includes survey measures of competitiveness, risk preferences, and locus of control; parental education includes a full set of dummies for the highest attained education of both mother and father; parental occupation includes a full set of dummies for the occupation of both mother and father (major ISCO group); local labor market includes a full set of local-labor markets dummies.

ticeship training is set up as a full-time undertaking, independent of the occupation chosen, rendering such considerations irrelevant at this stage). Moreover, the wage gap at this stage is substantively smaller than the wage gap among fully trained workers, which clearly points to the fact that additional factors become relevant later on, when apprentices have finished their training, enter the regular labor market and also engage in family formation. Finally, the fourth column shows that this gender difference also becomes insignificant once the full set of controls is taken into account. In this case it is even more obvious that this is mainly driven by the inflation of the associated standard error (which increases from 0.01 to 0.045).

Gender differences in actual occupational choices

In the final step of our analysis, we rerun the analysis above, but focus on the job content along the things-versus-people dimension of actual occupational choices instead of occupational aspirations.²⁷ By and large, these additional estimates mirror those associated with individuals’ occupational aspirations. There is, however, a small and statistically insignificant difference in the average value in the “things”-intensity of occupational aspirations ($\bar{c} = -0.433$) and actual occupational choices ($\bar{c} = -0.361$).

Again, the first specification of Table 7 simply regresses the things-orientation of an adolescent’s actual occupational choice on the female dummy. This yields an estimate of $\beta = -1.691$ with an standard error of 0.068; somewhat smaller, though not significantly so, than the corresponding estimate from Table 5. Thus, in the unconditional case we find a similar sized gender difference along the things-versus-people dimension for occupational aspirations as well as occupational choices. As evident from column 2, adding individual-level controls has only a small impact on the estimate of β , which becomes somewhat smaller than in the unconditional case ($\beta = -1.664$, robust standard error of 0.070). Again, this parallels the corresponding result related to occupational aspirations from Table 5.

²⁷ Focusing on actual choices implies that we must focus on a smaller, and potentially somewhat selective, sample. Specifically, both individuals who decided to go on with general education (i.e. baccalaureate school) and individuals who have not yet found an apprenticeship position, or who opted for an interim solution, are not included in this subsample.

Results differ somewhat between occupational choices and aspirations when we further add parental-level controls and the corresponding interaction terms with the female dummy, as done in column 3 of Table 7. In the case of actual occupational choices, both the F-test associated with the interaction terms between the female dummy and parents' education ($F = 2.437$, $p = 0.024$) and with the interaction terms between the female dummy and parental occupation ($F = 2.921$, $p < 0.001$) turns out to be statistically significant, suggesting that parental background does have some differential impact on the things-intensity of actual occupational choices between girls and boys, a finding that is also in line with existing empirical evidence (e.g. Humlum et al., 2019).

In general, however, the impact of parental background remains very limited, as the gender difference in occupational choices along the things-people dimension shows up across most attributes characterizing parental background. Thus one main conclusion from the analysis using individual-level data is that the gender difference in vocational interests is not driven by any of the controls at the individual or at the parental level, and that it already shows up in adolescents' aspirations.

A comparison of the implied effect sizes

To ease the comparison across the different estimates from tables 5 to 7, we also present a short description of the effect sizes which are implied by these estimates. First, our estimate based on the individual-level data from column 1 of Table 5 implies an effect size (i.e. a standardized mean difference between male and female adolescents) with regard to the task content in occupational aspirations of about -1.28, while the corresponding estimate from Table 7 implies an almost identical effect size of about -1.261 for the task content of actual occupational choices. Moreover, this corresponds surprisingly well with effect sizes related to vocational interests, and it is therefore also consistent with the intuitive notion that vocational interests and choices overlap to a considerable degree. In stark contrast, the estimates from Table 6 show that the effect sizes associated with either an occupation's demand level or the associated apprentice wage are much smaller. The comparatively small effect size associated with the demand level of occupational aspirations, about 0.281, is consistent with both the finding that girls slightly outperform boys academically and that the finding that there is no gender difference in mean intelligence, which further corroborates our general approach. With regard to recommended apprentice wages, the implied effect size equals about -0.248. Notably, this effect size is also much smaller than the effect size associated with an occupation's task content.

5. Conclusions

Occupational gender segregation remains at persistently high levels, for reasons not yet fully understood. In this paper, we add to this important discussion and argue that gender differences in vocational interests are among the most important proximate determinants of occupational gender segregation. More precisely, motivated by previous empirical evidence, mainly by psychologists, we use a unique combination of different data sources to show that male adolescents clearly tend to favor occupations which require creating and/or manipulating inanimate objects (i.e. "things"), while female adolescents prefer to work in occupations in which interacting with customers or patients is important (i.e. "people").

For our own empirical analysis, we use data on the cognitive requirements in 130 different learnable occupations in the Swiss apprenticeship system to classify them according to their broad task content along the things-versus-people dimension, using a data-driven process. Moreover, we validate our classification results using an alternative, fully independent source of data that describes the actual task-content of different occupations. We then show at the occupational level that this simple, unidimensional classification of the occupations explains a surprisingly large part of the observed variation in the proportion of females in an occupation (with an unusually large R-squared of 0.477). In other words, knowing whether an occupation is tilted towards working with things rather than people allows one to formulate a reasonable guess as to whether men rather than women predominantly choose this occupation. This finding also illustrates an issue that is not often acknowledged in the discussion about women's underrepresentation in the technical, i.e. things-oriented, occupations: this apparently goes hand-in-hand with women's overrepresentation in the social, people-oriented occupations, such as healthcare. Therefore, for example, policies explicitly aiming at increasing the proportion of women in the STEM occupations will, *ceteris-paribus*, necessarily decrease the number of females working in other occupations.

In the second part of our empirical analysis, we replicate this finding using individual-level data for a sample of adolescents from the German language area of the canton of Bern. Taken together, we find that variation in the task content of occupations along the things-versus-people is a very powerful predictor of whether males or females predominantly choose an occupation; actually, and consistent with the occupational-level results, it appears that this variable is likely one of the most important proximate predictors of occupational gender segregation – by any statistical measure applied. For example, a simple, univariate regression of the things-intensity of individuals' occupational aspirations (occupational choices) on a female dummy yields an estimate of -1.775 (-1.691), which approximately corresponds an effect size of about -1.28 (-1.26). Notably, these differences are considerably larger than gender differences in economic preferences, such as risk preferences or positive reciprocity.²⁸ It is also larger than the gender differences in personality measures, at least when judged on the

²⁸ For Switzerland, the data from Falk et al. (2018) yield an effect size of about 0.228 for risk preferences and an effect size of about 0.170 for positive reciprocity (effect sizes are much smaller for the other preference measures available in these data).

basis of individual personality traits.²⁹ As we mentioned above, however, this finding is well in line with previous empirical studies. For example, both [Su et al. \(2009\)](#) and [Lippa \(2010\)](#) report gender differences in vocational interests along the things-versus-people dimension that are as large as those reported in this paper. Moreover, in his overview, [Archer \(2019\)](#) reports differences of similar size in closely related concepts (i.e. systemizing and things orientation).

Finally, our findings also overlap with and complement a number of recent empirical studies which show that various characteristics of (the task content of) an occupation are important in explaining the segregation of men and women into different occupations (e.g. [Speer, 2017](#); [Delaney and Devereux, 2019](#); [Jiang, 2021](#)). Most importantly, perhaps, the results of our analysis align closely with the findings from [Breda and Napp \(2019\)](#), who show that girls tend to score relatively better in languages than in mathematics, and vice versa for boys, and from [Baker and Cornelson \(2018\)](#), who document gender differences in visuo-spatial and motor aptitudes (see also [Halpern, 2011](#); [Moè et al., 2018](#); [Archer, 2019](#), among others). Thus, and not really surprisingly, interests appear to considerably overlap with associated cognitive abilities (see also [Geary, 2021](#); [Cohen, 2004](#)). Persistent gender differences in relative abilities and vocational interests may also explain why explicit interventions targeted towards the reduction or elimination of these differences often turn out to be relatively ineffective (e.g. [Baker and Cornelson, 2019](#); [Li, 2018](#)). Overall, and based on our own results as well as on the results of related studies, it appears evident to us that gender differences in these traits are absolutely key for understanding why men and women still tend to prefer to work in distinct occupations (see also [Wang et al., 2013](#); [2015](#); [Haworth et al., 2010](#)).

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Additional tables and figures

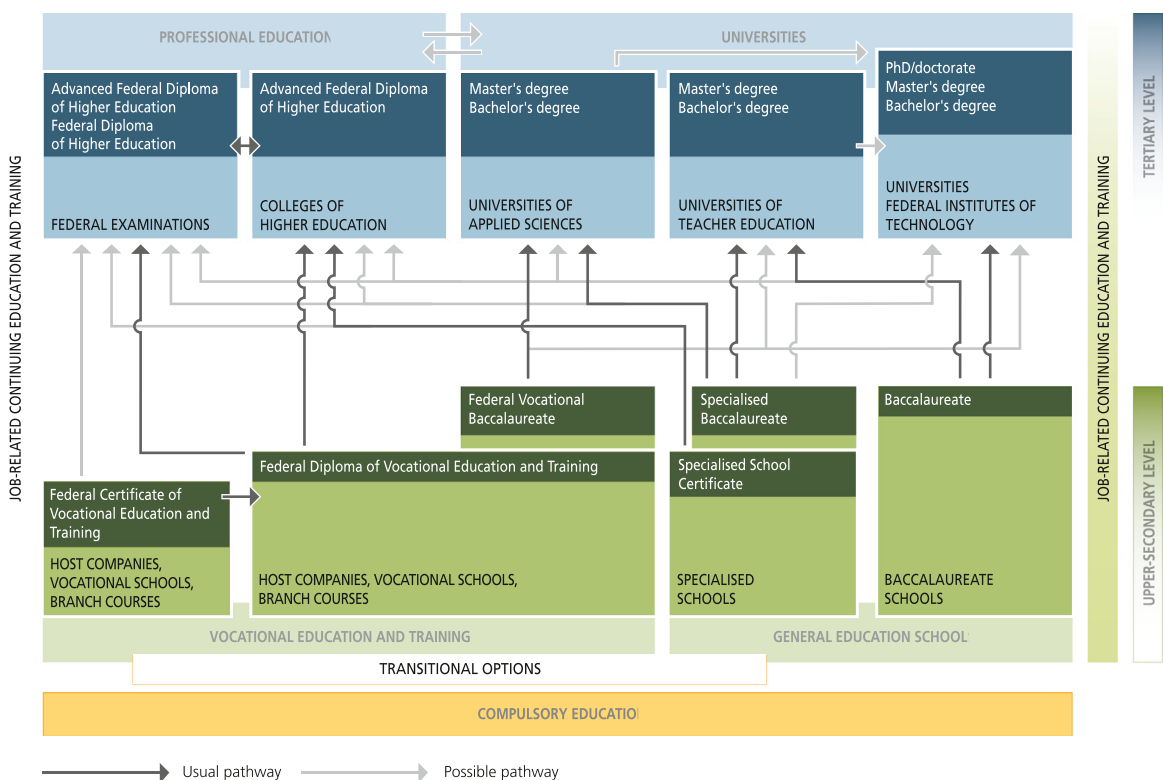


Fig. A.1. The Swiss educational system. Source: State Secretariat for Education, Research and Innovation (SERI).

²⁹ [Weisberg et al. \(2011\)](#) report an effect size of about -0.48 and -0.39 for the gender difference in agreeableness and in neuroticism, respectively (negative values indicate that women score higher on the traits than men; differences on the other traits are smaller). However, differences in personality traits are considerably larger when considered simultaneously, as argued by [Giudice et al. \(2012\)](#), who report an overall effect size of about 1.71 to 2.71.

Berufe im Vergleich





Fig. A.2. Two example profiles from the “Anforderungsprofile” website. Notes: The figure shows an example comparison in the cognitive requirements for two different occupations, healthcare assistant (“Fachmann/-frau Gesundheit”) and mechanical engineer (“Polymechniker/-in”), as available from the website (www.anforderungsprofile.ch). See also Table 1 in the main text.

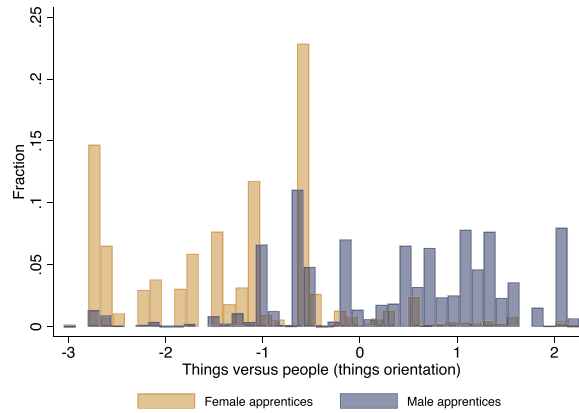


Fig. A.3. Proportion of apprenticeship contracts by an occupation’s task content along the things-people dimension. Notes: The figure shows the frequency distribution of C_j^{things} , i.e. the task content of an occupation along the things-people dimension, weighted by the number of female and male apprenticeship contracts, respectively.

Table A.1

The ten most popular apprenticeship occupations among male and female apprentices in the canton of Bern.

Rank	Occupation	f_j	F_j
<i>(a) Male apprentices</i>			
1.	Commercial employee (“Kaufmann”)	0.093	0.643
2.	Information technologist (“Informatiker”)	0.051	0.089
3.	Farmer (“Landwirt”)	0.042	0.156
4.	Mechanical engineer (“Polymechaniker”)	0.041	0.031
5.	Licensed electrician (“Elektroinstallateur”)	0.040	0.028
6.	Retail professional (“Detailhandelsfachmann”)	0.033	0.649
7.	Logistics expert (“Logistiker”)	0.033	0.102
8.	Carpenter (“Zimmermann”)	0.032	0.012
9.	Carpenter (“Schreiner”)	0.031	0.148
10.	Gardener (“Grtner”)	0.027	0.287
<i>(b) Female apprentices</i>			
1.	Commercial employee (“Kauffrau”)	0.213	0.643
2.	Healthcare assistant (“Fachfrau Gesundheit”)	0.138	0.902
3.	Retail professional (“Detailhandelsfachfrau”)	0.077	0.649
4.	Certified social care worker (“Fachfrau Betreuung”)	0.058	0.865
5.	Specialist in hotel housekeeping (“Hotelfachfrau”)	0.041	0.909
6.	Dental assistant (“Dentalassistentin”)	0.029	0.987
7.	Medical practice assistant (“Medizinische Praxisassistentin”)	0.027	0.996
8.	Hairdresser (“Coiffeuse”)	0.024	0.928
9.	Chef (“Kchin”)	0.024	0.438
10.	Pharmacy assistant (“Pharma-Assistentin”)	0.022	0.978

Notes: The table shows the ten most popular occupations among male and female apprentices in the canton of Bern as of August 2014 (the official German description of the occupation is given in parentheses, along with the English translation proposed by the State Secretariat for Education, Research and Innovation (SERI), where available). f_j denotes the proportion of male (female) apprentices choosing a specific occupation among all male (female) apprentices, while F_j denotes the proportion of female apprentices in the corresponding occupation.

Table A.2

Principal-components analysis of the occupational-level data describing the cognitive requirements in four different subjects.

Variable	PC-1	PC-2	PC-3	PC-4
Mathematics	0.3617	0.7195	0.4990	0.3202
Native language	0.5694	-0.3266	-0.3653	0.6601
Natural sciences	0.5661	0.3110	-0.4754	-0.5974
Foreign language	0.4739	-0.5281	0.6258	-0.3238
Proportion	0.616	0.287	0.075	0.023

Notes: The table shows the results from the principal-components analysis of the four main subjects covered by the data on cognitive requirements in learnable occupations. The upper part of the table shows the factor loadings for the four PCs, the lower part shows the proportion of variance explained by any of the four PCs.

Table A.3

Task content of occupations ranked highest/lowest along the overall level of cognitive requirements.

Rank	Occupation	C_j^{demand}
<i>(a) Ranked highest on overall cognitive demands</i>		
1.	Mediamatician (“Mediamatiker”)	3.600
2.	Optometrist (“Augenoptikerin”)	3.445
3.	Interactive media designer (“Interactive Media Designer”)	3.133
4.	Geomatics expert (“Geomatiker”)	3.028
5.	Information and documentation specialist (“Fachfrau Information und Dokumentation”)	2.977
6.	Druggist (“Drogistin”)	2.848
7.	Laboratory technician (“Laborantin”)	2.656
8.	Automation engineer (“Automatiker”)	2.476
9.	Bookseller (“Buchhndlerin”)	2.467
10.	Information technologist (“Informatiker”)	2.461
<i>(b) Ranked lowest on overall cognitive demands</i>		
1.	Mechanical assistant (“Mechanikpraktiker”)	-2.972
2.	Automobile assistant (“Automobil-Assistent”)	-2.793
3.	Timber worker (“Holzbearbeiter”)	-2.792
4.	Logistician (“Logistikerin”)	-2.722
5.	Dairy industry assistant (“Milchpraktiker”)	-2.573
6.	Tire work assistant (“Reifenpraktiker”)	-2.569
7.	Food production assistant (“Lebensmittelpraktiker”)	-2.564
8.	Horse keeper (“Pferdewartin”)	-2.548
9.	Metal construction practitioner (“Metallbaupraktiker”)	-2.542
10.	Building services assistant (“Haustechnikpraktiker”)	-2.463

Notes: The table shows the ten occupations ranked highest (lowest) in terms of their overall cognitive requirement level, C_j^{demand} . We show the female (male) description in German if an occupation is chosen by a majority of female (male) apprentices.

Table A.4

Disaggregated task content of occupations ranked highest/lowest along the things-versus-people dimension.

Rank	Occupation	C_j^{math}	C_j^{abstract}	C_j^{things}
<i>(a) The most things-oriented occupations</i>				
1.	Plant and apparatus engineer (“Anlagen- und Apparatebauer”)	1.713	0.547	2.260
2.	Design engineer (“Konstrukteur”)	1.434	0.650	2.084
3.	Mechanical engineer (“Polymechaniker”)	1.389	0.662	2.052
4.	Automatic technician (“Automatikmonteur”)	1.178	0.850	2.027
5.	Mechanical technician (“Produktionsmechaniker”)	1.178	0.850	2.027
6.	Gunsmith (“Bchsenmacher”)	1.380	0.594	1.974
7.	Insulation contractor (“Isolierspengler”)	1.698	0.235	1.933
8.	Automation engineer (Automatiker”)	1.056	0.828	1.884
9.	Carpenter (“Schreiner”)	1.515	0.023	1.538
10.	Industrial ceramist (“Industriekeramiker”)	0.599	0.918	1.517
<i>(b) The most people-oriented occupations</i>				
1.	Podiatrist (“Podologin”)	-1.293	-1.757	-3.049
2.	Healthcare assistant (“Fachfrau Gesundheit”)	-1.242	-1.474	-2.715
3.	Information and documentation specialist (“Fachfrau Information und Dokumentation”)	-2.784	0.150	-2.634
4.	Certified social care worker (“Fachfrau Betreuung”)	-1.060	-1.530	-2.590
5.	Druggist (“Drogistin”)	-1.582	-0.956	-2.538
6.	Pharmacy assistant (“Pharma-Assistentin”)	-1.504	-0.730	-2.234
7.	Hairdresser (“Coiffeuse”)	-1.660	-0.573	-2.233
8.	Bookseller (“Buchhndler”)	-2.403	0.185	-2.218
9.	Hairdresser (“Coiffeuse”)	-1.207	-0.929	-2.136
10.	Customer dialogue specialist (“Fachmann Kundendialog”)	-3.048	0.976	-2.072

Notes: The variable C_j^{things} is equal to the sum of C_j^{math} and C_j^{abstract} . Therefore, column 3 is identical to Table 1 in the main text. We show the female (male) description in German if an occupation is chosen by a majority of female (male) apprentices. “Hairdresser” appears twice in the table because there is both a two-year and a regular apprenticeship in that occupation.

Appendix B. Details concerning the occupational classification

In this appendix we present some additional robustness and validation checks concerning our classification of the task content of the various occupations, as discussed in Section 3.1 in the main text.




Fig. B.1. C_j^{abstract} , C_j^{math} and C_j^{things} , by occupation. Notes: Panels (a) and (b) plot our main task measure C_j^{things} against the two underlying principal components (i.e. mathiness and abstractness, C_j^{math} and C_j^{abstract} , respectively). Panel (c) plots these two variables directly against each other. Finally, panel (d) show the association between the proportion of female apprentices and C_j^{math} .

B1. Disaggregating our main task measure

One potential objection against our occupational classification is that we combine two distinct PCs into a single measure (and that one of these two PCs, i.e. C_j^{abstract} , is admittedly ambiguous in its interpretation). We thus first present some additional graphical evidence in support of this data-driven decision.

Panels (a) and (b) of Fig. B.1 plot C_j^{things} against C_j^{math} or against C_j^{abstract} , respectively, illustrating that both C_j^{math} and C_j^{abstract} are strongly correlated with our main task measure. Given that C_j^{things} equals the sum of C_j^{math} and C_j^{abstract} , these correlations hold by definition, however. More interestingly, therefore, and not immediately evident otherwise, the two figures also show that the strength of the association appears to vary along different values of either C_j^{math} or C_j^{abstract} . For example, there is apparently more variability in C_j^{things} for occupations with a weaker orientation towards mathematics. Panel (c) plots C_j^{abstract} against C_j^{math} , directly showing that there is more variation in C_j^{abstract} among occupations that are more heavily oriented towards languages – compared to occupations where mathematics and natural sciences are relatively more important. Finally, panel (d) plots the share of female apprentices against C_j^{math} , showing that the two variables are associated negatively with each other, i.e. the share of female apprentices tends to be lower in occupations oriented towards mathematics and natural sciences. Clearly, the two variables are also strongly correlated with each other, but the correlation is supposedly weaker than the one shown in Fig. 3 in the main text.

In an additional step, we show that regressions using the disaggregated measures of occupations’ task content yield similar, though somewhat muted, results (muted in the sense that the predictive value is lower). The first column of Table B.1

Table B.1
Occupational-level regressions, disaggregated task measure.

	F_j			
C_j^{math}	-0.210*** (0.019)	-0.207*** (0.017)	-0.209*** (0.021)	-0.206*** (0.018)
C_j^{abstract}		-0.187*** (0.028)		-0.186*** (0.028)
C_j^{demand}			0.050*** (0.017)	0.050*** (0.014)
Number of observations	130	130	130	130
R-Squared	0.396	0.477	0.444	0.525
Adjusted R-Squared	0.391	0.469	0.435	0.514

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

Table B.2
Validating our occupational classification using subcategories in cognitive requirements (mathematics only).

	Cognitive requirements in mathematics					
	Mean	Subcat-1 Algebra	Subcat-2 Geometry	Subcat-3 Measurement	Subcat-4 Calculus	Subcat-5 Statistics
C_j^{things}	12.137*** (0.188)	11.393*** (0.443)	17.391*** (0.602)	11.827*** (0.622)	13.652*** (0.405)	6.422*** (0.554)
C_j^{demand}	6.778*** (0.176)	7.052*** (0.333)	4.820*** (0.579)	6.447*** (0.552)	7.515*** (0.304)	8.057*** (0.329)
Number of observations	130	130	130	130	130	130
R-Squared	0.980	0.886	0.841	0.825	0.922	0.812
Adjusted R-Squared	0.980	0.884	0.838	0.823	0.921	0.809

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

uses C_j^{math} instead of C_j^{things} as the main regressor. The resulting point estimate of -0.21 is virtually the same as the corresponding estimate from Table 4 in the main text (i.e. -0.202). Consistent with Fig. B.1, however, there is a notable difference in the predictive value of the variable in explaining the share of female apprentices across all the occupations (i.e. the R-squared associated with this regression is considerably lower than when using C_j^{things} as regressor). Adding C_j^{abstract} as an additional regressor, we can see that both point estimates are statistically significant and, moreover, are of almost equal size (-0.207 and -0.187 , respectively). Indeed, we can not reject the null hypothesis that the two coefficients are statistically identical ($F = 0.36$, $p = 0.547$). This specification thus provides a data-based rationale for combining the two PCs into a single measure of task content.

Columns 3 and 4 of Table B.1 replicate the the first two specifications while also controlling for the overall demand level in an occupation. Again, the estimate from column 3 is close to the corresponding estimate from Table 4, but with lower predictive power. Also, in column 4, the estimates of the two task measures do not differ from each other statistically ($F = 0.34$, $p = 0.563$), paralleling the result from column 2.

B2. Additional validation checks based on the subcategories in cognitive requirements

In addition, the fact that the data describing the cognitive requirements in the different apprenticeship occupations contain subclassifications that we do not directly use in the construction of the two latent task variables (i.e. C_j^{demand} and C_j^{things}) allows us to implement further validation checks. Specifically, we next show that these two variables predict the subscores, both in mathematics and in first language, in a pattern that lines up with our occupational classification along the things–people dimension. Importantly, note that this finding is not simply a mechanical implication as we only use the aggregate scores to construct the latent variables (and thus the subscores only enter indirectly, through making up the aggregate score).

We start with the results for mathematics in Table B.2. The dependent variable in the table varies across columns, with the mean value in mathematics used in the first column, and with the different subcategories used in the remaining columns. In Table B.2 (as well as in and Table B.3 below), we use the original scaling of the dependent variables (which can take on values between 0 and 100). The first column of Table B.2 mainly serves as a benchmark for the other columns (i.e. we can compare whether the estimates for the subscores differ from these estimates). The interesting pattern emerging from Table B.2 is that the point estimates associated with C_j^{things} do indeed vary across the different columns. For each

Table B.3

Validating our occupational classification using subcategories in cognitive requirements (native language only).

	Cognitive requirements in native language					
	Mean	Subcat-1 Reading	Subcat-2 Listening	Subcat-3 Writing	Subcat-4 Conversation	Subcat-5 Speech
C_j^{things}	-4.758*** (0.187)	-0.567 (0.637)	-6.928*** (0.378)	-3.406*** (0.408)	-5.798*** (0.399)	-7.092*** (0.497)
C_j^{demand}	8.071*** (0.175)	8.361*** (0.406)	6.590*** (0.392)	8.560*** (0.326)	8.079*** (0.339)	8.764*** (0.315)
Number of observations	130	130	130	130	130	130
R-Squared	0.962	0.738	0.814	0.861	0.861	0.856
Adjusted R-Squared	0.962	0.734	0.811	0.859	0.859	0.854

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

mathematical subcategory, there is a positive coefficient on C_j^{things} , implying that things-oriented occupations tend to score higher on each mathematical subcategory; however, the effect varies substantively across columns, from a low of about 6.4 for “statistics” to a high of about 17.4 for “geometry”. In contrast, note that the coefficient associated with C_j^{demand} varies much less across the different columns.

Table B.3 presents analogous findings concerning native language (for example, German in the German-speaking part of the country). Again, we regress the various values on the subcategories on both C_j^{things} and C_j^{demand} . As for mathematics, the result from the first column is shown as a benchmark for the other columns. The five subcategories for the cognitive requirements in native language, such as “reading” or “writing”, are all negatively associated with C_j^{things} , but again to very different degrees. For example, “reading” is not significantly associated with C_j^{things} , while both “participation in discussions” and “coherent speech” are strongly associated with C_j^{things} . Again, this aligns well with the claim that things-intensive occupations are characterized by a lower level of immediate interaction with other people (e.g. customers or patients) than things-intensive occupations.

References

- Aeppli, M., Kuhn, A., Schweri, J., 2019. Frauen und Männer lernen andere Berufe. *EHB, skilled* 2/19 4–5.
- Archer, J., 2019. The reality and evolutionary significance of human psychological sex differences. *Biological Reviews* 94 (4), 1381–1415.
- Baker, M., Cornelson, K., 2018. Gender-based occupational segregation and sex differences in sensory, motor, and spatial aptitudes. *Demography* 55 (5), 1749–1775.
- Baker, M., Cornelson, K., 2019. Title IX and the spatial content of female employment – Out of the lab and into the labor market. *Labour Economics* 58, 128–144.
- Bertrand, M., 2011. New perspectives on gender. *Handbook of Labor Economics* 4b, 1543–1590.
- Blau, F.D., Kahn, L.M., 2017. The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature* 55 (3), 789–865.
- Bonin, H., Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2007. Cross-sectional earnings risk and occupational sorting: The role of risk attitudes. *Labour Economics* 14 (6), 926–937.
- Breda, T., Napp, C., 2019. Girls' comparative advantage in reading can largely explain the gender gap in math-related fields. *Proceedings of the National Academy of Sciences* 116 (31), 15435–15440.
- Buchmann, M., Buchs, H., Busch, F., Gnehm, A.-S., Hevenstone, D., Klarer, U., Müller, J., Müller, M., Sacchi, S., Salvisberg, A., vonOw, A., 2020. *Stellenmarkt-Monitor Schweiz 1950 – 2019* [dataset]. Distributed by FORS, Lausanne.
- Buser, T., Niederle, M., Oosterbeek, H., 2014. Gender, competitiveness, and career choices. *Quarterly Journal of Economics* 129 (3), 1409–1447.
- Buser, T., Peter, N., Wolter, S.C., 2017. Gender, competitiveness, and study choices in high school: Evidence from Switzerland. *American Economic Review* 107 (5), 125–130.
- Buser, T., Peter, N., Wolter, S.C., 2022. Gender, willingness to compete and career choices along the whole ability distribution. *Experimental Economics* forthcoming.
- Charles, M., 2017. Venus, mars, and math: Gender, societal affluence, and eighth graders' aspirations for STEM. *Socius* 3, 1–16.
- Charles, M., Bradley, K., 2009. Indulging our gendered selves? Sex segregation by field of study in 44 countries. *American Journal of Sociology* 114 (4), 924–976.
- Charles, M., Grusky, D.B., 2004. *Occupational ghettos: The worldwide segregation of women and men*. Stanford University Press, Stanford, CA.
- Charness, G., Gneezy, U., 2012. Strong evidence for gender differences in risk taking. *Journal of Economic Behavior & Organization* 83 (1), 50–58.
- Chernyshenko, O.S., Stark, S., Nye, C.D., 2019. Interest measurement. in c. d. nye and j. rounds, editors, vocational interests in the workplace: Rethinking behavior at work. *SIOP Organizational Frontiers Series* 80–96. Routledge
- Coenen, J., Borghans, L., Diris, R., 2021. Personality traits, preferences and educational choices: A focus on STEM. *Journal of Economic Psychology* 84, 102361.
- Baron-Cohen, S., 2004. *The essential difference*. Penguin UK.
- Cortes, P., Pan, J., 2018. Occupation and gender. In: Averett, L.M.A., Hoffman, S.D. (Eds.), *The Oxford handbook of women and the economy*. Oxford University Press, Oxford, UK. Chapter 18, pages 425–452.
- Crosen, R., Gneezy, U., 2009. Gender differences in preferences. *Journal of Economic Literature* 47 (2), 448–474.
- Del Giudice, M., 2019. Measuring sex differences and similarities. *Gender and sexuality development: Contemporary theory and research*.
- Delaney, J.M., Devereux, P.J., 2019. Understanding gender differences in STEM: Evidence from college applications. *Economics of Education Review* 72, 219–238.
- Dittrich, M., Leipold, K., 2014. Gender differences in time preferences. *Economics Letters* 122 (3), 413–415.

- Donnay, D., Morris, M., Schaubhut, N., Thompson, R., 2005. Strong interest inventory manual. Consulting Psychology Press., Mountain View, CA.
- Donnay, D.A., Borgen, F.H., 1996. Validity, structure, and content of the 1994 strong interest inventory. *Journal of Counseling Psychology* 43 (3), 275.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., Sunde, U., 2018. Global evidence on economic preferences. *Quarterly Journal of Economics* 133 (4), 1645–1692.
- Falk, A., Hermle, J., 2018. Relationship of gender differences in preferences to economic development and gender equality. *Science* 362 (6412), eaas9899.
- Fouarge, D., Kriechel, B., Dohmen, T., 2014. Occupational sorting of school graduates: The role of economic preferences. *Journal of Economic Behavior & Organization* 106, 335–351.
- Geary, D. C. (2021). *Male, female: The evolution of human sex differences*. American Psychological Association, 3rd edition.
- Gelblum, M., 2020. Preferences for job tasks and gender gaps in the labor market. Unpublished manuscript, Harvard University.
- Del-Giudice, M., Booth, T., Irwing, P., 2012. The distance between Mars and Venus: Measuring global sex differences in personality. *PLoS one* 7 (1), e29265.
- Gorsuch, R.L., 2014. Factor analysis. Taylor and Francis, 2nd edition..
- Graziano, W.G., Habashi, M.M., Woodcock, A., 2011. Exploring and measuring differences in person–thing orientations. *Personality and Individual Differences* 51 (1), 28–33.
- Greenberg, D.M., Warrier, V., Allison, C., Baron-Cohen, S., 2018. Testing the empathizing–systemizing theory of sex differences and the extreme male brain theory of autism in half a million people. *Proceedings of the National Academy of Sciences* 115 (48), 12152–12157.
- Groen, Y., Fuermaier, A., Tucha, L., Koerts, J., Tucha, O., 2018. How predictive are sex and empathizing–systemizing cognitive style for entry into the academic areas of social or physical sciences? *Cognitive Processing* 19 (1), 95–106.
- Halpern, D.F., 2011. Sex differences in cognitive abilities. Psychology Press, 4th edition..
- Haworth, C.M., Dale, P.S., Plomin, R., 2010. Sex differences in school science performance from middle childhood to early adolescence. *International Journal of Educational Research* 49 (2–3), 92–101.
- Holland, J.L., 1959. A theory of vocational choice. *Journal of Counseling Psychology* 6 (1), 35–45.
- Houston, J.M., Harris, P.B., Howansky, K., Houston, S.M., 2015. Winning at work: Trait competitiveness, personality types, and occupational interests. *Personality and Individual Differences* 76, 49–51.
- Humlum, M.K., Nandrup, A.B., Smith, N., 2019. Closing or reproducing the gender gap? parental transmission, social norms and education choice. *Journal of Population Economics* 32 (2), 455–500.
- McIntyre, M.M., Graziano, W.G., 2016. Seeing people, seeing things: Individual differences in selective attention. *Personality and Social Psychology Bulletin* 42 (9), 1258–1271.
- McIntyre, M.M., Graziano, W.G., 2019. A snapshot of person and thing orientations: How individual differences in interest manifest in everyday life. *Personality and Individual Differences* 136, 160–165.
- Jaik, K., Wolter, S.C., 2019. From dreams to reality: market forces and changes from occupational intention to occupational choice. *Journal of Education and Work* 32 (4), 320–334.
- James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. *An Introduction to Statistical Learning*. Springer.
- Jiang, X., 2021. Women in STEM: Ability, preference, and value. *Labour Economics* 70, 101991.
- John, K., Thomsen, S.L., 2014. Heterogeneous returns to personality: the role of occupational choice. *Empirical Economics* 47 (2), 553–592.
- Kahn, S., Ginther, D., 2017. Women and STEM. NBER Working Paper 23525..
- Kamas, L., Preston, A., 2015. Can social preferences explain gender differences in economic behavior? *Journal of Economic Behavior & Organization* 116, 525–539.
- Kleinjans, K.J., 2009. Do gender differences in preferences for competition matter for occupational expectations? *Journal of Economic Psychology* 30 (5), 701–710.
- Kuhn, A., Wolter, S.C., 2022. The strength of gender norms and gender-stereotypical occupational aspirations among adolescents. *Kyklos* forthcoming.
- McLee, B., Lawson, K.M., Hale, S.M., 2015. Longitudinal associations between gender-typed skills and interests and their links to occupational outcomes. *Journal of Vocational Behavior* 88, 121–130.
- Li, H.H., 2018. Do mentoring, information, and nudge reduce the gender gap in economics majors? *Economics of Education Review* 64, 165–183.
- Lippa, R., 1998. Gender-related individual differences and the structure of vocational interests: The importance of the people–things dimension. *Journal of Personality and Social Psychology* 74 (4), 996–1009.
- Lippa, R.A., 2010. Gender differences in personality and interests: When, where, and why? *Social and Personality Psychology Compass* 4 (11), 1098–1110.
- Lippa, R.A., Preston, K., Penner, J., 2014. Women's representation in 60 occupations from 1972 to 2010: More women in high-status jobs, few women in things-oriented jobs. *PLoS One* 9 (5), e95960.
- Lordan, S.G., Pischke, J., 2022. Does Rosie like riveting? Male and female occupational choices. *Economica* 89 (353), 110–130.
- Moè, A., Jansen, P., Pietsch, S., 2018. Childhood preference for spatial toys. gender differences and relationships with mental rotation in STEM and non-STEM students. *Learning and Individual Differences* 68, 108–115.
- Morris, M.L., 2016. Vocational interests in the united states: Sex, age, ethnicity, and year effects. *Journal of Counseling Psychology* 63 (5), 604.
- Nauta, M.M., 2010. The development, evolution, and status of Holland's theory of vocational personalities: Reflections and future directions for counseling psychology. *Journal of Counseling Psychology* 57 (1), 11.
- Niederle, M., Vesterlund, L., 2011. Gender and competition. *Annual Review of Economics* 3 (1), 601–630.
- Nye, C.D., Bhatia, S., Prasad, J.J., 2019. Vocational interests and work outcomes. in: c. d. nye and j. rounds, editors, *vocational interests in the workplace. rethinking behavior at work*. SIOP Organizational Frontiers Series, chapter 5 97–128. Routledge.
- Olivetti, C., Petrongolo, B., 2016. The evolution of gender gaps in industrialized countries. *Annual Review of Economics* 8, 405–434.
- Prediger, D.J., 1982. Dimensions underlying holland's hexagon: Missing link between interests and occupations? *Journal of Vocational Behavior* 21 (3), 259–287.
- SCCRE, 2018. Swiss education report 2018. Aarau: Swiss Coordination Centre for Research in Education (SCCRE)..
- Schmitt, D.P., Realo, A., Voracek, M., Allik, J., 2008. Why can't a man be more like a woman? Sex differences in big five personality traits across 55 cultures. *Journal of Personality and Social Psychology* 94 (1), 168–182.
- SERI, 2022. Vocational and professional education and training in Switzerland. Facts and figures 2022. Bern: State Secretariat for Education, Research and Innovation SERI..
- Speer, J.D., 2017. The gender gap in college major: Revisiting the role of pre-college factors. *Labour Economics* 44, 69–88.
- Stoet, G., Geary, D.C., 2015. Sex differences in academic achievement are not related to political, economic, or social equality. *Intelligence* 48, 137–151.
- Stoet, G., Geary, D.C., 2022. Sex differences in adolescents' occupational aspirations: Variations across time and place. *PLoS one* 17 (1), e0261438.
- Su, R., Rounds, J., 2015. All STEM fields are not created equal: People and things interests explain gender disparities across STEM fields. *Frontiers in Psychology* 6, 189.
- Su, R., Rounds, J., Armstrong, P.I., 2009. Men and things, women and people: a meta-analysis of sex differences in interests. *Psychological Bulletin* 135 (6), 859–884.
- Svedholm-Häkkinen, A.M., Lindeman, M., 2016. Testing the Empathizing-Systemizing theory in the general population: Occupations, vocational interests, grades, hobbies, friendship quality, social intelligence, and sex role identity. *Personality and Individual Differences* 90, 365–370.
- Taber, B.J., Hartung, P.J., Borges, N.J., 2011. Personality and values as predictors of medical specialty choice. *Journal of Vocational Behavior* 78 (2), 202–209.
- Thelwall, M., Bailey, C., Tobin, C., Bradshaw, N.A., 2019. Gender differences in research areas, methods and topics: Can people and thing orientations explain the results? *Journal of Informetrics* 13 (1), 149–169.
- Voyer, D., Voyer, S.D., 2014. Gender differences in scholastic achievement: A meta-analysis. *Psychological Bulletin* 140 (4), 1174–1204.

- Wang, M.-T., Degol, J., Ye, F., 2015. Math achievement is important, but task values are critical, too: examining the intellectual and motivational factors leading to gender disparities in stem careers. *Frontiers in Psychology* 6, 36.
- Wang, M.-T., Eccles, J.S., Kenny, S., 2013. Not lack of ability but more choice: Individual and gender differences in choice of careers in science, technology, engineering, and mathematics. *Psychological science* 24 (5), 770–775.
- Weisberg, Y.J., DeYoung, C.G., Hirsh, J.B., 2011. Gender differences in personality across the ten aspects of the big five. *Frontiers in Psychology* 2, 178.
- Wettstein, E., Schmid, E., Gonon, P., 2017. *Swiss Vocational and Professional Education and Training (VPET)*. Bern: hep Verlag.
- Stewart-Williams, S., Halsey, L.G., 2021. Men, women and STEM: Why the differences and what should be done? *European Journal of Personality* 35 (1), 3–39.
- Wooldridge, J.M., 2010. *Econometric analysis of cross section and panel data*. MIT press.
- Wright, D.B., Eaton, A.A., Skagerberg, E., 2015. Occupational segregation and psychological gender differences: How empathizing and systemizing help explain the distribution of men and women into (some) occupations. *Journal of research in personality* 54, 30–39.
- Yang, Y., Svedholm-Barth, J.M., 2015. Gender differences in STEM undergraduates' vocational interests: People–thing orientation and goal affordances. *Journal of Vocational Behavior* 91, 65–75.
- Zickar, M.J., Min, H., 2019. A history of vocational interest measurement. in c. d. nye and j. rounds, editors, *vocational interests in the workplace. rethinking behavior at work*. SIOP Organizational Frontiers Series, chapter 3 59–79. Routledge