Field Choice, Skill Specificity, and Labor Market Disruptions*

Valerie Smeets
Aarhus University & CEPR

Lin Tian
INSEAD & CEPR

Sharon Traiberman
NYU & NBER

March 19, 2025

Abstract

We study how college students' field of study choices shape the economy's skill distribution and its resilience to labor market disruptions. Our framework models forward-looking students whose field choice determines their comparative advantage across occupations, influencing their ability to adjust to labor market shocks. Using Danish administrative data, we account for unobserved heterogeneity in both student preferences and worker skills to address selection concerns in field and occupation choices. We find substantial variation in specificity, with graduates from generalized fields responding to negative income shocks nearly six times more than those from specialized fields. While students cluster into broad STEM and non-STEM groups, within-cluster elasticities are higher than previously estimated. Counterfactual analysis suggests that easing capacity constraints in university programs leads to substantial reallocation of students, with minimal effects on national output. Additionally, simulated shocks highlight field-switching as a critical mechanism for adjusting to labor market disruptions, substantially enhancing economic outcomes. Importantly, unobserved differences among students significantly drive these responses, with the most adaptable individuals experiencing the largest gains.

Keywords:

JEL: F14, F16, J24, I24

*Contact: vas@econ.au.dk, lin.tian@insead.edu, st1012@nyu.edu. We thank Pol Antras, Chris Conlon, Rafael Dix Carneiro, Esteban Rossi-Hansberg, Martin Rotemberg, Dan Trefler, Dan Waldinger, Matt Wiswall, Jonathan Vogel, and numerous seminar and conference participants for helpful comments and suggestions. We thank Daniel Rock for providing the data on AI technological shock. Access to the Danish register data used in this study has been provided by ECONAU, Aarhus University.

1 Introduction

Labor markets are undergoing profound transformations driven by globalization, immigration, and technological advancements, reshaping occupational structures and skill requirements. Economies' resilience to these disruptions critically depends on the adaptability of workers and their ability to respond effectively (Heckman et al., 1998; Adao et al., 2024). Extensive research has examined how incumbent workers adjust to economic shocks (Grigsby, 2021) and how new entrants decide on skill acquisition along the vertical dimension—i.e., whether to pursue higher education (Charles et al., 2018; Adda and Dustmann, 2023). However, economists have only recently begun to investigate the horizontal dimension: how students choose specific fields of study, and how these specialization decisions shape labor market outcomes and the overall distribution of skills (Deming and Noray, 2020; Arcidiacono et al., 2025).

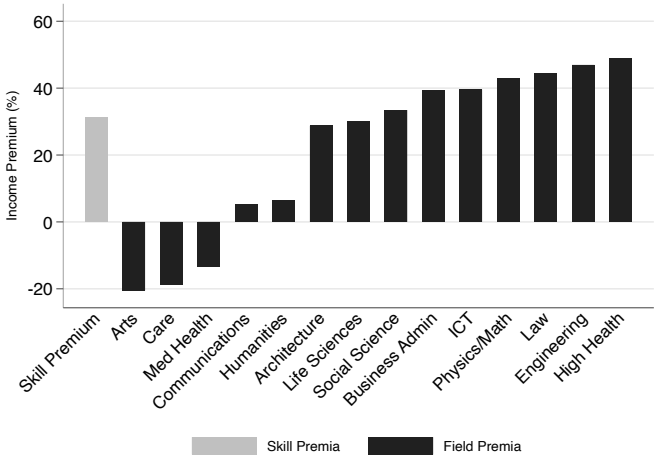
This horizontal dimension of educational choices is critical for policymaking and for fully understanding labor market dynamics, as the choice of field significantly influences graduates' employment trajectories, career adaptability, and ultimately, their resilience to economic shocks. Indeed, Figure 1 highlights substantial earnings variation across different fields, often exceeding the overall college wage premium.[1] Different fields equip workers with distinct skill sets, shaping their labor market outcomes (Altonji et al., 2012a), and the degree of specialization within these fields critically determines graduates' flexibility to adjust to disruptions (Kinsler and Pavan, 2015). Nevertheless, there remains little work that models the joint determination of field and career decisions.

To address this gap, we develop a comprehensive framework that jointly models educational decisions of new entrants and occupational reallocation of incumbent workers. Our model captures how individual preferences and abilities shape field-specific career paths and comparative advantages, highlighting two distinct margins of adjustment to shocks: the reallocation of incumbent workers across occupations and the adaptation of new entrants through changes in field selection. We use detailed Danish administrative data that links higher education application records with long-panel career histories of the entire Danish workforce from 1996 to 2019. Denmark's centralized admission process—which assigns students to programs via a deferred acceptance algorithm based on ranked priority lists—provides a unique measure of unobserved student preferences. This rich dataset enables us to quantify the responsiveness of different fields of study to labor market shocks, account

[1] Stark differences in earnings across fields have also been found in the United States (Bleemer and Mehta, 2022a) and Chile (Hastings et al., 2013).

for heterogeneity in student and worker characteristics, and perform counterfactual analyses on how policy changes and labor demand shocks affect overall human capital supply and long-term labor market outcomes.

Figure 1
College Premia and Field Premia



Notes: Black bars are field premia for individuals with college education, calculated as the raw, unadjusted, average earnings across all employed individuals whose college's field of study is f relative to those employed with base field, education, $\frac{\overline{Earn}_f}{\overline{Earn}_{Educ}}$. Skill premium is calculated as the average earnings of all workers with a college degree relative to workers without. In all regressions, workers are tracked in their first job after completing the degree of their highest level of education. All earnings are conditional on positive earnings.

Building on this dataset, we first document several facts that motivate our model. We begin by showing that workers from different fields of study cluster into specialized sets of occupations. We then apply the fixed effects decomposition method of Porzio et al. (2022) to assess how much of this occupational heterogeneity arises from new cohorts entering the labor market versus incumbents switching jobs—both channels matter, though their relative contributions differ across occupations. Finally, a Kitagawa-Blinder-Oaxaca decomposition shows that a substantial fraction of occupational reallocation among new cohorts is driven by changes in field of study, rather than by differences in occupational entry within the same field. These findings indicate that students' field choices evolve over time, potentially in response to shifting labor demand, highlighting the need for a framework that jointly models education decisions and labor market dynamics.

Motivated by these facts, we develop a dynamic model of educational choices and labor supply that comprises two blocks: (i) education decisions of new entrants, and (ii) career decisions of incumbent workers. In the education block, students decide which field of study to pursue by

weighing expected labor market returns (i.e., income) against nonpecuniary benefits. Our framework allows for heterogeneity along three dimensions: non-pecuniary preferences for fields, the relative preference between income and non-pecuniary factors, and the value of a field of study as a function of observable characteristics, such as test scores and high school tracks. Importantly, field choices determine the distribution of future comparative advantages across occupations, thereby establishing the foundation for students' future productivity and outcomes on the labor market.

In the career block, once students graduate and enter the labor market, each new entrant draws a vector of productivities across occupations from a distribution determined by their field of study and other observable characteristics, thereby defining their comparative advantage profile.[2] Incumbent workers then choose occupations each period à la Artuç et al. (2010), facing partially stochastic switching costs that, together with their initial comparative advantage profiles, drive the observed occupational reallocation in response to labor market shocks.

To quantify the interaction between educational choices and labor market outcomes, we estimate the career choice and the field choice blocks of the model. For the career choice block, we adopt a two-stage procedure following Traiberman (2019), building on Arcidiacono and Miller (2011). In the first stage, we apply the Expectation-Maximization (EM) algorithm to recover key parameters such as the conditional distribution of workers' comparative advantages and their occupational transition rates. In the second stage, we regress these transition rates on wage differentials—following Artuç et al. (2010)—to identify the remaining parameters. In the field choice block, we exploit Denmark's centralized admission process—which provides complete ranked preference data—to recover latent heterogeneity in students' field choices using a novel estimation strategy, which built on (Agarwal and Somaini, 2018; Luflade, 2019; Fack et al., 2019). Finally, we leverage time variation in the model-implied returns to fields of study to estimate how responsive students are to changes in expected future income.

We next discuss the estimation results for each block. To examine how different fields translate into career opportunities, we construct a measure of each field's comparative advantages across occupations. We find that some fields are highly specialized, mapping narrowly to a small set of occupations, while others are much more general. To quantify the economic implications of this specificity, we estimate the increase in the occupational exit rate following a 1% decline in earnings. On average, a 1% decline in the net present value (NPV) of an occupation increases the exit rate by

[2]Abilities may vary both within and across educational backgrounds; for instance, social science programs may provide better quantitative skills but weaker communication skills compared to humanities programs on average, but this is not true of all graduates in these programs.

11%. This average obscures substantial heterogeneity across fields: workers in the most generalized fields are nearly six times more responsive to shocks than those in the most specialized fields.

Turning to students' field choice decisions, we identify three key findings. First, students are responsive to changes in future income—a 1% increase in the NPV of a field raises entry into that field by 0.14 percentage points, corresponding to an elasticity of 0.92. Second, heterogeneity in preferences plays a central role in shaping field choices. Some students exhibit strong preferences for particular fields, making them highly elastic within related disciplines yet unlikely to switch more broadly. For example, students with a strong preference for education are more likely to substitute into care or health-related fields, whereas STEM-oriented students primarily switch within STEM or into business. Finally, explicitly modeling this heterogeneity is crucial for capturing realistic substitution patterns. We show that the implied substitution matrix—which describes how students reallocate when a given field experiences a shock—closely resembles observed ranking preferences. By contrast, a model without this heterogeneity, akin to a simple logit model that assumes independence of irrelevant alternatives, produces implausible substitution patterns that fail to match observed choices.

We close the model by incorporating the labor demand side to assess the general equilibrium (GE) effects of counterfactual policies and labor market shocks. Specifically, we extend a standard multi-sector framework in which each sector employs different mixes of occupations, allowing for heterogeneous substitution patterns in response to external shocks such as foreign price changes and AI-driven technological advances. Using the GE model, we first examine the impact of capacity constraints in higher education by simulating a policy that reduces restrictions on students' field choices. We consider two scenarios: partial flexibility, in which we remove capacity constraints but maintain high school tracking—thereby limiting each student to fields accessible under their track—and full flexibility, where students may enter any program.[3] Our results indicate that relaxing these constraints decreases GDP by 0.3% under partial flexibility and raises it by 1.7% under full flexibility. Though field switching is substantial—9% of students reallocate their field of study under partial flexibility and 21% under full flexibility—the net effect on occupational reallocation remains modest, explaining the more limited aggregate income gains.

Second, we evaluate the economic impacts of two labor market disruptions—international trade shocks and AI-driven technological shifts—under varying degrees of educational flexibility. Our

[3]In Denmark, upper secondary education is organized into distinct tracks, each granting access to a set of higher education programs. Though the system has become more flexible over time, we maintain these track-based constraints under *partial flexibility*. We provide more details in Section 2.3 and Appendix E.1.

results reveal substantial long-run increases in welfare and income under both shocks, with gains significantly larger when students can freely adjust their field choices. Allowing major-switching flexibility leads to extensive reallocation across fields, enhancing overall economic adaptability. Importantly, these gains vary considerably across individuals with different preference types. Students with greater elasticity of field choice experience larger income improvements, highlighting that intrinsic responsiveness to labor market changes significantly influences the economy's ability to adjust effectively to disruptions.

Our paper is related to several strands of literature. First, we build on the literature that employs GE models to study the dynamic response of workers to labor market shocks. The labor market side of the model draws from studies that examine how workers with heterogeneous abilities and/or horizontally differentiated skills make career decisions in the presence of mobility frictions. These models have been widely used to study trade-induced reallocation (Artuç et al., 2010; Dix-Carneiro, 2014; Traiberman, 2019), labor market responses to immigration (Llull, 2018; Burstein et al., 2020; Monras, 2020; Khanna and Morales, 2017), technological advances (Dvorkin and Monge-Naranjo, 2019; Humlum, 2021; Adao et al., 2024), and labor market sorting and search patterns (Lise and Postel-Vinay, 2020). Our key contribution is to extend this literature by endogenizing skill acquisition through forward-looking educational choices. This allows us to analyze not only how incumbent workers reallocate in response to shocks but also how the skill distribution evolves as new entrants select fields of study in anticipation of future labor market conditions. Embedding these choices within a GE framework enables us to examine both short-run labor market adjustments and long-term shifts in workforce composition, offering new insights into the propagation of shocks and the role of policy in shaping skill supply.

We also contribute to the large literature on the the joint determination of human capital and labor market returns. While much of the prior literature focuses on vertical education decisions, e.g., no college, vocational school, or college (Blanchard and Olney, 2017; Heckman et al., 2018; Adda and Dustmann, 2023), a growing body of work examines variation in returns across fields of study (Gemici and Wiswall, 2014; Kirkeboen et al., 2016; Bleemer and Mehta, 2022b; Andrews et al., 2022) and the role of major specificity in shaping career trajectories (Kinsler and Pavan, 2015; Leighton and Speer, 2020).[4] Our paper is most closely related to studies that estimate models of human capital accumulation alongside labor market processes (Heckman et al., 1998; Arcidiacono,

[4] Altonji et al. (2012b) and Altonji et al. (2016) provide extensive surveys, including discussions of structural modeling issues.

2004; Beffy et al., 2012) and builds on the seminal work by Keane and Wolpin (1997), which integrates labor demand into college decisions.[5] We extend this literature by linking field choice to a dynamic labor market model with multiple occupations and sectors, allowing us to study how occupational heterogeneity shapes labor market responses to shocks. Specifically, we show that both incumbent worker transitions and new cohort decisions are key margins of adjustment, which interact as incumbent reallocation affects expected earnings for new entrants, while new entrants influence future labor supply. Furthermore, many shocks—such as trade, technological change, and immigration—primarily impact a subset of occupations directly, with broader effects emerging through equilibrium reallocation. For instance, STEM-intensive fields like R&D may be highly exposed to trade or artificial intelligence, while fields such as medicine may be less affected. Our framework unifies field choice and dynamic occupational transitions, providing a tool to analyze the long-term effects of targeted labor market disruptions and policy responses.

Finally, our paper contributes to the expanding literature on the demand for fields of study (Robles and Krishna, 2012; Wiswall and Zafar, 2015; Agarwal and Somaini, 2018), by nesting field choice within a broader labor supply framework. Methodologically, we are close to the literature exploiting centralized assignment mechanisms as a source of identification for preferences over college majors (Kapor et al., 2024; Fack et al., 2019; Larroucau and Rios, 2022). Unlike much of the existing literature, which primarily examines educational decisions in isolation, our approach links them to occupational reallocation, allowing us to assess how shocks influence both individual choices and aggregate labor market dynamics.

The paper proceeds as follows. Section 2 describes the data and institutional background. Section 3 presents key stylized facts motivating our model. Sections 4–6 outline the model, estimation strategy, and results, focusing on labor supply. Labor demand is discussed in Section 7, where we close the model and conduct counterfactual analyses. Section 8 concludes with directions for future research.

2 Data and Institutional Background

In this section, we describe the data sources and definitions used in the empirical analysis. We then provide a brief description of the Danish education system and show prominent features in students'

[5] While our paper focuses on higher education, a substantial literature highlights the importance of early human capital accumulation (e.g., Caucutt and Lochner, 2020).

applications and field choices.

2.1 Data Sources

We employ several administrative datasets covering the Danish economy from 1996 to 2019. The first is the *Integrated Database for Labour Market Research* (IDA)—a panel dataset with information on the universe of Danish workers over 15, both native and non-native. IDA includes information on employment status, earnings, ISCO-based occupation, age, gender and other demographic variables. Throughout the text "earnings" refers to the earnings concept in this register—which are total annual earnings in workers' main job. Workers are assigned an employer and occupation based on their employment status in November of each year.

Our main source of information regarding individuals' education comes from the *Education Register* (UDDA). The data include the type of education completed, the level of education and the year of graduation. The education register also includes a mapping from all education programs to the International Standard Classification of Education (ISCED) which allows us to classify education programs into fields of study. We supplement the education register with information regarding student applications to higher education programs from the *Coordination Register* (KOT). The data are provided by the Coordinated Admissions System, which is run by the Ministry of Higher Education. For each student, we know the set of programs to which they applied, whether they received admission in a program, the system under which they received admission, their high-school GPA and track. We hand-link the education programs in KOT to UDDA and calculate admission cutoffs based on all applications data.[6]

In addition to these datasets, we bring in standard firm level data from the *Firm Accounting Statistics Register* (FIRE). This allows us to identify workers' industries. Finally, we use the *Danish Foreign Trade Statistics Register* (UHDI). These data are product-firm-origin/destination level customs data on all international transactions. They contain data on value of imports and exports at the CN8 level, and often also contain quantity and unit value information.

2.2 Defining Fields and Occupations

We define the level of education based on the ISCED system, and divide students into three levels: high school graduates, short-cycle education, and college and above. The last group encompasses

[6]We also need to determine the effective GPA cutoff threshold for each program with a binding testing cutoff, as this is not available (digitally at least). Appendix E.2.1 contains a complete description of how we concorded KOT to UDDA and calculated admission cutoffs.

both Professional Bachelor's and University degrees and is the primary focus of our analysis.[7] Our definition of fields is based on the Danish variant of the ISCED-F classification of fields. The ISCED-F classification breaks fields into 27 broad categories. We aggregate several codes—especially in the natural sciences—due to their small size. We also break apart the health code into programs at the Professional Bachelor's level and the University level. Appendix Table 11 contains our list of educational fields and their ISCED-F equivalent codes.

Occupations in Denmark are recorded according to the International Standard Classification of Occupation (ISCO), developed by the ILO. There is a break in 2009 when codes were updated. We use the concordance of Humlum (2021) to create a time-consistent series of occupations. For data quality issues, we focus on 2 digit occupational codes. After removing or pooling small occupations, we are left with 22 occupations. When the data breaks, there are several quality issues in the data. Appendix E.3 contains a complete description of how we concord and aggregate occupations over time and how we deal with variation in the quality of occupational codes (e.g., imputations). Appendix Table 12 contains our list of occupations and their ISCO-08 equivalent codes. Finally, in appendix E.3.1 discuss how we use the ONET Database to construct occupational tasks.

2.3 The Danish Education System

The Higher Education System in Denmark

Our analysis focuses on higher education which includes two main types of programs: Professional Bachelor's degrees and University Bachelor's degrees. Both of these programs are classified as Level 6 according to the ISCED, and are commonly pooled in studying skilled workers in Denmark (Hummels et al., 2014). Professional Bachelor's degrees, also referred to as Academy Bachelor's degrees, typically require three years of study and are designed to provide students with practical, job-oriented skills tailored to meet labor market demands. These programs cover fields such as business, education, engineering, and various forms of vocational training, emphasizing technical expertise and immediate employability. University Bachelor's degrees, on the other hand, follow a more academically focused path, comparable to a standard American four-year degree but completed in three years.[8] These programs are oriented toward theoretical knowledge and research, preparing students for further academic pursuits or careers requiring in-depth expertise in their chosen fields.

[7] While we focus on this group in estimation, in all of counterfactuals we take account of all workers, treating the level of educational background as exogenous and fixed.

[8] While the Bachelor's is awarded after three years, well over 90% of students pursue a Master's program. In our analysis, we treat this as one continuous 5 year program.

An increasingly large share of Danish students pursue a degree from a Professional Bachelor's or University (henceforth "college"). Appendix figure 9 plots the share of 30 year olds at each level of education. At the start of our sample, in 1995, around 18% of the sample has some form of college, and this more than doubles to 41% by the end of the sample.[9]

In Denmark, as in much of Europe, the higher education system emphasizes early specialization, differing significantly from the U.S. model of majors and minors. Students apply directly to specific programs, such as engineering or business, rather than entering with a broad area of study. Unlike the U.S. model that offers general education requirements and flexibility in course selection, the curriculum in Denmark is more predefined, with less room for electives outside the chosen field.

High school tracks play a critical role in shaping students' options for further education. Danish upper secondary education consists of several tracks, each influencing the types of programs students can pursue beyond high school (see Appendix E.1 for more information on higher education and high school tracks in Denmark). While the system has become more fluid over time, with for example options for students to take supplementary courses to broaden their eligibility, these tracks still determine much of the access to higher education. To this end, in our estimation, we use upper secondary background, including specialization when available, to construct students' choice sets for higher education. We also use high school tracks as a control when estimating students' returns and preferences over different fields of study.

The Application System

Entry into both Professional Bachelor's programs and University Bachelor's programs occurs through a centralized admissions process. Students can apply to up to eight programs (each application being a combination of field and institution). Each program has its own GPA threshold for entry, determined by student demand and program supply, and regulated by the Danish Ministry of Higher Education and Science. Some programs also have specific prerequisites, such as sufficient mathematics preparation. Students in Denmark are provided with significant resources to help them make informed education choices and improve their information about alternatives educational paths.[10] The *UddannelsesGuiden* (Education Guide) website provides detailed information on program-specific prerequisites, historical GPA cutoffs, admissions numbers, labor market outcomes,

[9] We condition on our sample, and our use of 30 year olds will include students who dropped out of college, and students who are still studying. Nevertheless, the estimates line up reasonably well with the reported national statistics from Denmark, available at `https://www.statbank.dk/HFUDD11`.

[10] This is line with the fact that information plays a crucial role in educational choices, and that students often face significant information gaps when making education decisions (Wiswall and Zafar, 2015; Kapor et al., 2020).

demand trends across fields of study, and subject matter expectations.[11] Study guidance centers (*Studievalg*) further support students through counseling services. Nevertheless, while students inevitably face some gaps in information and cannot fully predict the returns to different fields of study, significant efforts are made to provide them with comprehensive guidance and resources.

The Danish admissions system uses a Deferred Acceptance mechanism, ensuring that if a student is ineligible for their first choice, it does not affect their chances of being considered for subsequent choices. This structure encourages students to truthfully rank their preferences (Fack et al., 2019). However, some students may still choose not to apply to programs where they believe they are ineligible or incur psychological costs from the process. In our estimation, we assume students apply within a "realistic" set of alternatives, consistent with the "score bounds" approach of Kapor et al. (2024). While the majority of admissions to higher education programs take place through GPA cutoffs (also referred to as the "Quota 1" system), a fraction of admissions take place based on supplementary exams or other criteria for students who were rejected from their Quota 1 application (the "Quota 2" system). In our analysis, we only focus on Quota 1 admissions. It is also not uncommon for students to wait between high school and college if they do not receive their top choice. Although the system allows for some fluidity, quotas remain binding for many programs, and delaying applications can result in penalties. Given these conditions, most students apply to multiple programs and many are admitted to or attend their second (or first) priority.

Figure 2 displays the unconditional distribution of first priorities for students across fields of study, as well as the growth in fields over time. The most popular fields are in Care and Medium Health, followed closely by Business, Social Sciences and Education. STEM fields, when including High Health, account for roughly one quarter of students, while Arts and Humanities account for around one tenth. The most notable changes in preferences involve ICT and Education. ICT, which begins on a very small base, has seen explosive growth, as have Business and Arts to a lesser extent.

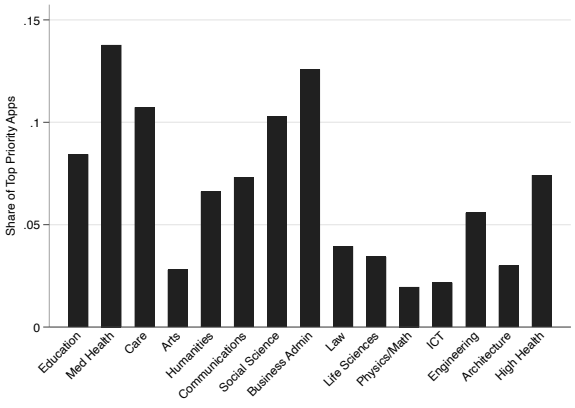
Table 1 provides insight into students' preferences by examining the relationship between first-choice fields (columns) and second-choice fields (rows). Two patterns stand out. First, the data show strong diagonal dominance across fields, indicating that most students select a second-choice program within the same field as their first choice.[12] Second, preferences vary significantly by

[11] See https://www.ug.dk/

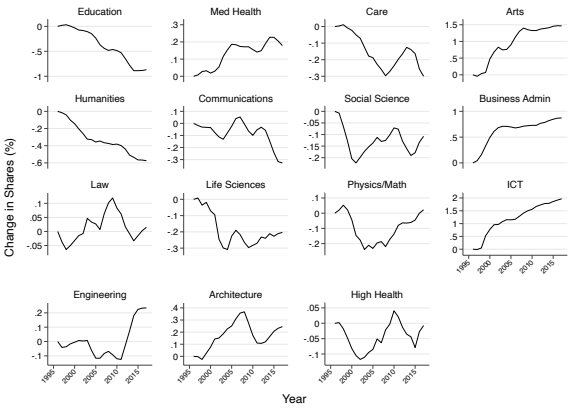
[12] In our subsequent analysis, we do not distinguish between students choosing a similar program within the same institution (e.g., Physics versus Mathematics at Aarhus) or applying to the same program at different institutions. We do not find a clear pattern that suggests students engage in one kind of behavior or another: 37% of the time students apply across programs within the same institution, and the other 63% of the time they apply for similar programs across institutions.

Figure 2
Summary Statistics on First Priority

(a) Unconditional Distribution



(b) Growth Over Time



Note: Panel (a) plots the share of high school graduates who apply to college education through the KOT system–regardless of final admission status–who apply to each field of education. The sample is pooled across years, and only fields with at least 100 applications over the course of the entire sample period are considered. Panel (b) plots the change in shares over time relative to the base year, 1995. Time series plots are plotted using a 3 year MA smoother.

discipline. In STEM fields such as life sciences, physics/mathematics, and engineering, there is notable clustering, with a large share of students selecting second or third choices within these closely related fields. In contrast, preferences in Education and Health appear more narrowly focused, with students showing a lower likelihood of considering programs outside their initial field of interest. This suggests a strong tendency toward specialization in these disciplines.

3 Motivating Model Ingredients

In this section, we present three stylized facts about the relationship between fields of study and career choices. Our analysis focuses on students and workers with higher education. The purpose of this section is to motivate our model, which incorporates both a career choice component and a field choice component. The facts presented below highlight that occupational reallocation is influenced by two margins: how incumbents reallocate themselves and how new students adjust their human capital investment decisions.

Fact 1: Fields of study are associated with specialization in different occupations.

We first demonstrate that students from different fields of study tend to work in a narrow band of occupations, though the extent of this specialization varies. To measure how specialized a field is,

	First Priority														
Second Priority	Education	Care	Med Health	Arts	Humanities	Social Science	Communications	Business Admin	Law	High Health	Life Sciences	Physics/Math	ICT	Engineering	Architecture
Education	75.92			4.54	4.19										
Care	8.38	88.43	5.03												
Med Health	3.51	4.02	85.57							7.56	4.80				
Arts				46.98											7.54
Humanities				13.27	51.48	13.89	16.81		5.72						
Social Science				9.40	15.40	49.09	11.60	9.46	16.07	4.10		7.62	4.69		
Communications				9.98	15.48	7.46	42.61	10.59	5.84				3.79		
Business Admin						9.32	14.11	66.77	22.30			4.62	6.98	4.13	
Law						4.74			39.98						
High Health										67.36	11.37	3.57			
Life Sciences										9.25	50.37	15.56	4.60	7.53	
Physics/Math											10.61	39.84	4.85	6.39	
ICT												4.17	46.64	9.18	
Engineering											11.37	13.35	19.15	53.37	14.73
Architecture				5.85										10.11	55.75

Table 1
$P(Priority_{r+1}|Priority_r)$

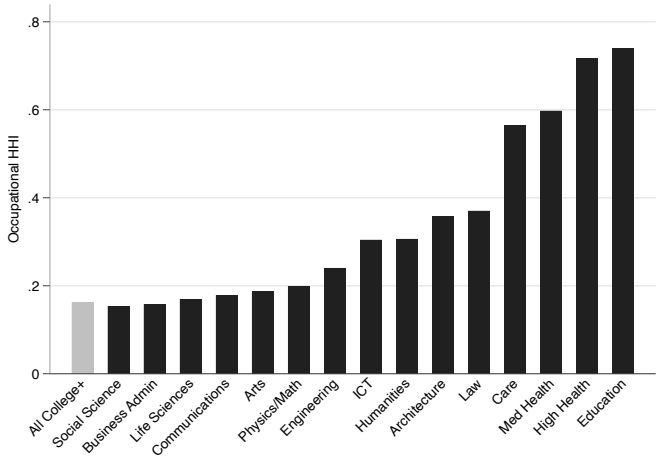
Notes: The columns display students' first-choice fields and the rows display students' second-choice fields. The definition of fields is based on the Danish variant of the ISCED classification fields and is discussed in Section 2.2. Cells are censored below 3.5% for readability.

we calculate the Herfindahl-Hirschman Index (HHI) across occupations for each field:

$$HHI_m = \sum_{o \in \mathcal{O}} s_{o|m}^2,$$

where $s_{o|m}$ is the share of workers from education field m working in occupation o. We calculate shares pooling across the whole economy. Figure 3 plots the HHI for each field of study, as well as the HHI calculated by pooling all fields (in grey).

Figure 3
Field of Study-Occupation HHIs



Notes: The black bars represent the HHI for different fields of study. The definition of fields is based on the Danish variant of the ISCED classification fields and is discussed in Section 2.2. The grey bar represents the HHI calculated by pooling workers from all fields.

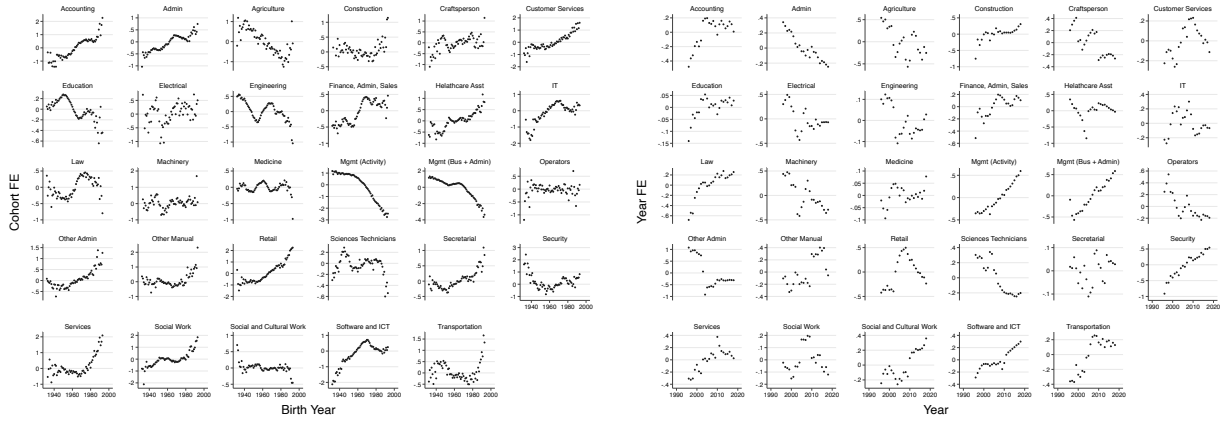
We highlight two takeaways from this graph. First, fields of study are generally highly specialized. The pooled HHI—unconditional on field of study—is about .2, while the average HHI conditional on field is nearly twice as high. Second, specialization varies significantly across fields. For example, Social Science has an HHI slightly below the pooled average, while health and education fields exhibit near-complete specialization. This variation in specialization suggests that workers from different fields may experience differing levels of vulnerability to economic shocks, with those from highly specialized fields potentially facing greater challenges in adapting to shifting labor market conditions.

Fact 2: Occupational reallocation involves both new entrants and incumbents.

We now examine the role of new entrants in occupational reallocation. We show that reallocation occurs through two channels: incumbent workers switching occupations (or exiting the labor force) and new cohorts entering a different mix of occupations.

Figure 4
Reallocation Decomposition

(a) Cohort Fixed Effects (b) Year Fixed Effects



Notes: The figure plots the estimated cohort and year fixed effects from a projection of these variables onto log occupation shares (see (1)). Plots are all MA3 smoothed series. The null hypothesis of perfect labor market fluidity is a horizontal line at 0 for all cohort plots.

To quantify these channels, we adopt the approach of Porzio et al. (2022). Let $s_{o|c,t}$ be the share of workers in cohort c working in occupation o at time t. We project this share on cohort and time fixed effects for each occupation:

$$logs_{o|c,t} = \varphi_{oc} + \psi_{ot} + \varepsilon_{oct}. \tag{1}$$

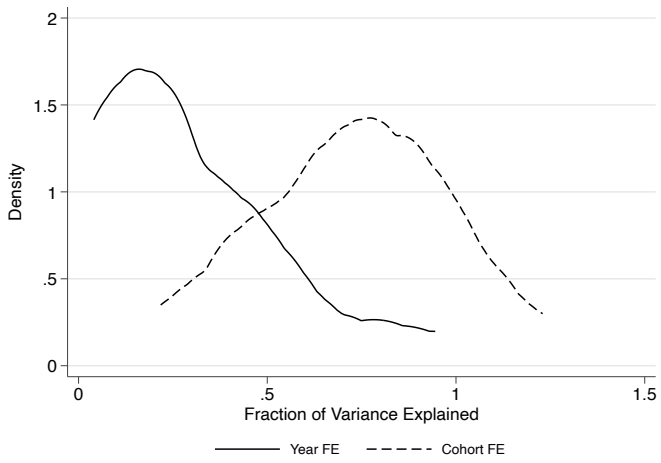
If workers are fully mobile across occupations (or equally mobile across generations), the cohort effects would be irrelevant, as reallocation would primarily reflect common demand or supply shocks that are absorbed by the year fixed effects. Conversely, if workers are completely immobile, reallocation would occur entirely through new cohorts entering different occupations, making cohort effects the primary driver of changes in occupational shares.

Figure 4 plots the cohort and year fixed effects, where null effects—indicating minimal reallocation through that channel—would appear as flat lines bunched around zero. Instead, both cohort and year effects are important, though their relevance varies across occupations. In some cases, cohort

effects are minimal, while in others, time effects are less important. This heterogeneity highlights the need for a model that incorporates both channels of occupational adjustment.

To further assess their relative importance, Figure 5 reports the ratio of the variance in fixed effects to the variance of the fitted log shares: $\sigma^2_{\varphi_{oj}}/\sigma^2_{\widehat{\log s_{o|c,t}}}, j \in \{c,t\}$. The distribution of these ratios suggests that cohort effects generally explain a larger share of the variation, underscoring again the role of new cohort entry in occupational reallocation.[13]

Figure 5
Variance Explained by Cohort and Year Effects



Notes: This figure plots the density distribution of the fraction of explained variance in log employment shares across occupations that is attributable to cohort and year fixed effects. Specifically, for each fixed effect f, it plots $\hat{\sigma}^2_{\hat{f}}/\hat{\sigma}^2_{\log s_{o|ct}}$ across occupations o.

Fact 3: New cohorts reallocate through field of study choices.

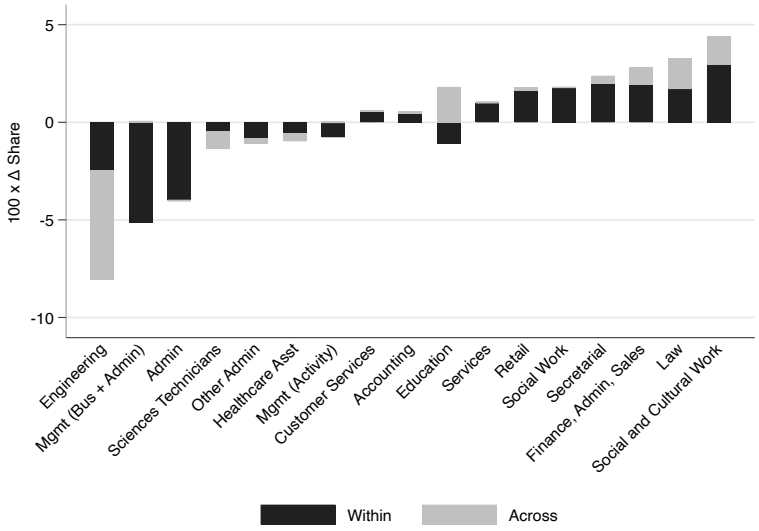
Our final fact documents the extent to which occupational reallocation across cohorts occurs through changes in the composition of fields of study. To quantify this, we perform a Kitagawa-Oaxaca-Blinder decomposition of changes in initial occupation shares for new cohorts into within- and across-field components. Specifically, we use the follwoing decomposition:

$$\Delta s_{ot} = \underbrace{\sum_m \Delta s_{o|m} \times s_{m,t-1}}_{\text{Within}} + \underbrace{\sum_m \Delta s_m \times s_{o|m,t-1}}_{\text{Across}} + \underbrace{\sum_m \Delta s_m \times \Delta s_{o|m}}_{\text{Growth Covariance}}. \tag{2}$$

[13]This ratio can technically be larger than one due to the covariance of cohort and year fixed effects, though this term is generally negligible.

Figure 6 plots this decomposition for occupations with at least a 5% change in share.[14] The light gray bars capture the across-field component, representing changes in occupational shares driven by shifts in field choice. For most occupations, this channel explains a substantial portion of the observed changes. Some exceptions arise, such as management and administration, where field choice appears less important, likely reflecting a weaker link between field choice and occupation.

Figure 6
Decomposing the New Cohort Entry Margin



Notes: Plots the decomposition of occupational reallocation into across- and within- major components according to the Kitagawa-Oaxaca-Blinder decomposition described in (2). Occupations are sorted based on total growth over 1998 to 2018. Only occupations with at least 1% change plotted. Appendix 10 contains all occupations and second order terms.

Summary: Together, these facts provide a clear view of how the occupational structure of the economy changes: incumbents and new cohorts both play a large role in shaping occupation reallocation. Moreover, new cohorts adjust not only their initial occupation choices but also shift their field of study over time. This pattern underscores the importance of modeling both labor supply decisions and field choice to capture the full range of occupational adjustment mechanisms. With this in mind, we now turn to our model of labor supply and field choice.

[14]The covariance term is generally small and is omitted here. The full decomposition for all occupations is provided in Appendix Figure 10.

4 A Model of Labor Supply and Field Choice

In this section we describe students' education and labor supply decisions, with labor demand left to Section 7. Our labor supply model combines three main ingredients: first, students make forward-looking choices over which field to pursue; second, workers with different educational backgrounds make forward-looking occupational choices; third, we incorporate latent heterogeneity in preferences over fields, as well as latent comparative advantage across occupations.

4.1 Environment

We consider a dynamic, small, and open economy. Time, indexed by t, is discrete. There is a measure $\bar{L}$ of workers, differentiated by their field of study $m \in \mathcal{M}$ and their labor market type $k \in \{1, ..., K\}$. The index k serves as a flexible, non-parametric way to capture unobserved heterogeneity across individuals. It is known to the worker but unobserved by the econometrician, and can influence both worker productivity and the costs associated with switching occupations, as detailed later. Each worker supplies one unit of labor inelasticially to an occupation $o \in \mathcal{O}$. Workers exit the labor force at an exogenous rate δ and are replaced by a new cohort. New cohorts are characterized by their latent preference type l, their test scores s, and other covariates X.

The model unfolds in two stages. Upon completing high school, individuals enter the *field choice block*, where they select a field of study based on anticipated returns and preferences, both of which may depend on their initial characteristics. After graduation, they enter the *career choice block*, where they hold a set of comparative advantages across occupations and make a forward-looking decision on their initial occupation. In each subsequent period, incumbent workers may switch occupations, subject to switching costs.

We next describe the model in reverse order: starting with the occupational decisions of incumbent workers, followed by the entry decisions of new cohorts, and finally the field of study choices.

4.2 Labor Markets

4.2.1 Incumbent Workers

At the *beginning* of period t, incumbent workers who spent the previous period in occupation o choose an occupation o'. Workers are characterized by (k, m) and supply human capital $H_o(k, m)$ to occupation o. Switching occupations entails a moving cost $C_{oo'km}$, which depends on the

worker's type and field, and the origin-destination pair. Workers also receive a stochastic occupation preference shock ε_{ot} in each period.

To parameterize human capital, we project each occupation onto Z tasks. Each occupation is represented by a vector $\zeta_o = (\zeta_{1,o}, \cdots, \zeta_{Z,o}) \in \mathbb{R}^Z_{++}$, the positive orthant of $\mathbb{R}^Z$, where $\zeta_{z,o}$ denotes the loading on task z in occupation o. A higher value of $\zeta_{z,o}$ indicates that task z is more important for output in that occupation. Each worker type has task-specific abilities $\alpha_{k,z}$, and their log output in an occupation is the dot product of *alpha* and ζ. We additionally include a field effect, α_{om}, which scales the human capital supply of workers in each occupation based on their field of study, and a type effect α_k, which captures the absolute advantage of type-k workers.[15] Combining, we have:

$$\log H_o(k, m) = \alpha_{om} + \alpha_k + \sum_{z=1}^{Z} \alpha_{k,z}\zeta_{z,o}. \tag{3}$$

A few comments are in order. First, workers' comparative advantage across occupations arises from differences in their labor market type k.[16] Second, the α_{om} terms allow workers in different fields of study to have distinct comparative advantages, even conditional on their type k. For example, workers with an education background in STEM may have large α_{om} values in technical occupations, such as engineering. Finally, the conditional distribution of types k depends on the chosen field of study m, a relationship we discuss in detail later.

Returns to an occupation are determined by both workers' human capital and competitively determined skill prices w_{ot}. The worker's Bellman equation in period t is given by,

$$v(k, m, o, \varepsilon) = \max_{o'} w_{o't} H_{o'}(k, m) - C_{oo'km} + \varepsilon_{o't} + \beta E_t V(k, m, o'), \tag{4}$$

where $V \equiv \int v(\cdot, \varepsilon) dG(\varepsilon)$ is the value function v integrated over shocks, and the effective discount factor is β, which is the product of the time discount rate, $\tilde{\beta}$ and the probability of exit, δ.

We assume preference shocks, ε_o, are conditionally independent of other variables and follow a Gumbel distribution with mean zero and scale parameter ν. Under these assumptions, the value

[15] Technically we use one index for all productivity differences across workers—absolute and comparative advantage. This is without loss of generality. For example, one could have workers with the same comparative advantage but different levels at each task simply by having two types with parameters shifted up by a constant for one type.

[16] As an example, consider a simple economy with two occupations, two tasks, and two types. Suppose the task vector for the first occupation is $\zeta_1 = (1, 0)$ and for the second occupation $\zeta_2 = (0, 1)$. Similarly, let worker abilities be $\alpha_1 = (1, 0)$ and $\alpha_2 = (0, 1)$. In this case, type 1 workers would produce output only in occupation 1, while type 2 workers would only produce output in occupation 2. Despite this horizontal differentiation, neither type has an absolute advantage.

function V simplifies to:

$$V_t(k,m,o) = \nu \log\left(\sum_{o'} \exp\left(\frac{-C_{oo'km} + w_{o't}H_{o'}(k,m) + \beta E_t V_{t+1}(k,m,o')}{\nu}\right)\right). \tag{5}$$

Moreover, the corresponding policy function, which determines the transition probabilities across occupations, is:

$$\lambda_t(o'|o,k,m) = \frac{\exp\left(\frac{-C_{oo'km}+w_{o't}H_{o'}(k,m)+\beta E_t V_{t+1}(k,m,o')}{\nu}\right)}{\sum_{o''}\exp\left(\frac{-C_{oo''km}+w_{o''t}H_{o''}(k,m)+\beta E_t V_{t+1}(k,m,o'')}{\nu}\right)}. \tag{6}$$

(6) can be used to interpret ν as governing the elasticity of occupational switching with respect to fundamentals. At the extremes, as $\nu \to 0$, the transition probabilities collapses to a deterministic outcome where all workers choose the occupation with the highest expected payoff. Conversely, as $\nu \to \infty$, workers allocate themselves uniformly across occupations, becoming indifferent to differences in deterministic payoffs.

4.2.2 New Entrants

Like incumbents, new entrants are characterized by a (k,m) pair. Upon entering the labor market, they chose an initial occupation, according to the utility function,

$$v_t^e(k,m) = \max_o -C_{o,km}^e + \varepsilon_o + w_{ot}H_o(k,m) + \beta E_t V_{t+1}(k,m,o), \tag{7}$$

where $C_{o,km}^e$ denotes entry costs specific to occupation o for type (k,m) and ε_o is a stochastic preference shock at the initial entry, assumed to follow the same Gumbel distribution as for incumbents. The resulting choice probabilities for new entrants are given by:

$$\lambda_t^e(o|k,m) = \frac{\exp\left(\frac{-C_{o,km}^e+w_{ot}H_o(k,m)+\beta E_t V_{t+1}(k,m,o)}{\nu}\right)}{\sum_{o'}\exp\left(\frac{-C_{o',km}^e+w_{o't}H_{o'}(k,m)+\beta E_t V_{t+1}(k,m,o')}{\nu}\right)}. \tag{8}$$

The expected value of entering the labor market with type (k,m) is given by:

$$V_t^e(k,m) = \ \nu \log \sum_{o'} \exp\left(\frac{-C_{o,km}^e + w_{o't}H_{o'}(k,m) + \beta E_t V_{t+1}(k,m,o')}{\nu}\right).$$

4.3 Field Choice

We begin by describing the returns to a field and then outline how students construct their ranked choice lists. Students enter the field choice block with a test score, s, and a latent preference type,

$l \in 1, ..., L$. They may also have additional covariates, such as their high school background or gender, which we collect in the vector X. For simplicitiy, we omit these covariates in the following discussion, though all outcomes are implicitly conditioned on them on these characteristics.

Returns to Field

In deciding on a field of study, students consider the expected labor market returns associated with each option, denoted $V_t(m|s)$. These returns depend on two key factors. First, there is the common productivity effect, α_{om} in (3). This term reflects the average productivity associated with each field in different occupations. For example, degrees in Physics/Math may offer relatively high returns in technical occupations than the return in those occupations to communications degree holders. Second, each field of study induces a conditional distribution over labor market types $k \in K$, denoted $\pi(k|m, s)$.[17] This stochastic term allows for heterogeneity in outcomes among students within the same field, as students may possess varying talents that influence their career trajectories. Given this conditional distribution, the expected labor market return for a student with test score s choosing field m is:

$$V_t(m|s) = \sum_k V^e_{t+5}(k, m)\pi(k|m, s),$$

where $V^e_{t+5}(k, m)$ is the lifetime payoff to the student upon graduation with a labor market type k and a field m.

In addition to expected labor market returns, students' field choices reflect non-pecuniary factors. Each type-l student has a time-invariant utility for a field denoted by ϑ_{lm}, which captures intrinsic preferences unrelated to earnings. The relative importance that students assign to future earnings varies across types too and is governed by the parameter θ_l. Students may also prefer fields that offer a larger number of programs. This is captured by a love-of variety effect, denoted by $\upsilon \log(N_m)$, where N_m is the number of programs available in field m. For example, the "Math/Physics" field may encompass multiple mathematics and physics programs across institutions, which some students may view as desirable.[18] Finally, students receive a conditionally independent Gumbel(0,1) preference shock over fields, ε^m_m, which introduces idiosyncratic variation in their choices.

[17] An alternative interpretation is that the college education production function is given by $\pi(k|m, s)$, where field choice m and score s, along with other observable characteristics, serve as inputs. The output is a distribution of labor market types k, which, as described in Section 4.2, is linked to worker productivities and occupational switching costs.

[18] As the focus of our paper is not on institutional quality, nor on within-field variation, we treat programs within a field of study as symmetric, though the average quality or popularity at the field level is partially captured by the cut-offs in the admission process, which we describe below.

Combining these components, the utility for students of preference type l with test score s assigns to field m is:

$$\tilde{U}_{msl} = \vartheta_{lm} + \theta_l V_t(m|s) + \upsilon \log(N_m) + \varepsilon_m^m. \tag{9}$$

Henceforth, we use $\tilde{U}$ for the payoff including the shock, ε, and use U to denote the payout without the shock:

$$U_{msl} = \vartheta_{lm} + \theta_l V_t(m|s) + \upsilon \log(N_m).$$

Ranked List Determination

Upon receiving their preference shocks ε_m^m, students construct a ranked list of fields of length R. Field assignments are determined using the deferred acceptance algorithm, and we assume that students truthfully reveal their preferences over their choice sets.[19] Students' choice sets, $\mathcal{C} \subset \mathcal{F}$, may depend on their test scores and high school background. For example, students with low scores may avoid applying to highly selective programs, while those without prior math coursework may exclude math-intensive fields.

For ease of notation, let m_r denote the field ranked in position r, with m_1 being the top-ranked field, and so on. Under the Gumbel assumption, the probability of observing a ranked list $(m_1, m_2, ...,)$ is given by,

$$P(m_1, m_2, ..., m_R; s, \mathcal{C}) = \prod_{r=1}^{R} \frac{\exp(U_{lm_r t}(s))}{\sum_{m \in \mathcal{C} \setminus m_{-r}} \exp(U_{lmt}(s))}.$$

This expression is a concatenated probability of the usual logit choice probabilities, with the choice set adjusted at each step.

Our demand system is tractable and encompasses many models as special cases. First, in the absence of unobservable heterogeneity or occupational choice component in the labor market, the model reduces to one where workers assess the net present value of earnings based on cross-sectional average earnings in each field. Second, when $\beta = 0$, the model simplifies to one where workers focus solely on their initial earnings in the job market.

[19] Since the ranked list is truncated, students may have an incentive to include a "safety" choice. However, given our level of aggregation, few students are actually constrained by the list length, so we abstract from this issue.

5 Estimation

We present our estimation strategy in the same order as the model: first we discuss the career choice block of the model, then turn to the field choice block of the model.

5.1 Labor Market Parameters

5.1.1 Labor Supply of Incumbent Workers

Our estimation of labor supply parameters for incumbent workers follows Traiberman (2019), which was built on Arcidiacono and Miller (2011) and Artuç et al. (2010), and is extended to incorporate field differences. We adopt a two-stage estimator: first, we use the EM algorithm to estimate the parameters governing human capital supply, the distribution of latent labor market types, and the occupation transition probabilities; second, we estimate the remaining structural parameters—the switching cost parameters (C) and ν, scale of the Gumbel distribution that governs the elasticity of occupation switching.

Turning to the first stage, we assume that individual i's observed log wages at time t are the model's wages plus a Gaussian noise term.[20] Let $f(w_{it}|k,o,m,t;\alpha)$ be the Gaussian PDF of log wages conditional on type, field, occupation, and year, with human capital parameters α. Let q_k be the probability of being type k. The likelihood contribution of individual i is given by,

$$\tilde{\mathcal{L}}_i = \sum_{k=1}^{K} q_k \times \lambda^e(o_1, m|k) f(w_{i1}|k,o,m,1;\alpha) \times \left[\prod_{t=2}^{T} f(w_{it}|k,o,m,t;\alpha) \times \lambda_t(o_t|o_{t-1},m,k)\right],$$

where $\lambda^e(o_1,m|k)$ is the probability that a type-(m,k) individual starts in occupation o at $t=1$, while $\lambda_t(o_t|o_{t-1},m,k)$ is the probability that a type-(m,k) individual transitions to occupation o_t at time t, given they were in occupation o_{t-1} at $t-1$.

The full likelihood, given that q_k is unobserved, is the product over $\tilde{\mathcal{L}}_i$. Directly optimizing this function poses two key challenges: (1) the likelihood is not log-linear; and (2) λ's are outcomes of the solution of the structural model, which requires computing future continuation values and their expectations.

To address the first challenge, we employ the Expectation-Maximization (EM) algorithm. This requires defining an auxiliary function—the mean of the likelihood conditional on the distribution of k, as if it were known. The second challenge is more specific to our setting. As shown by Arcidiacono

[20] In Traiberman (2019), this noise term is treated as ex post iid heterogeneity. This can be included in the model, as it shifts the value of an occupation. Nevertheless, since this ex post shock has no effect on selection, we ignore it in our present setting.

and Miller (2011), we can bypass this issue by treating λ as a parameter to be estimated directly, then mapping it back to the structural parameters in the second stage. Specifically, Arcidiacono and Miller (2011) demonstrate that the distribution of heterogeneity and the parameters governing state payoffs can be consistently estimated by maximizing the following function (which we continue to call the log likelihood in a slight abuse of terminology):

$$\hat{\mathcal{L}} = \sum_{i=1}^{N} \sum_{k=1}^{K} \left[\sum_{t=1}^{T} \log f(w_{it}|k, o_{it}, m_i, t; \alpha) + \log \hat{\lambda}_t(o_{it}|o_{i,t-1}, m_i, k) \right] q_{ik}, \tag{10}$$

where q_{ik} is the probability that individual i is type k, and the hats atop the λ terms indicate that these are estimated directly from the data rather than derived from the model. Appendix F.1 contains details of implementation. Briefly, the EM algorithm proceeds by iterating on a guess of weights q_{ik}, which are updated according to Bayes' Rule, and a guess of parameters, which are updated by maximizing (10) treating q_{ik} as given. We also discuss several smaller details including an implementation of Wooldridge (2005)'s method for dealing with initial conditions, computational details for constructing $\hat{\lambda}$, and how we implement bootstrapping for the calculation of standard errors.

In the second stage of the estimation procedure, we apply the Hotz-Miller inversion (Hotz and Miller, 1993) to recover the switching cost parameters and the scale of the Gumbel distribution. After estimating the first stage, $w_{ot}H_o(k,m)$ is observed for each (o,k,m), as are the switching probabilities, $\lambda_t(o'|o,k,m)$.[21] As shown in Appendix F.1.2, equation (6) can be used to derive the following regression equation:

$$\log\left[\frac{\lambda_t(o'|o,k,m)}{\lambda_t(o|o,k,m)}\right] + \beta \log\left[\frac{\lambda_{t+1}(o'|o',k,m)}{\lambda_{t+1}(o'|o,k,m)}\right] = -\frac{C_{oo'km}(1-\beta)}{\nu} + \frac{1}{\nu}\left(w_{o't}H_o(k,m) - w_{ot}H_o(k,m)\right) + \xi_{oo'kmt+1}, \tag{11}$$

where $\xi_{oo'kmt+1}$ is a time $t+1$ expectational error. Compared to (6), this formulation differences out the current and next period's continuation values, addressing the remaining selection issue. Under rational expectations, the error is uncorrelated with time t variables, allowing (11) to be estimated via ordinary least squares.

[21] Even though these are generated regressors, we will avoid the cumbersome use of hat notation. Instead, it should be understood that these are estimated.

5.1.2 Labor Supply of New Cohorts

The only remaining parameter to estimate for new cohorts is the entry cost, C^e. We can apply a similar strategy as above to recover these parameters. Without loss of generality, we normalize the entry cost into occupation 1 ("Management") to zero. In Appendix F.1.2, this leads to the following estimating equation:

$$C^e_{o,km} = \Big(w_{ot}H(o,k,m) - w_{1t}H(1,k,m)\Big) + \beta C_{1,o,km} - \nu\left\{\log\frac{\lambda^e_t(o|k,m)}{\lambda^e_t(1|k,m)} + \beta\log\frac{\lambda_{t+1}(o|o,k,m)}{\lambda_{t+1}(o|1,k,m)}\right\} \quad (12)$$

We estimate C_e by taking the time average of this expression.

5.2 Field Choice

To estimate the field choice parameters, our approach involves two key steps. First, we derive a likelihood function under the assumption that each student's latent preference type l is known, or at least known with probability q_{il}. Second, we outline the procedure for estimating these types.

To carry out the first step, we begin by deriving a recursive representation of payoffs, which enables us to effectively leverage the panel structure of our application data. An output of labor market estimation is q_{ik}, the probability that individual i is type k. This can be projected onto observables in order to recover the education production function, $\pi(k|m,s)$. With this in hand, the payoff to a field m at time t for type l individuals is,

$$U_{tlm}(s) = \vartheta_{lm} + \theta_l\left[\sum_k \pi(k|m,s)V^e_{t+5}(k,m)\right] + \upsilon\log(N_{mt}), \quad (13)$$

It is useful to define:

$$\bar{U}_{tlm}(s) \equiv U_{tlm}(s) - \upsilon\log(N_{mt}) = \vartheta_{lm} + \theta_l\left[\sum_k \pi(k|m,s)V^e_{t+5}(k,m)\right]. \quad (14)$$

The right hand side of (14) is *not* directly observable, because V^e is never directly estimated in the labor market estimation—recall the strategy for estimating labor supply parameters relied on differencing out these continuation values.[22] To resolve this issue, we treat the utility in the last period, i.e., $\bar{U}_{Tlm}(s)$ as another parameter to be estimated directly, where T refers to the last period for which we observe field choices. In Appendix F.1.2, we show that for any $t = T - \tau$, (14) can be

[22]The key challenge of estimating V^e directly is that we do not observe the final continuation value of V. Even differences across majors, $V^e(k,m) - V^e(k,m')$ cannot be directly observed either, as there is no simple representation of the difference that does not require knowledge of some terminal value.

iterated forward and rewritten using the following representation:

$$\bar{U}_{T-\tau,lm}(s) = \vartheta_{lm}(1-\beta^{T-\tau}) + \theta_l \underbrace{\left[\sum_{\tau'=\tau}^{T-1} \beta^{\tau'-\tau} V^F_{T-\tau'+5}(m,s)\right]}_{XV^F_{\tau,T}(m,s)} + \beta^{T-\tau}\bar{U}_{Tlm}(s), \tag{15}$$

where V^F is an entry value term that *only* depends on observable data.[23] For ease of notation, we collect these observable terms into a single term and denote it as $XV^F_{\tau,T}(m,s) \equiv \sum_{\tau'=\tau}^{T-1} \beta^{\tau'-\tau} V^F_{T-\tau'+5}(m,s)$. Notice that if $\tau = T$, this collapses to the final period utility. Otherwise, this is a discounted sum of observed payoffs, using the discounted final period utility, i.e., $\beta^{T-\tau}\bar{U}_{Tlm}(s)$, as a control for the continuation value.

With (15) in hand, we proceed to derive the likelihood function. To construct the likelihood conditional on latent preference type l, we define the relevant elements: let i index individuals, let $\mathcal{C}_i$ denote their choice set (determined by calendar time, GPA, and high school track), and let $(m_1, m_2, \ldots, m_R)$ represent their ranked choice list. To develop intuition for the likelihood function, we begin with an illustrative example of an individual i observed in the final period of available data, $T-6$. For this individual's *top* ranked choice, conditional on their type latent l_i and score being in decile s_i, the corresponding likelihood contribution is given by:

$$L_{i,T-6,1,l} = \frac{N^{\upsilon}_{m_{r_i}} \exp\left(U_{m_r l_i T-6}(s_i)\right)}{\sum_{m' \in \mathcal{C}_i} N^{\upsilon}_{m'} \exp\left(\bar{U}_{m' l_i T-5}(s_i)\right)}.$$

Once a choice has been made, the number of available options must be adjusted accordingly. To simplify the notation, let $N_{i,t,r,m}$ denote the number of options individual i has at time t, in position r within field m, given their prior choices. For instance, suppose there were initially 40 nursing programs available. If the individual selected a nursing program as their first priority, then 39 nursing programs remain available for their second priority. Using this notation, the likelihood contribution for the full ranked list is expressed as follows:

$$L_{i,T-6}(m_1, m_2, ..., m_R, s_i) = \prod_{r=1}^{R} \frac{N^{\upsilon}_{i,T-5,r,m_{r_i}} \exp\left(\bar{U}_{m_r l_i T-5}(s_i)\right)}{\sum_{m' \in \mathcal{C}_i} N^{\upsilon}_{i,T-5,,r,m'} \exp\left(\bar{U}_{m' l_i T-6}(s_i)\right)}.$$

Finally, by taking the logarithm of the likelihood expression and summing across all types and

[23] Specifically,

$$V^F_{t+5}(m,s) = \sum_k \pi(k|m,s)\left\{-(1-\beta)C^e_{okm} + w_{o,t+5}H(1,k,m) - \nu \log \lambda^e_{t+5}(o|k,m) - \beta\nu \log\left(\frac{\lambda_{t+6}(o|o,k,m)}{\lambda^e_{t+6}(o|k,m)}\right)\right\}.$$

individuals, we obtain the log-likelihood contribution for the final observation period:

$$LL_{T-6} = \sum_l \sum_i \sum_r q_{il} \left(\upsilon \log N_{i,r,m_{r_i}} + \bar{U}_{m_r l_i T-6}(s_i) - \log \left(\sum_{m' \in \mathcal{C}_i} N_{i,r,m'}^{\upsilon} \exp\left(\bar{U}_{m' l_i T-6}(s_i) \right) \right) \right).$$

A similar argument applies to earlier periods. Summing across all periods leads to the following log-likelihood conditional on the distribution of latent types:

$$\begin{aligned} LL = \sum_t \sum_l \sum_i \sum_r q_{il} \Bigg(& \upsilon \log N_{irm_{r_i}} + \vartheta_{lm}(1-\beta^{T-6-t}) + \theta_l X V^F_{t,T-6}(m,s) + \beta^{T-6-t} U_{T-6,lm}(s) - \\ & \log \Big(\sum_{m' \in \mathcal{C}_i} N_{irm'}^{\upsilon} \exp\Big(\vartheta_{lm}(1-\beta^{T-6-t}) + \theta_l X V^F_{t,T-6}(m,s) + \beta^{T-6-t} U_{T-6,lm}(s) \Big) \Big) \Bigg) \end{aligned} \tag{16}$$

The objective can be maximized by iterating on the first order condition, as in Berry et al. (1995). Details, including the exact expression from the FOCs, and their connection to various data moments, can be found in Appendix F.1.2.

In order to recover the distribution of latent preference types l, we once again rely on the EM algorithm. To update the probability that individual i of type l, conditional on a guess of parameters, we use Bayes' rule:

$$q_{itl}^{(g+1)} = \frac{\exp(LL_{itl}^{(g)}) q_l^{(g)}}{\sum_{l'} \exp(LL_{itl'}^{(g)}) q_{l'}^{(g)}}.$$

We iterate on updated values of q_{il} and the model parameters until the likelihood ceases to grow within some tolerance.

5.3 Identification

We discuss the intuition for identification by following the same sequence as the estimation procedure. First, the identification of wage parameters comes from the panel structure of the wage data, resembling a fixed-effects identification strategy in a panel dataset. If there were no occupations, the EM algorithm that groups workers into different labor market types would be tantamount to clustering workers by their long-run average quantile in the wage distribution.[24] When occupations are introduced, identification leverages wage patterns among occupational switchers and group them based on their revealed comparative advantages. Specifically, if a worker switches occupations

[24] More precisely, given our functional forms, we would be fitting a Gaussian mixture model to the data, and assuming that the Normal from which each worker draws is fixed at the beginning of their working life. This is similar in spirit to Bonhomme et al. (2019), and indeed the EM algorithm is often how k-means clustering is implemented in practice.

yet maintains a stable position in the income distribution, this suggests they possess an absolute advantage in both occupations. Conversely, if their wage rank declines (relative to peers in the new occupation) following the switch, this indicates they have a comparative advantage in their initial occupation. These scenarios are illustrated in Appendix Figure 12 for the case of two types.[25]

Second, the identification of switching cost parameters follows straightforwardly from the regression specification and relies on the relationship between occupational switching and changes in wage differentials. In particular, the response of *net* flows to changes in wage *differentials* identifies ν. For example, if flows from occupation o to o' remain unchanged when the wage gap $w_{o'} - w_o$ varies, this implies a small ν. On the other hand, the overall level of flows—*gross* flows—pins down the average level of switching costs, $C_{oo'km}$. Finally, conditional on the observed volume of switching, differences in C across occupation pairs are identified by variation in switching patterns relative to wage differentials. For example, if most switching occurs between o and o' rather than o and o'', despite similar wage gaps, this suggests that $C_{oo'}$ is smaller than $C_{oo''}$.

Lastly, while all parameters of the field choice parameters are jointly identified via likelihood maximization, useful intuition one can be drawn from the separate steps of the EM algorithm. First, given a distribution of types, the identification of θ_l relies on changes in field choices in response to shifts in returns to the observed component of that field, XV^F. Similarly, the fixed effects, ϑ_{lm}, capture long-term sorting patterns of students and how students of the same latent type make decisions when faced with different choice sets.

Next, the distribution of latent types is identified using information from the ranked list of choices. Conditional on a set of parameters, students with similar lists are grouped together. This becomes clearer with an illustrative example: Suppose there were no clustering and only one type of student. In this case, since preference shocks are conditionally independent, an individual's second choice would be independent of their first choice, i.e., $P(m_{r_2}|m_{r_1}) = P(m_{r_2})$. If, instead, there is dependence between first and second choices, students must be clustered in a way that ensures conditional independence given their type.[26] The ranked lists are hence *crucial* for identification. Without them, observing only a single choice per individual would make it impossible to separately identify the distribution of types, l, from the preference shock.

[25] This intuition is incomplete, however, as we also account for occupational histories to ensure that workers with similar career paths are also grouped together.

[26] See Greene and Hensher (2010) or Conlon et al. (2024) for more on the identification of unobservable preference parameters from ordered choice data.

6 Estimation Results

6.1 Labor Supply

In this section we focus on measuring the specificity of different fields in the labor market, while Appendix C contains all parameter estimates, as well as standard errors. We separate the discussion of the labor market block of the model into two sets of parameters: those governing human capital and those governing switching costs. Recall that human capital, as defined in (3), depends on both an individual's field of study and their labor market type k. In our estimation, we classify individuals into four types, i.e., $K = 4$.

To examine how human capital influences the specificity of a field, we construct, using α_{om}, a measure that captures the extent to which some fields are concentrated in a narrow set of occupations, while others distribute workers across many occupations. Maximizing (10) identifies α_{om} only up to differences, not in absolute levels. This is intuitive—identifying comparative advantage inherently requires defining both a base occupation and a base field as reference points. To capture the degree of field specificity, we borrow from the field of international trade and construct a measure of specificity based on the Revealed Comparative Advantage (RCA) measure of Balassa (1965). First we construct RCA:

$$RCA_{om} = \frac{\exp\left(\alpha_{om}/\nu\right) / \sum_{o'} \exp\left(\alpha_{o'm}/\nu\right)}{\frac{1}{M-1}\sum_{m'}\left[\exp\left(\alpha_{om'}/\nu\right) / \sum_{o'} \exp\left(\alpha_{o'm'}/\nu\right)\right]} - 1. \tag{17}$$

This measure represents the probability that a worker with field m enters occupation o, assuming only α_{om} determines occupational returns, relative to the average probability of entering that occupation across all other fields. For ease of interpretation, we subtract 1. A negative value indicates that workers from field m are less likely than average to enter occupation o, while a positive value suggests they are more likely. The magnitude further reflects the strength of this sorting relative to other fields.

Table 2 presents three summary statistics of RCA for each field.[27] The first column reports the maximum RCA value, which we consider our preferred measure of specificity. The second and third columns provide additional perspectives: the skewness of RCA_{om} and the slope of RCA_{om} on its rank within each field. These measures convey consistent insights. Fields such as Health, Education, and Care exhibit high specificity, indicating they are closely tied to a narrow set of occupations.

[27] Appendix Figure 11 plots the RCA values across occupations for each field.

Table 2
Specificity of Different Fields

	Max	Skew	β
High Health	0.78	1.96	2.30
Architecture	0.49	1.59	1.75
Education	0.45	0.96	1.60
Care	0.43	-0.11	1.73
Law	0.38	-0.00	1.79
Communication	0.34	0.84	1.23
Business Admin	0.31	-0.02	1.62
Physics/Math	0.29	-0.21	1.32
Life Sciences	0.28	0.34	1.24
ICT	0.27	0.32	1.34
Engineering	0.27	0.17	1.48
Medium Health	0.24	-0.34	1.60
Humanities	0.24	0.28	1.17
Arts	0.23	0.35	1.44
Social Science	0.21	-0.39	1.52

Notes: Displays various statistics (calculated across occupations) for each field's Revealed Comparative Advantage, as calculated in (17). *Max* is the maximum RCA difference across occupations; *Skew* is the skewness; β is the slope coefficient from a regression of RCA on its rank, $RCA_{om} = \beta_0 + \beta \times rank_o(RCA_{om})$.

STEM fields fall in the middle, while Humanistic and Social Science fields appear far more general, spreading their graduates across a wider range of occupations. Notably, for a few fields, the RCA skewness is negative. This implies that rather than being very well matched to a specific occupation, there might be fields that match very *poorly* to a few occupations, but are otherwise quite general skills.

Turning to the parameters governing switching, Table 3 displays estimates for this block of the model. The first row is the value of $1/\nu$ we obtain, scaled so that average wages are equal to 1. The second row is the value of C/w averaged across observed switches.[28] The first column is estimated by ignoring worker heterogeneity, as in Artuç et al. (2010), while the second is our model specification. Controlling for human capital dramatically raises the measured responsiveness of switching to wage differentials, and consequently leads to much smaller estimates of switching costs.

To give economic meaning to ν, we focus on the semi-elasticity of exiting occupation in response to a temporary shock to wages in that occupation:

$$\frac{d(1-\pi_{oo})}{d\log w_o} = -\frac{w_o}{\nu}\pi_{oo}(1-\pi_{oo}). \tag{18}$$

This semi-elasticity is the percentage point increase in exit from occupation o given a temporary

[28] For each individual, we take the q_{ik} weighted average across types.

Table 3
Labor Supply Parameters

Parameter	(1)	(2)
$\bar{w}/\nu$	0.249	1.67
	(.100)	(.160)
C/w	21.32	2.60

Notes: The first row presents the estimates for $1/\nu$ in (5), scaled by the mean wage. The second row presents the estimates for C/ν in (5). The first column reports the estimates using Artuç et al. (2010), while the second reports the main specification. In the first column, standard errors are clustered by year, ignoring first stage error in estimating average returns. In the second column, standard errors are calculated using a bootstrapping procedure that adjusts for both first stage errors and clusters by year (see Appendix F.4).

shock to wages. As workers of different fields differ in their exit rates, earnings, and sorting into occupations, we can calculate the semi-elasticity in (18) for different fields of study. More precisely, for each worker, we calculate the semi-elasticity given their current earnings, current occupation, and the probability they are type k. We define the mean semi-elasticity, $\mathcal{E}_m$, as the mean over all workers and years in a given field.[29]

Table 4 shows the value of $\mathcal{E}_m$ for each field of study. While correlated with our previous measure of specificity, there are clear differences. For example, Business Administration is highly elastic, but in the middle of the pack using our specificity measure. The differences arise as $\mathcal{E}_m$ depends not just on m but also on the correlation between m and labor market type, k, and other parameters of the model. A striking feature of Table 4 is the substantial heterogeneity in semi-elasticities across fields of study. For example, in the most responsive field—business administration—a 1% decline in wages raises the exit rate by nearly 6 times as in the least responsive field—medium health. The economy-wide average semi-elasticity is 0.13, meaning that a 1% temporary decline in earnings raises the exit rate by 0.13 percentage points. While this figure may seem small, it reflects the small magnitude of a 1% temporary shock. To provide further context, consider a permanent wage shock, which can be approximated by $w_o/(1-\beta)$.[30] In this case, a 1% permanent decline in earnings would

[29] Formally, let

$$\mathcal{E}_{imt} = \frac{1}{\nu} \sum_k q_{ik} \hat{w}_{imkt} \times \hat{\lambda}_{o(i)o(i)mkt} \times (1 - \hat{\lambda}_{o(i)o(i)mkt}).$$

Then we can define the mean semi-elasticity as,

$$\mathcal{E}_m = \frac{1}{NT} \sum_{i \in m} \sum_t \mathcal{E}_{imt}.$$

[30] Technically, the value function changes both directly through the wage decline and indirectly through ripple effects on switching probabilities. For simplicity, we focus only on the direct effect.

Table 4
Elasticity values by Education Field

Field of Study	$\mathcal{E}_m$
Business Admin	0.281
Social Science	0.254
Engineering	0.237
ICT	0.224
Physics/Math	0.200
Law	0.192
Life Sciences	0.187
Architecture/Construction	0.177
Arts	0.140
Humanities	0.137
Communications	0.136
High Health	0.084
Care	0.063
Education	0.052
Med Health	0.048

Notes: This table calculates the partial elasticity of occupational switching for each field of study. The partial elasticity of occupational switching is defined as the percentage point change in the exit rate from the most popular occupation in a given field in a response to a one percent decline in the earnings in that occupation.

raise the exit rate by approximately 1.7 percentage points, illustrating the more substantial impact of persistent wage changes.

6.2 Field Choice

In this subsection, we focus on three aspects of field choice estimation: the sensitivity of field choice with respect to net present value, heterogeneity in preferences by type, and substitution patterns between fields following a shock to net present values.[31] For our estimation, we set $L = 4$, which strikes a balance between computational efficiency and capturing meaningful heterogeneity in preferences and decision patterns.

Turning to the first set of results, Table 5 presents our estimates of θ for all types, along with their distribution. Additionally, we compute the mean semi-elasticity of students' first-choice field with respect to a shock. Specifically, for each student i, we calculate:

$$\mathcal{E}^f_{il} \equiv \frac{d\pi_{m(i)l}}{d\log V^e} = \pi_{im(i)l}(1 - \pi_{im(i)l}),$$

where $m(i)$ and $s(i)$ are the field choice and score (GPA) of student i, respectively. Since V^e is measured in differences rather than levels, we adopt a simple normalization to approximate its

[31] The complete set of parameter estimates is provided in Appendix C.

Table 5
Field Choice Parameters

Parameter	Type 1	Type 2	Type 3	Type 4
θ	0.030	0.144	0.052	0.031
	(0.001)	(0.022)	(0.004)	(0.001)
Q_l	0.183	0.197	0.301	0.320
$\mathcal{E}_m$	0.071	0.352	0.115	0.074

Notes: The table reports the results of field choice estimation by latent preference type l. The first row reports the estimates of θ_l in (13), with bootstrapped standard errors in parentheses below. The second row reports the share of each type in the population. The third row reports the mean semi-elasticity of students' first choice field with respect to a shock, as calculated in (19).

level: we assume the entry cost into Retail equals one-quarter of a year's earnings, reflecting the average time required to find a retail job.[32] Our measure of sensitivity is the weighted average of this semi-elasticity across students, where the weights reflect each student's probability of belonging to a given type:

$$\mathcal{E}_l^f = \frac{1}{\sum_i q_{il}} \sum_i q_{il} \times \mathcal{E}_{il}^f. \tag{19}$$

This measure represents the percentage point decrease in the share of students choosing a particular field in response to a 1% decline in that field's NPV.

Table 5 highlights two key findings. First, students appear relatively responsive to shocks. The type-weighted average of the final row— approximately 0.41—implies that a 10% decrease in the net present value of a field leads to a 4 percentage point decline in the probability of choosing that field. While this effect may seem modest, it is substantially larger than many estimates in the existing literature. For example, Beffy et al. (2012) estimates that a 10% increase in returns to STEM fields only increases entry by a quarter of a percentage point. Similarly, Wiswall and Zafar (2015) estimate that the elasticity of the log odds ratio of switching majors with respect to a shock to be 1.6, while the equivalent number in our estimation would be 2.5. Several factors may explain these differences. The Danish education system may provide students with better information about field-specific returns, enabling them to respond more effectively to economic incentives. Additionally, our analysis considers a larger set of fields, and as we will discuss below, much of the observed

[32] Our results are not particularly sensitive to this normalization. For example, we have also approximated V^e using the net present value of labor market returns—$\sum_o \pi_{om} \frac{w_o + \nu \log(\pi_{oo})}{1-\beta}$—and find similar results. This robustness arises because occupations with high entry costs are precisely those that students rarely enter (e.g., management). We have also used $w_o/(1-\beta)$ to assess the impact of an income shock while holding occupational choices fixed, and the results remain qualitatively unchanged. Appendix G provides details on our approximation of V^e, including how we take averages over time.

reallocation occurs within broad categories such as STEM.[33]

Our second key finding is the substantial heterogeneity in students' sensitivity to shocks. For instance, the most responsive students—identified as Type 2—are 18 times more sensitive to shocks than the least responsive group, identified as Type 3. This variation in sensitivity has important implications for counterfactual analyses: shocks that disproportionately affect more flexible types will have markedly different economic consequences, both in magnitude and in nature, compared to those that primarily impact less flexible types.

Beyond differences in sensitivity, we also find notable variation in preferences across types. Table 6 presents the distribution of first-priority choices by student type, with the final column showing the unconditional distribution in the population. The four types align well with common intuition and the priority matrix discussed in Section 2.3. Type 1 exhibits a strong preference for Education, Nursing, and Care fields. Type 2 corresponds to a "health care" profile, displaying a clear preference for High Health professions (e.g., doctors, dentistry) and Medium Health roles (e.g., nursing, technicians). Type 3 is a STEM-oriented profile. Interestingly, despite this STEM focus, a significant share of students in this group still select Business or Care as their first priority.[34] Lastly, Type 4 corresponds to a humanities/social science profile, with a clear preference for those fields.

Finally, before turning to counterfactuals, we first examine aggregate substitution patterns between fields of study. To do this, we calculate, for each field m, the distribution of students' choices following a shock to that field's NPV. For example, if the returns to Education increase by 1%, we observe the fraction of students who then choose Medium Health, Care, and other fields. In a standard multinomial choice model without heterogeneity, it is well known that these substitution patterns are independent of any parameters, and are instead proportional to initial conditional probabilities:

$$\frac{d\pi_{m'}/d\log V_m}{d(1-\pi_m)/d\log V_m} = -\pi_{m'}/(1-\pi_m).$$

[33]Interestingly, we do not find that the differences are driven by our use of ranked choice data, which is a different source of variation than in many other papers. If we only estimate one type, and ignore the ranked lists, we estimate similar mean semi-elasticities—but very different substitution patterns. On the other hand, we do find that controlling for choice set is very important, and assuming that all students have access to all fields of study lowers our estimates considerably. At the extreme of ignoring all observable heterogeneity, and simply inverting aggregate choice probabilities, our estimates are close to 0.

[34]There are two key reasons for this pattern. First, STEM is not a particularly popular choice overall, as reflected in the last column of Table 6. Consequently, even among students classified as Type 3, many still select fields like Business. Second, students within the same type do not necessarily face identical choice sets, either due to differences in their test scores or their high school background. As we will see below, examining substitution patterns between fields reveals a slightly different picture of student preferences.

Table 6
Distribution of First Priority by Student Type

	Type				Population
	1	2	3	4	
Education	50.7	0.8	1.4	1.8	10.4
Medium Health	7.7	57.7	0.9	0.6	13.2
Care	31.2	1.8	15.7	0.9	11.1
Arts	0.6	0.2	0.7	5.9	2.2
Humanities	2.0	0.4	1.0	21.0	7.4
Communication	1.8	0.7	9.9	13.7	7.8
Social Science	2.0	1.2	2.7	27.5	10.2
Business Admin	0.5	0.5	25.9	10.5	11.4
Law	0.4	0.6	4.7	7.5	4.0
Life Sciences	0.4	6.0	6.5	1.1	3.6
Physics/Math	0.2	0.6	4.7	1.1	1.9
ICT	0.0	0.0	3.4	1.5	1.5
Engineering	0.3	0.5	14.1	2.2	5.1
Architecture	0.4	0.2	6.0	3.2	2.9
High Health	1.9	28.7	2.5	1.5	7.2

Notes: Columns (1)-(4) report the share of individuals of each type (weighted by the probability of being that type, q_{il}) choosing each field as their first choice. The final column reports the population shares of each top priority.

However, with agent heterogeneity, this relationship may no longer hold. Instead, the substitution patterns are given by:

$$\frac{d\pi_{m'}/d\log V_m}{d(1-\pi_m)/d\log V_m} = -\frac{\sum_i \sum_l q_{il}\pi_{im'}\pi_{im}\theta_l V_{im}^e}{\sum_i \sum_l q_{il}(1-\pi_{im})\pi_{im}\theta_l V_{im}^e},$$

where the subscript i indexes individual student characteristics. Aside from heterogeneity in test scores and choice sets, students all differ due to their latent preference types. Table 7 presents the resulting substitution matrix, with the shocked field displayed in the columns and the alternative fields in the rows (hence, each column sums to 100%).

In line with the previous discussion, the substitution patterns reveal clear "clustering" of preferences. Two striking examples are the Education and Engineering columns. In the case of Education, nearly all substitution occurs toward Social Work and Health fields, with only a small share of students shifting to other fields. This is in line with the fact that most students opting for Education are Type 1 students, who also have strong preferences for Care and Medium Health. Importantly, these patterns could not be captured without our model of heterogeneous preferences. Similarly, following a decline in the NPV of Engineering, 25% of students who switch out move into Physics, Math, or Life Sciences—despite these fields accounting for only 5.5% of total students.

	Initial Field														
Substitute Field	Education	Medium Health	Care	Arts	Humanities	Communication	Social Science	Business Admin	Law	Life Sciences	Physics/Math	ICT	Engineering	Architecture	High Health
Education	.	7.53	60.13	2.63	3.58	2.87	3.04	1.61	1.76	0.78	0.72	0.79	0.86	1.11	1.08
Medium Health	12.11	.	16.30	1.40	0.92	1.32	1.22	1.11	1.39	19.51	1.23	0.47	0.53	0.65	47.82
Care	63.93	17.40	.	2.04	1.51	1.55	1.33	0.58	0.71	0.43	0.31	0.28	0.38	0.56	0.53
Arts	1.24	1.04	2.43	.	7.73	4.24	7.36	1.96	3.18	0.28	0.49	1.83	0.65	1.32	0.13
Humanities	5.58	2.27	4.43	26.17	.	15.14	35.76	4.28	13.47	0.92	1.73	1.86	1.29	2.75	1.41
Communication	4.48	3.97	5.47	18.27	21.36	.	16.53	24.37	13.92	3.68	7.39	5.02	7.34	6.45	1.21
Social Science	3.99	3.78	3.39	25.68	39.49	14.65	.	11.68	25.09	2.58	3.66	5.66	3.36	7.22	6.57
Business Admin	1.36	2.42	1.73	11.02	9.76	33.95	10.64	.	23.27	9.31	25.28	21.25	26.19	22.26	1.50
Law	0.76	1.82	0.91	6.02	8.69	7.97	12.80	12.19	.	2.61	3.22	5.05	5.18	6.38	2.19
Life Sciences	0.76	12.03	1.20	0.91	0.83	4.99	0.97	6.66	3.35	.	11.30	8.59	14.52	9.14	29.61
Physics/Math	0.40	1.17	0.63	0.58	0.92	2.23	1.42	7.09	2.01	5.00	.	6.49	11.46	6.28	4.05
ICT	0.10	0.11	0.21	1.47	0.71	1.35	0.70	5.43	1.41	1.59	4.92	.	9.03	7.69	0.22
Engineering	0.55	0.83	0.71	0.83	0.74	2.16	1.17	11.86	3.47	10.52	26.76	26.85	.	25.21	2.91
Architecture	0.62	0.63	0.82	1.73	1.89	2.93	3.71	7.47	3.52	3.53	7.69	13.70	15.47	.	0.76
High Health	4.12	45.00	1.65	1.25	1.85	4.66	3.35	3.72	3.46	39.25	5.30	2.15	3.75	2.98	.

Table 7
Substitution Matrix $\left(\frac{d\pi_{m'}/d\log V_m}{d(1-\pi_m)/d\log V_m}\right)$

Notes: Displays the estimated unconditional substitution patterns of students in each field who face a 1% shock to the NPV of that field. Columns are initial field, rows are substitute field. Each cell is the fraction of switchers moving from the initial field to a new field.

This suggests that the model effectively captures realistic clustering patterns.[35] Appendix Table 13 presents the substitution matrix without estimating any type heterogeneity, and using only students' first priority to recover preferences. We still allow for observable heterogeneity in test scores and in choice sets (arising from different backgrounds). In this case, we find unrealistic substitution patterns. Most strikingly, across all fields, nearly 10% of students who switch fields following a shock switch into Education. This contrasts dramatically with Table 7 where for most fields, this number is less than 2%, often less than 1%. This finding in particular highlights the importance of carefully modeling preference heterogeneity, and shows the value in our method.

The substitution patterns demonstrated above have important implications for our counterfactual analysis. While students respond to income shocks, their choices tend to remain concentrated within a narrow set of fields. Hence, even sizable shocks may lead to limited reallocation. In the next and final section, we use counterfactual simulations to explore our model's implications for both reallocation and welfare in response to realistic economic shocks.

7 Counterfactual Analysis

This section first outlines the remaining component of the model, which closes the model by specifying the labor demand side, and then describes the two counterfactual experiments. The first experiment assesses the effects of relaxing capacity constraints in higher education, examining how expended access influences labor market outcomes. The second evaluates the impact of simulated labor market shocks under different policy scenarios.

7.1 Closing the Model

We briefly describe the labor demand side of the model. Additional technical details are provided in Appendix H. In what follows, let $s \in \mathcal{S}$ index industries, and denote Denmark and Foreign by D and F, respectively.

The economy features a final non-traded good that agents consume. This final good is constructed as a CES aggregate of industrial outputs, combining both domestic and foreign varieties. At the

[35] Relatedly, notice that shocks to STEM fields (Physics, ICT, Engineering) do not lead to a large exodus to Care, despite the fact that these were clustered together by preference type. This reflects differences in the choice set between those students who could opt into STEM fields in the first place.

top level of this nested CES structure, the final good aggregator, Y_t, is defined as:

$$Y_t = \left(\sum_s \mu_s^{\frac{1}{\sigma}} Y_{st}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{20}$$

where Y_{st} represents industrial outputs, and μ_s is an exogenous demand shifter for industry s. Industrial outputs are further disaggregated into non-traded combinations of domestic and foreign varieties:

$$Y_{st} = \left((Y_{st}^D)^{\frac{\eta-1}{\eta}} + (Y_{st}^F)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \tag{21}$$

where Y_{st}^D and Y_{st}^F are domestic and foreign varieties, respectively, with corresponding price indices P_{st}^D and P_{st}^F. The elasticity parameters σ and η satisfy the standard condition $\eta > \sigma > 0$.

In addition to domestic demand, there is external demand for Danish exports, captured by the following demand curve:

$$X_{st} = A_s^F \left(P_{st}^D \right)^{-\eta}, \tag{22}$$

where A_s^F is an exogenous demand shifter for exports.

Domestic Production

Domestic production in each industry s is carried out by a large number of perfectly competitive production units. Each unit produces domestic output Q_{st} by combining production and non-production services. We classify occupations into two groups: production and non-production, denoted by $\mathcal{O}^P$ and $\mathcal{O}^N$, respectively. The production process follows the CES function:

$$Q_{st} = A_{st} \left[\left(A_s^N \right)^{\frac{1}{\alpha}} \left(L_{st}^N \right)^{\frac{\alpha-1}{\alpha}} + \left(\left(A_s^P \right)^{\frac{1}{\rho}} \left(L_{st}^P \right)^{\frac{\rho-1}{\rho}} + \left(A_s^M \right)^{\frac{1}{\rho}} \left(Q_{st}^M \right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}\frac{\alpha-1}{\alpha}} \right]^{\frac{\alpha}{\alpha-1}}, \tag{23}$$

where L_{st}^N denotes the total output of non-production occupation services, L_{st}^P represents the total output of production occupation services, and Q_{st}^M is the aggregated imported material inputs. The variables A_{st}, A_s^N, A_s^P, and A_s^M denote the sector-specific total factor productivity (TFP) and the productivities of non-production services, production services, and imported inputs, respectively. The parameter $\alpha > 0$ represents the elasticity of substitution between L_{st}^N and the production aggregate, while $\rho > 0$ denotes the elasticity of substitution between L_{st}^P and Q_{st}^M.

Both production and non-production inputs are modeled as Cobb-Douglas aggregates over

occupations. For $f = \{N, P\}$, we have:

$$L_{st}^{f} = \prod_{o \in \mathcal{O}^f} \left(\frac{L_{sot}^{f}}{\beta_{so}} \right)^{\beta_{so}}, \tag{24}$$

where $\sum_{o \in \mathcal{O}^f} \beta_{so} = 1$, and L_{sot} represents the efficiency units of human capital within each occupation and industry. Since workers are indifferent across sectors within an occupation, the total human capital in an occupation is given by: $L_{ot} = \sum_s L_{sot}$. Finally, consistent with the labor supply model, workers allocate their entire income within the period it is earned. Thus, aggregate expenditure equals aggregate income:

$$E_t = \sum_o L_{ot} W_{ot}.$$

To conduct our counterfactual analysis, we calibrate the full set of demand-side parameters directly from aggregate labor shares or rely on values derived from existing literature. Further details on this calibration process are provided in Appendix H.

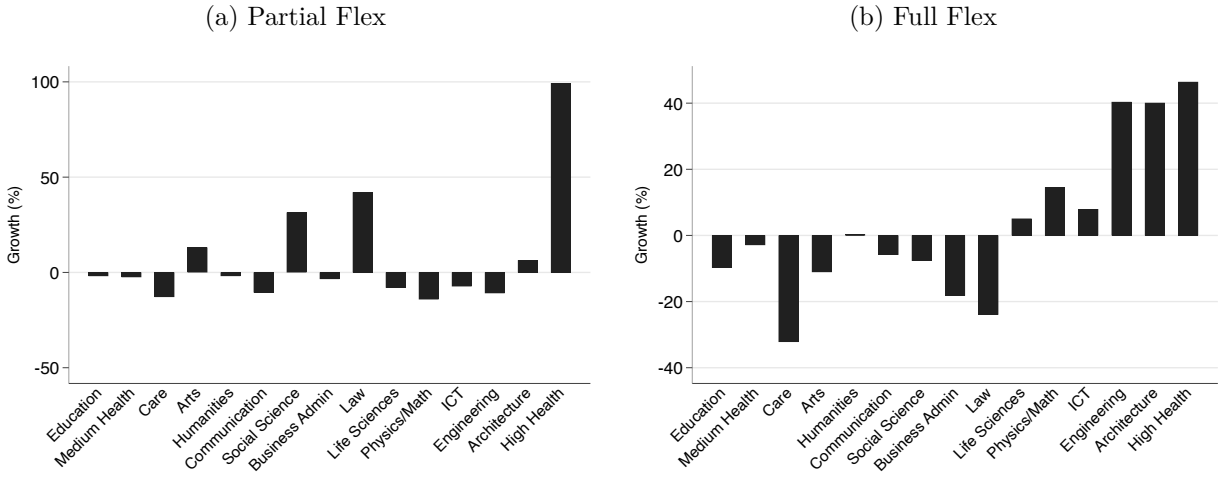
7.2 Capacity Constraints in General Equilibrium

In our first counterfactual, we examine how Danish capacity constraints and high school tracking affect worker allocation and GDP. We explore two scenarios. First, we remove capacity constraints but maintain high school tracking, allowing students to apply only to programs where the highest scoring peers in their track qualified ("Partial Flex"). Second, we remove both capacity constraints and tracking entirely ("Full Flex"). Both scenarios align with relevant policy changes—for instance, Denmark's expanded supplemental coursework (*Gymnasiale suppleringskurser*) to increase student flexibility.

Two caveats apply. First, we ignore the costs of increasing capacity. Second, we assume all students successfully complete their chosen programs, which may overstate graduation potential in certain majors—albeit this concern is partially alleviated by controlling for high school performance. In Appendix (XX), we offer one further scenario—we do maintain cutoffs and tracking specifically for medicine, a field that attracts many students but our model may be ill-suited for.

Figure 7 shows significant reallocation across majors in both cases. Under Partial Flex, 9% switch fields, compared to 21% in Full Flex. Appendix Table 14 reports conditional switching shares—the proportion of switchers who reallocate from column field m to row field m'. Switching aligns closely with student preference rankings and our estimated preference heterogeneity. Typically, reallocation

Figure 7
Field Reallocation Without Capacity Constraints



Notes: ...

occurs within similar fields: students in humanities and social sciences shift towards higher-paying fields (e.g., law), and STEM fields experience internal movement as well.

Policy makers designing educational systems face a potential tension between respecting student preferences and meeting national objectives. We are agnostic on the the source of these objectives, but there are a few possibilities. Strong student preferences may reflect incomplete information, possibly justifying a paternalistic approach. Alternatively, fiscal considerations may require policies that maximize national productivity to sustain education subsidies. Without endorsing a particular welfare function, we use our model to evaluate the economic and distributional consequences of removing constraints.

Table 8
GE Consequences of Lifting Capacity Constraints

	Partial Flex	Full Flex
%Δ GDP	-0.25	1.67
%Δ s.d. log Income	0.57	- 2.56
Field Reallocation	25.89	75.84
Occupational Reallocation	0.99	6.13

Notes: Table reports the changes for row variables between steady state after a policy simulation and the initial steady state. Occupational and field reallocation is measured as the average change in the absolute value of the growth rate of each occupation or field. I.e., for a discrete variable indexed by j, $Reallocation = \frac{1}{J}\sum_j |\log s_j^{SS_1} - \log s_j^{SS_0}|$, where s are shares.

Table 8 summarizes our findings. Removing capacity constraints alone (Partial Flex) slightly reduces GDP by 0.25%, whereas removing tracking and constraints (Full Flex) boosts GDP by 1.67%.

Neither change is dramatic, highlighting that current constraints have limited overall productivity impacts.

This muted GDP response is explained partly by limited occupational reallocation. To see this, we measure the average absolute percentage changes in occupation shares:

$$Reallocation = 100 \times \frac{1}{|\mathcal{O}|} \sum_{o \in \mathcal{O}} |\log(s_{o,cf}) - \log(s_{o,ss})|,$$

where $s_{o,j}$ is the share of college educated workers in occupation o under steady state ($j = ss$) and the counterfactual ($j = cf$). Occupational reallocation, shown in the third row of Table 8, is small relative to field reallocation. To examine distributional outcomes, we simulate the long-run incomes of a large cohort of 120,000 individuals.[36]

The Partial Flex scenario results in a slight increase in income inequality, whereas Full Flex significantly reduces inequality. To make sense of these results in partial equilibrium, note that, the net effect of relaxing capacity constraints depends largely on the correlation between existing capacity constraints and students' potential earnings. If constraints primarily restrict entry into lower-paying fields, relaxing them could increase inequality by permitting lower-scoring students into these less lucrative careers. This is especially true since students below the cutoffs have lower test scores, potentially exacerbating earnings disparities. On the other hand, if constraints limit access to higher-paying fields, easing these constraints could raise earnings for marginal students, thereby compressing the income distribution. We find no systematic correlation between the probability of being accepted to a program and earnings in steady state. This discussion abstracts from the general equilibrium forces at play, such as wage adjustments caused by reallocation that affect all workers.

Focusing on standard deviation may obscure richer distributional consequences. Figure 8 provides greater detail by plotting simulated individuals' changes in long-run income against their initial steady-state income under capacity constraints. The majority of individuals, clustered around the middle of the distribution, experience modest income growth positively correlated with initial earnings, consistent with the small inequality increase under Partial Flex. However, both tails show substantial convergence: earnings significantly increase for the lowest earners and decrease

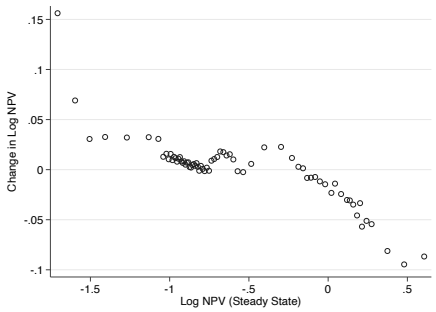
[36] We simulate individuals' field choices and earnings, computing each individual's discounted lifetime earnings:

$$NPV_i = (1 - \beta) \times \sum_{t=0}^{50} \beta^t w_{it}.$$

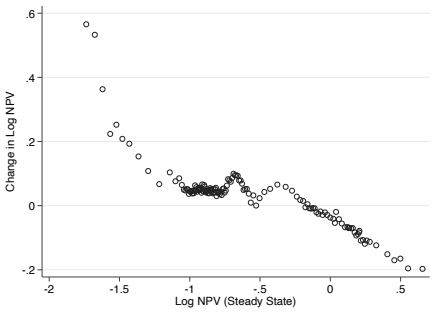
somewhat for top earners. This convergence effect is especially pronounced in the Full Flex scenario, dominating overall distributional outcomes.

Figure 8
Removing Capacity Constraints and Inequality

(a) Partial Flex



(b) Full Flex



Notes: For a simulation of $N = 100,000$ individuals, plots the change in steady state lifetime earnings, $(1-\beta)\frac{1}{50}\sum_{t=1}^{T}\beta^t w_{it}$, across steady states, against the steady state lifetime earnings in the original steady state.

In sum, relaxing capacity constraints improves students' welfare with limited consequences for aggregate output. Our analysis also uncovers an important interaction between high school tracking and tertiary education constraints: simultaneously removing both types of constraints boosts GDP, while only removing the latter does little. Nevertheless, caution is warranted; if tracking reflects inherent student abilities, our full flexibility scenario would be too blithe. Achieving this flexibility in practice may incur considerable costs.

7.3 Labor Market Shocks

This section examines two labor market shocks: an international trade shock that reduces the cost of imported goods and offshored inputs, and an AI-driven shock that alters demand for certain skills. We analyze their impacts on welfare, income, occupation and field-of-study reallocation, and intergenerational inequality across policy scenarios. The short run reflects effects on incumbent workers during the transition, while the long run represents steady-state outcomes.

International Trade Shocks

The international trade shock lowers imported input costs, reducing demand for domestic labor in occupations substitutable with offshored inputs. Starting from the projected 2018 steady state, we model a 20% price reduction, reflecting Denmark's observed decline in foreign prices over two

decades. This shock affects all sectors but disproportionately impacts occupations tied to offshored inputs.

As shown in Panel A of Table 9, in the short run, gains from trade for incumbent workers are substantial, with welfare increasing by 2.98% and income by 9.04% when there is no field reallocation. Allowing reallocation of new entrants in the transition raises these slightly to 3.01% in welfare and 9.47% in income.[37] The differences become more pronounced in the long-run steady state. A standard model with only occupational reallocation yields gains from trade of 3.47% in welfare and 13.28% in income. Allowing some field adjustments, as in the *Current System*, raises these to 3.94% and 23.12%. Even larger gains emerge under *full flexibility*, where gains from trade rise by 21% in welfare (to 4.20%) and 75% in income (to 23.23%). On inter-generational inequality, results reveal that increased flexibility leads to higher inter-generational inequality. Intuitively, incumbents, unable to change their earlier field choices, are disadvantaged relative to new entrants who can adapt to evolving opportunities. Specifically, under no field reallocation, inequality is 3.89%, but this rises to 12.5% under the *Current System* and 12.6% with *full flexibility*.

AI Shock

The AI shock is designed to capture shifts in relative labor demand across occupations. We use an external measure predicting declines in occupational demand and adjust sector-occupation productivity parameters to ensure wage shares remain unchanged. The model incorporates these changes through a two-stage process, first for non-production workers and then for production workers. For details on the full algorithm and calibration, see Appendix I.2.

As shown in Panel B of Table 9, short-run gains for incumbent workers are modest. Without field reallocation, welfare rises by 3.53% and income by 12.19%. Allowing some reallocation increases these slightly to 3.59% in welfare and 12.34% in income. In the long run, the gains are larger, especially with greater flexibility. With only occupational reallocation, welfare rises by 3.83% and income by 15.56%. Introducing some field reallocation, as in the *Current System*, raises these to 4.51% and 28.29%. Under *full flexibility*, gains reach 4.77% in welfare and 28.18% in income. On inter-generational inequality, the pattern mirrors the trade shock. Without field reallocation, inequality is 3.31%, rising to 14.4% under both the *Current System* and *full flexibility*, again reflecting the disadvantage faced by incumbents who are unable to adjust their earlier field choices.

[37] Since incumbents cannot switch fields, these gains are driven entirely by general equilibrium effects, such as wage adjustments and occupation switching.

Table 9
Results across Policy Scenarios: International Trade Shock vs. AI Shock

	Panel A: International Trade Shock			Panel B: AI Shock		
	No Field Reallocation (1)	Current System (2)	Fully Flex System (3)	No Field Reallocation (1)	Current System (2)	Fully Flex System (3)
Short Run						
%△ Welfare	2.98	3.01	3.01	3.53	3.59	3.59
%△ Income	9.04	9.47	9.47	12.19	12.34	12.33
Long Run						
%△ Welfare	3.47	3.94	4.20	3.83	4.51	4.77
%△ Income	13.28	23.12	23.23	15.56	28.29	28.18
Occupation Reallocation	5.26%	18.16%	21.6%	8.50%	22.99%	25.7%
Field Reallocation	0%	19.1%	24.1%	0%	21.8%	26.0%
Long Run vs. Short Run						
Inter-generational Inequality	3.89%	12.5%	12.6%	3.31%	14.4%	14.4%

Notes: Average percentage differences in welfare is defined as: $\frac{1}{N}\sum_{i=1}^{N}\frac{U_i^{CF}-U_i^{Base}}{U_i^{Base}}$, where U_i^{CF} and U_i^{Base} denote individual welfare under the counterfactual and baseline scenarios, respectively. The percentage change in lifetime income is calculated as: $\frac{1}{N}\sum_{i=1}^{N}\frac{NPV_i^{CF}-NPV_i^{Base}}{NPV_i^{Base}}$, where NPV_i^{CF} and NPV_i^{Base} denote the NPV of lifetime income for individual i under the counterfactual and baseline scenarios, respectively. Inter-generational inequality is defined as $\frac{1}{N}\sum_{i=1}^{N}\frac{NPV_i^{SteadyState}-NPV_i^{Incumbent}}{NPV_i^{Incumbent}}$, where $NPV_i^{Incumbent}$ is the NPV of incumbents in the transition and $NPV_i^{SteadyState}$ is the NPV of a new entrant in the new steady state.

Type-Specific Impacts

The results above underscore that the education choice margin plays a significant role in shaping the labor market's response to economic shocks. In particular, greater flexibility in educational pathways can amplify the gains from trade and AI-driven changes by improving the alignment between workforce skills and labor market demands.

When assessing the impact of labor market shocks, existing literature has generally focused on the heterogeneity created by the differential exposure of different fields to economic shocks. However, our findings also point to another important source of inequality: differences in students' inherent latent preference types, i.e., ν. These preferences govern individuals' willingness to switch fields in response to anticipated labor market changes, influencing how they are affected by economic shocks. In Table 10, we show that Type 2 individuals experience the largest improvement in lifetime income following both trade and AI shocks. This group is primarily composed of students in "Health care" fields, which are less exposed to trade and AI-driven disruptions. More importantly, they also possess the highest elasticity of field choice, ν, meaning they are particularly responsive to changes in expected labor market outcomes. This flexibility allows them to mitigate potential losses, driving their stronger income growth.

In contrast, Types 1, 3, and 4 exhibit smaller improvements in income. These groups share similarly low levels of elasticity, limiting their capacity to adjust in response to changing economic conditions. However, there are meaningful differences in their outcomes that reflect the nature of the shocks and the characteristics of their chosen fields. Type 1 individuals, who are concentrated in Education, remain relatively insulated from both trade and AI-related disruptions, contributing to their modest income gains. Meanwhile, Type 4 individuals, who are concentrated in Business and Management, experience slightly better outcomes than Type 3 individuals. Despite their similarly low elasticity, this advantage likely stems from the relatively stable nature of occupations that tend to be taken up by students with Business degrees, which are less exposed to disruption than the Social Science fields that characterize Type 3.

Table 10
Impact of International Trade and AI Shocks on Labor Types

Latent Preference Type	**Type 1**	**Type 2**	**Type 3**	**Type 4**
International Trade Shock	25.25	30.51	21.87	18.61
AI Shock	32.87	38.41	26.01	21.80

Note: The values are average percentage changes in lifetime income for each latent preference type l. See Table 5 and Section 6.2 for discussions on the differences of the types.

These results highlight that both the nature of economic shocks and the responsiveness of workers are crucial in understanding the distributional consequences of trade and technological changes. Moreover, removing institutional barriers that limit field switching also potentially plays a vital role in ensuring some groups are better positioned to thrive amid labor market disruptions.

8 Conclusion

In this paper, we model how students adjust their education decisions in response to globalization. Understanding how responsive students are to shocks is a key input to policymakers who wish to use education policy as a lever to address labor market disruptions such as trade, immigration, or automation. Our model captures selection on unobservable comparative advantage and preferences, while remaining tractable enough for estimation on standard administrative datasets.

We find that fields of study tend to be quite specialized, but there is substantial heterogeneity with some fields imparting quite general skills. Students cluster strongly on fields. However, within their clusters we find that students are more sensitive to shocks to future income than many past studies. We show that our estimation strategy captures observed rankings by students well, while

models that ignore latent heterogeneity do quite poorly.

We use our model to perform several counterfactual exercises. First, we find that relaxing capacity constraints on fields would lead to a huge amount of reallocation of students across fields. However, because much of this reallocation is within clusters, the induced occupational reallocation is small. As such, effects on output are small. When we look at shocks, we find that accounting for field choice leads to much larger gains from trade and technology shocks, but it also increases dispersion in outcomes, especially between incumbents and new entrants.

Our research leaves substantial room for future work. For example, in our model we treat one's high school background as fixed. It is plausible, indeed likely, that these decisions are also endogenous. We have also abstracted from institution quality, and other factors that might matter for students' decisions, such as their parents' backgrounds. Finally, the largest determinant of students' choices are embedded in the time-invariant component of preferences—essentially a residual. Opening up this black box—whether it reflects information frictions, secondary school inputs, or "deeper" preferences—could shed further light on what kinds of policy interventions in educational institutions will have impacts on labor market outcomes.

References

Adao, Rodrigo, Martin Beraja, and Nitya Pandalai-Nayar, "Fast and Slow Technological Transitions," *Journal of Political Economy Macroeconomics*, 2024, *2* (2), 183–227.

Adda, Jerome and Christian Dustmann, "Sources of wage growth," *Journal of Political Economy*, 2023, *131* (2), 456–503.

Agarwal, Nikhil and Paulo Somaini, "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism," *Econometrica*, 2018, *86* (2), 391–444.

Altonji, J.G., P. Arcidiacono, and A. Maurel, "Chapter 7 - The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects," in Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, eds., *Handbook of the Economics of Education*, Vol. 5, Elsevier, 2016, pp. 305–396.

Altonji, Joseph G., Erica Blom, and Costas Meghir, "Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers," *Annual Review of Economics*, 2012, *4* (1), 185–223.

_ , _ , and _ , "Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers," *Annual Review of Economics*, 2012, *4* (1), 185–223.

Andrews, Rodney J, Scott A Imberman, Michael F Lovenheim, and Kevin M Stange, "The returns to college major choice: Average and distributional effects, career trajectories, and earnings variability," Technical Report, National Bureau of Economic Research 2022.

Arcidiacono, Peter, "Ability sorting and the returns to college major," *Journal of econometrics*, 2004, *121* (1-2), 343–375.

_ and Robert A. Miller, "CCP Estimation of Dynamic Models with Unobserved Heterogeneity," *Econometrica*, November 2011, *79* (6), 1823–1868.

_ , Esteban Aucejo, Arnaud Maurel, and Tyler Ransom, "College Attrition and the Dynamics of Information Revelation," *Journal of Political Economy*, 2025, *133* (1), 53–110.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren, "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, June 2010, *100* (3), 1008–45.

Atalay, Enghin, "How Important Are Sectoral Shocks?," *American Economic Journal: Macroeconomics*, October 2017, *9* (4), 254–80.

Balassa, Bela, "Trade Liberalisation and "Revealed" Comparative Advantage," *The Manchester School*, 1965, *33* (2), 99–123.

Beffy, Magali, Denis Fougere, and Arnaud Maurel, "Choosing the field of study in postsecondary education: Do expected earnings matter?," *Review of Economics and Statistics*, 2012, *94* (1), 334–347.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–90.

Blanchard, Emily and William Olney, "Globalization and human capital investment: Export composition drives educational attainment," *Journal of International Economics*, 2017, *106* (C), 165–183.

Bleemer, Zachary and Aashish Mehta, "Will Studying Economics Make You Rich? A Regression Discontinuity Analysis of the Returns to College Major," *American Economic Journal: Applied Economics*, April 2022, *14* (2), 1–22.

_ and _ , "Will Studying Economics Make You Rich? A Regression Discontinuity Analysis of the

Returns to College Major," *American Economic Journal: Applied Economics*, April 2022, *14* (2), 1–22.

Bonhomme, Stephane, Thibaut Lamadon, and Elena Manresa, "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 2019, *87* (3), 699–739.

Broda, Christian and David E. Weinstein, "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, May 2006, *121* (2), 541–585.

Burstein, Ariel, Gordon Hanson, Lin Tian, and Jonathan Vogel, "Tradability and the Labor-Market Impact of Immigration: Theory and Evidence From the United States," *Econometrica*, 2020, *88* (3), 1071–1112.

Caucutt, Elizabeth M. and Lance Lochner, "Early and Late Human Capital Investments, Borrowing Constraints, and the Family," *Journal of Political Economy*, 2020, *128* (3), 1065–1147.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo, "Housing Booms and Busts, Labor Market Opportunities, and College Attendance," *American Economic Review*, October 2018, *108* (10), 2947–94.

Conlon, Chris, Julie Mortimer, and Paul Sarkis, "Estimating Preferences and Substitution Patterns from Second-Choice Data Alone," Technical Report, New York University 2024. Accessed: 2024-08-14.

Deming, David J and Kadeem Noray, "Earnings Dynamics, Changing Job Skills, and STEM Careers*," *The Quarterly Journal of Economics*, 06 2020, *135* (4), 1965–2005.

Dix-Carneiro, Rafael, "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 2014, *82* (3), 825–885.

Doraszelski, Ulrich and Jordi Jaumandreu, "Measuring the Bias of Technological Change," *Journal of Political Economy*, 2018, *126* (3), 1027–1084.

Dvorkin, Maximiliano and Alexander Monge-Naranjo, "Occupation Mobility, Human Capital and the Aggregate Consequences of Task-Biased Innovations," Working Papers 2019-13, Federal Reserve Bank of St. Louis 2019.

Fack, Gabrielle, Julien Grenet, and Yinghua He, "Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions," *American Economic Review*, April 2019, *109* (4), 1486–1529.

Feenstra, Robert C., "New Product Varieties and the Measurement of International Prices," *The American Economic Review*, 1994, *84* (1), 157–177.

Gemici, Ahu and Matthew Wiswall, "EVOLUTION OF GENDER DIFFERENCES IN POST-SECONDARY HUMAN CAPITAL INVESTMENTS: COLLEGE MAJORS," *International Economic Review*, 2014, *55* (1), 23–56.

Greene, William H and David A Hensher, *Modeling ordered choices: A primer*, Cambridge University Press, 2010.

Grigsby, John, "Skill Heterogeneity and Aggregate Labor Market Dynamics," Mimeo, Princeton University 2021.

Hastings, Justine S., Christopher A. Neilson, and Seth D. Zimmerman, "Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile," Working Paper 19241, National Bureau of Economic Research July 2013.

Heckman, James J., John Eric Humphries, and Gregory Veramendi, "Returns to Education: The Causal Effects of Education on Earnings, Health, and Smoking," *Journal of Political Economy*, 2018, *126* (S1), S197–S246.

Heckman, James J, Lance Lochner, and Christopher Taber, "Explaining Rising Wage

Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," *Review of Economic Dynamics*, 1998, *1* (1), 1–58.

Hotz, V. Joseph and Robert A. Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, July 1993, *60* (3), 497–529.

Humlum, Anders, "Robot Adoption and Labor Mafrket Dynamics," Working Paper, University of Chicago Booth School of Business 2021.

Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, June 2014, *104* (6), 1597–1629.

Kapor, Adam J., Christopher A. Neilson, and Seth D. Zimmerman, "Heterogeneous Beliefs and School Choice Mechanisms," *American Economic Review*, May 2020, *110* (5), 1274–1315.

Kapor, Adam, Mohit Karnani, and Christopher Neilson, "Aftermarket frictions and the cost of off-platform options in centralized assignment mechanisms," *Journal of Political Economy*, 2024, *132* (7), 2346–2395.

Keane, Michael and Kenneth Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy*, 1997, *105* (3), 473–522.

Khanna, Gaurav and Nicolas Morales, "The IT Boom and Other Unintended Consequences of Chasing the American Dream - Working Paper 460," Working Papers 460, Center for Global Development August 2017.

Kinsler, Josh and Ronni Pavan, "The Specificity of General Human Capital: Evidence from College Major Choice," *Journal of Labor Economics*, 2015, *33* (4), 933–972.

Kirkeboen, Lars J, Edwin Leuven, and Magne Mogstad, "Field of study, earnings, and self-selection," *The Quarterly Journal of Economics*, 2016, *131* (3), 1057–1111.

Larroucau, Tomas and Ignacio Rios, "Dynamic College Admissions," 2022. Working paper.

Leighton, Margaret and Jamin D. Speer, "Labor market returns to college major specificity," *European Economic Review*, 2020, *128*, 103489.

Lise, Jeremy and Fabien Postel-Vinay, "Multidimensional Skills, Sorting, and Human Capital Accumulation," *American Economic Review*, August 2020, *110* (8), 2328–76.

Llull, Joan, "Immigration, wages, and education: A labour market equilibrium structural model," *The Review of Economic Studies*, 2018, *85* (3), 1852–1896.

Luflade, Margaux, "The Value of Information in Centralized School Choice Systems," *Mimeo*, 2019.

Monras, Joan, "Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis," *Journal of Political Economy*, 2020, *128* (8), 3017–3089.

Moscarini, Giuseppe and Kaj Thomsson, "Occupational and Job Mobility in the US*," *The Scandinavian Journal of Economics*, 2007, *109* (4), 807–836.

Porzio, Tommaso, Federico Rossi, and Gabriella Santangelo, "The Human Side of Structural Transformation," *American Economic Review*, August 2022, *112* (8), 2774–2814.

Robles, Verónica C. Frisancho and Kala Krishna, "Affirmative Action in Higher Education in India: Targeting, Catch Up, and Mismatch," NBER Working Papers 17727, National Bureau of Economic Research, Inc January 2012.

Simonovska, Ina and Michael E Waugh, "The elasticity of trade: Estimates and evidence," *Journal of International Economics*, 2014, *92* (1), 34–50.

Traiberman, Sharon, "Occupations and Import Competition: Evidence from Denmark," *American*

Economic Review, December 2019, *109* (12), 4260–4301.

Wiswall, Matthew and Basit Zafar, "Determinants of college major choice: Identification using an information experiment," *The Review of Economic Studies*, 2015, *82* (2), 791–824.

Wooldridge, Jeffrey M, "Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity," *Journal of applied econometrics*, 2005, *20* (1), 39–54.

A Additional Tables

Table 11
Field of Study Classification

Field	Field Name	ISCED-F Code	College	Notes
1	Education	11	✓	
2	Arts	21	✓	
3	Humanities	22, 23	✓	
4	Social Science	31	✓	
5	Communication	32	✓	
6	Business Admin	41	✓	
7	Law	42	✓	
8	Life Sciences	51, 52	✓	
9	Physics/Math	53, 54	✓	
10	ICT	61	✓	
11	Engineering	71, 72	✓	Combined only for highest level of education
12	Manufacturing	72		
13	Architecture	73	✓	
14	Agriculture	80		
15	Medium Health	91	✓	Separated between university and professional bachelor's
16	Care	92	✓	
17	Personal Services	101, 102		
18	High Health	91	✓	Separated between university and professional bachelor's
19	Security	103		
20	Transport	104		

Notes: This table presents the authors' classification of fields of study, aggregated from the International Standard Classification of Education Fields of Education and Training (ISCED-F) codes. The ISCED-F is a framework developed by UNESCO to classify educational programs and related qualifications by fields of study, facilitating international comparability of education statistics and policies. Not all fields are part of the College (ISCED 6+) system, some are exclusively present in the short-cycle training (ISCED 5). Checkmarks denote the fields we call College.

Table 12
Occupation Classification

Occupation	Occupation Name	DISCO-08 2-Digit Codes	Notes
1	Management (Business + Admin)	10, 11, 12	
2	Management (Activity)	13, 14	
3	Engineering	21	
4	Medicine	22	
5	Education	23	
6	Finance, Admin, and Sales	24	
7	Software and ICT	25	
8	Social and Cultural Work	26	
9	Law	26	Split out of 3-digit code 261
10	Science Technicians	31	
11	Healthcare Assistance	32	
12	Admin	33, 34	
13	IT	35	
14	Secretarial	41	
15	Customer Services	42	
16	Accounting	43	
17	Other Admin	44	
18	Services	51, 94	
19	Retail	52	
20	Social Work	53	
21	Security	54	
22	Agriculture	60, 92	
23	Machinery	72	
24	Electrical	74	
25	Craftsperson	71, 73, 75	
26	Operators	81, 82	
27	Transportation	83	Includes 9621
28	Other Manual Work	91, 96	Excludes 9621
29	Construction	93	

Notes: This table presents the authors' classification of occupations, aggregated from the Danish version of the International Standard Classification of Occupations (DISCO-08). DISCO-08 is Denmark's national implementation of ISCO-08, a system developed by the International Labour Organization (ILO) to classify and compare occupations across countries based on the type of work performed and required skill levels. The classification enables standardized occupational data analysis in labor market research.

	Initial Field														
Substitute Field	Education	Medium Health	Care	Arts	Humanities	Communication	Social Science	Business Admin	Law	Life Sciences	Physics/Math	ICT	Engineering	Architecture	High Health
Education	.	15.15	19.35	10.47	11.97	13.24	9.57	11.49	8.91	8.60	7.13	8.07	9.03	8.83	7.50
Medium Health	18.53	.	24.78	15.69	13.98	15.37	11.65	15.44	11.71	10.71	9.11	11.79	10.89	11.37	9.49
Care	18.53	19.59	.	11.51	10.91	13.17	7.55	12.88	7.47	6.85	5.57	6.03	6.01	6.12	5.40
Arts	1.88	2.68	2.51	.	2.14	2.16	2.19	2.06	2.15	1.57	1.59	2.13	1.43	1.64	1.98
Humanities	9.73	9.52	9.50	8.93	.	10.71	10.23	7.31	8.33	7.13	7.57	3.47	5.24	4.85	8.90
Communication	9.96	9.21	9.33	8.65	11.30	.	9.24	8.99	8.35	5.28	5.23	3.04	4.09	3.74	6.57
Social Science	9.59	9.16	7.53	10.33	12.97	12.12	.	15.99	15.59	11.27	13.47	7.95	10.22	9.29	16.65
Business Admin	11.03	12.21	11.21	11.45	11.27	12.76	13.79	.	14.23	10.80	11.35	9.26	10.22	9.82	12.12
Law	3.53	3.82	3.17	5.06	4.51	4.37	6.28	6.13	.	4.41	4.38	5.17	5.27	4.98	5.64
Life Sciences	3.15	3.29	2.65	2.90	3.41	2.45	4.37	2.62	3.48	.	6.06	7.00	9.21	6.97	5.99
Physics/Math	1.25	1.35	0.94	1.14	1.59	1.01	2.76	1.91	2.02	3.59	.	2.69	3.90	2.99	4.24
ICT	0.67	1.09	0.70	1.16	0.70	0.58	0.99	1.19	1.03	1.90	1.56	.	3.45	3.30	1.53
Engineering	3.56	4.35	2.22	3.13	3.65	2.02	5.42	4.21	4.43	12.50	10.87	16.22	.	16.59	9.46
Architecture	2.17	2.63	1.73	2.14	2.10	1.54	2.91	2.59	2.49	5.57	4.81	8.80	9.84	.	4.52
High Health	6.41	5.96	4.37	7.44	9.50	8.51	13.05	7.18	9.82	9.82	11.30	8.37	11.19	9.51	.

Table 13
Substitution Matrix $\left(\frac{d\pi_{m'}/d\log V_m}{d(1-\pi_m)/d\log V_m}\right)$

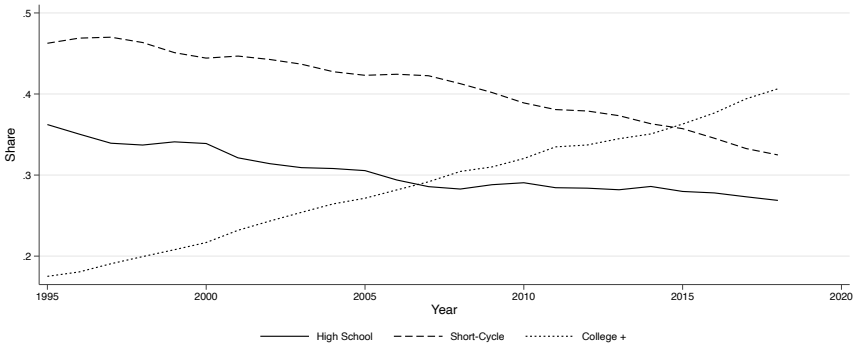
Notes: Displays the estimated substitution patterns of students in each field who face a 1% shock to the NPV of that field. These are calculated assuming *one* type of student. Columns are initial field, rows are substitute field. Each cell is the fraction of switchers moving from the initial field to a new field.

Table 14
Reallocation Table

	Initial Field														
Counterfactual Field	Education	Medium Health	Care	Arts	Humanities	Communication	Social Science	Business Admin	Law	Life Sciences	Physics/Math	ICT	Engineering	Architecture	High Health
Education			12.51				0.45	0.17		1.56		0.43			
Medium Health	12.73		4.01			0.75			8.18	0.17		1.10			
Care	2.06							0.09	0.68	0.30	0.35				
Arts	11.24	0.50	3.48		5.12	4.11	4.50	3.94	10.49	3.18	3.08	4.31	3.64	4.29	11.11
Humanities	7.30	1.00	1.64	1.62		8.22	1.35	5.44	10.19	4.64	6.32	5.35	7.39	3.43	
Communication	3.18	0.13	5.22	0.54	0.34		0.45	0.69	1.85	2.64	3.25	2.53	3.52	1.29	
Social Science	26.03	10.90	3.69	35.14	42.49	33.25		32.06	51.23	18.27	26.15	27.19	23.38	35.19	
Business Admin	2.62	0.38	10.18	2.70	2.90	4.73	1.80		8.33	5.82	7.35	5.50	8.78	4.72	38.89
Law	14.79	2.26	6.49	52.43	43.34	33.50	83.33	39.70		19.00	21.88	23.77	18.94	36.05	27.78
Life Sciences	2.25	3.63	4.96	0.54	0.51	1.00	0.45	1.04	4.01		5.30	4.75	6.52	3.00	
Physics/Math	0.19		1.21			0.12	1.35	0.58	1.23	0.64		0.30	0.69		
ICT	0.56	0.25	5.17	2.16	1.02	1.74	0.90	1.39	2.78	2.36	2.91		2.71	0.43	
Engineering	1.31		1.16	0.54	0.34	0.50	0.90		0.93	1.64	2.91	1.19			5.56
Architecture	7.49	0.63	4.96	1.08	1.37	3.11	3.15	2.43	4.32	4.64	5.64	5.05	6.64		16.67
High Health	8.24	80.33	35.30	3.24	2.56	8.97	1.80	12.73	4.63	28.45	14.19	19.76	14.78	11.16	

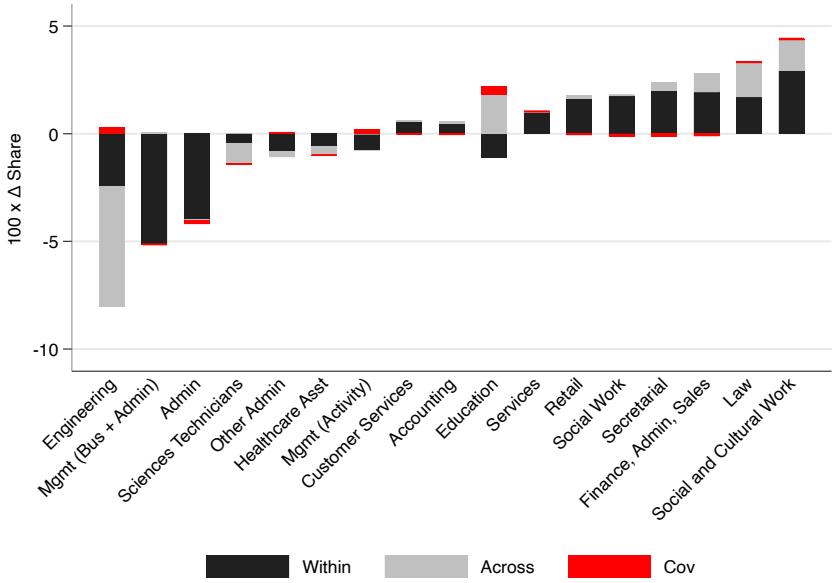
B Additional Figures

Figure 9
Education Level of New Cohorts at Labor Market Entry



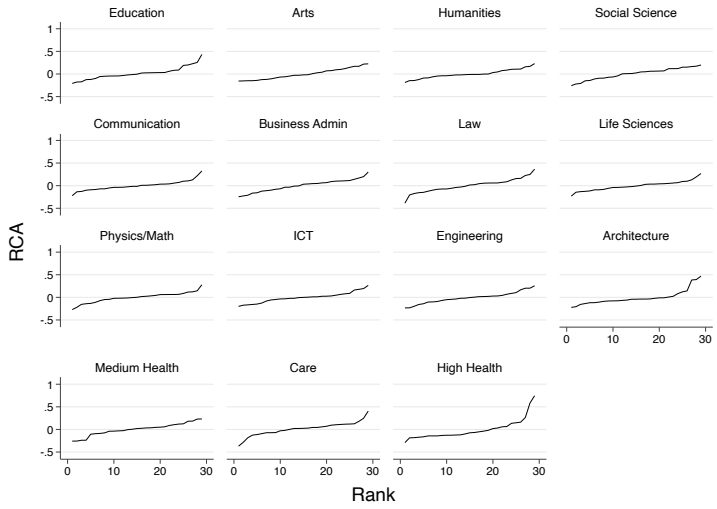
Notes: Plots the times series of the shares of highest attained level of education across workers in their first year on the labor market. Year of labor market entry is defined as the first year of employment after the worker achieves their highest level of education (i.e., work during schooling is not counted).

Figure 10
Decomposition for All Occupations



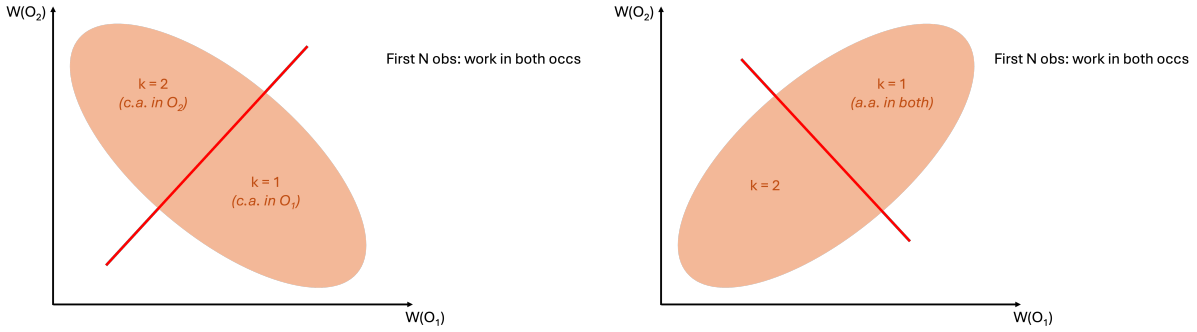
Notes:

Figure 11
RCA by Field



Notes: Plots RCA for each occupation as calculated in (17) against the rank of RCA across occupations. The slope of these lines is reported as β.

Figure 12
Intuition for the EM Algorithm



(a) Negative Correlation

(b) Positive Correlation

Notes: Panels (a) and (b) plot theoretical initial wages (x-axis) against next period wages (y-axis) for workers switching occupations.

C Complete Parameter Estimates

Table 15
α_k

Task	Type			
	1	2	3	4
0	0.000	-0.773	-0.251	0.132
	(0.000)	(0.020)	(0.014)	(0.010)
1	0.000	-0.040	-0.030	-0.037
	(0.000)	(0.002)	(0.001)	(0.001)
2	0.000	-0.026	-0.008	-0.006
	(0.000)	(0.002)	(0.000)	(0.000)
3	0.000	0.081	0.044	0.031
	(0.000)	(0.002)	(0.001)	(0.001)
4	0.000	0.023	-0.035	-0.017
	(0.000)	(0.002)	(0.001)	(0.001)

Notes: Displays the absolute advantage (task 0) and comparative advantage parameters (tasks 1-4) of differerent types. The first type is normalized in all parameters to 0.

Table 16

α_{om}

Occupation	Field														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	0.000	-0.130	0.004	0.226	0.018	0.268	0.145	0.138	0.211	0.070	0.257	0.115	-0.013	-0.182	0.243
	(0.000)	(0.021)	(0.014)	(0.010)	(0.011)	(0.008)	(0.011)	(0.009)	(0.014)	(0.016)	(0.007)	(0.011)	(0.013)	(0.011)	(0.013)
2	0.000	-0.303	-0.113	0.124	-0.114	0.131	0.106	-0.008	0.076	-0.023	0.130	-0.002	-0.284	-0.207	0.141
	(0.000)	(0.020)	(0.011)	(0.010)	(0.012)	(0.007)	(0.011)	(0.014)	(0.015)	(0.024)	(0.005)	(0.008)	(0.012)	(0.004)	(0.009)
3	0.000	-0.067	-0.000	0.192	0.051	0.248	0.147	0.155	0.241	0.125	0.280	0.157	0.044	-0.011	0.201
	(0.000)	(0.017)	(0.017)	(0.018)	(0.019)	(0.016)	(0.026)	(0.016)	(0.017)	(0.018)	(0.016)	(0.015)	(0.021)	(0.032)	(0.022)
4	0.000	0.082	0.061	0.258	0.112	0.260	0.159	0.332	0.364	0.209	0.434	0.216	0.102	-0.265	0.458
	(0.000)	(0.049)	(0.029)	(0.030)	(0.031)	(0.030)	(0.039)	(0.023)	(0.024)	(0.064)	(0.024)	(0.028)	(0.022)	(0.035)	(0.021)
5	0.000	-0.364	-0.108	-0.058	-0.118	-0.030	-0.057	-0.039	-0.020	-0.080	-0.010	-0.060	-0.223	-0.235	-0.036
	(0.000)	(0.006)	(0.004)	(0.005)	(0.004)	(0.005)	(0.011)	(0.004)	(0.004)	(0.012)	(0.006)	(0.007)	(0.007)	(0.002)	(0.005)
6	0.000	-0.162	-0.084	0.106	-0.073	0.154	0.113	0.044	0.094	-0.075	0.126	0.015	-0.026	-0.174	0.163
	(0.000)	(0.012)	(0.011)	(0.007)	(0.009)	(0.006)	(0.009)	(0.010)	(0.011)	(0.024)	(0.009)	(0.012)	(0.011)	(0.014)	(0.017)
7	0.000	-0.099	0.040	0.156	0.018	0.214	0.147	0.078	0.120	0.087	0.223	0.103	-0.008	-0.124	0.098
	(0.000)	(0.021)	(0.014)	(0.014)	(0.013)	(0.013)	(0.018)	(0.014)	(0.016)	(0.013)	(0.012)	(0.015)	(0.017)	(0.027)	(0.023)
8	0.000	-0.137	0.019	0.215	0.096	0.202	0.234	0.151	0.185	0.004	0.201	0.152	0.052	-0.081	0.091
	(0.000)	(0.014)	(0.014)	(0.011)	(0.012)	(0.014)	(0.014)	(0.015)	(0.015)	(0.024)	(0.019)	(0.022)	(0.022)	(0.012)	(0.019)
9	0.000	0.012	-0.017	0.190	0.029	0.175	0.236	0.022	0.120	-0.039	0.149	0.007	0.038	-0.023	0.072
	(0.000)	(0.058)	(0.032)	(0.028)	(0.028)	(0.030)	(0.029)	(0.031)	(0.031)	(0.036)	(0.031)	(0.029)	(0.051)	(0.037)	(0.035)
10	0.000	0.003	-0.029	0.177	0.026	0.221	0.192	0.145	0.187	0.145	0.261	0.140	-0.142	-0.065	0.136
	(0.000)	(0.020)	(0.025)	(0.032)	(0.020)	(0.017)	(0.040)	(0.019)	(0.022)	(0.026)	(0.017)	(0.018)	(0.017)	(0.027)	(0.047)
11	0.000	-0.069	0.112	0.224	0.155	0.236	0.262	0.231	0.236	0.279	0.343	0.225	-0.039	-0.053	0.261
	(0.000)	(0.061)	(0.049)	(0.047)	(0.054)	(0.042)	(0.046)	(0.043)	(0.050)	(0.071)	(0.044)	(0.050)	(0.043)	(0.062)	(0.052)
12	0.000	-0.074	0.044	0.258	0.067	0.279	0.245	0.124	0.221	0.102	0.280	0.157	0.069	-0.108	0.131
	(0.000)	(0.013)	(0.009)	(0.010)	(0.009)	(0.009)	(0.012)	(0.013)	(0.017)	(0.016)	(0.010)	(0.010)	(0.010)	(0.009)	(0.023)
13	0.000	-0.091	0.049	0.201	0.058	0.250	0.148	0.086	0.130	0.100	0.238	0.151	0.045	-0.028	0.097
	(0.000)	(0.020)	(0.024)	(0.021)	(0.018)	(0.018)	(0.048)	(0.025)	(0.027)	(0.026)	(0.018)	(0.025)	(0.027)	(0.032)	(0.039)
14	0.000	-0.025	0.070	0.212	0.162	0.316	0.274	0.144	0.177	0.162	0.289	0.193	0.073	-0.099	0.138
	(0.000)	(0.022)	(0.013)	(0.016)	(0.012)	(0.012)	(0.018)	(0.018)	(0.022)	(0.024)	(0.015)	(0.017)	(0.014)	(0.014)	(0.027)
15	0.000	-0.006	0.075	0.114	0.091	0.291	0.195	0.070	0.024	0.086	0.225	0.211	0.063	-0.055	0.066
	(0.000)	(0.042)	(0.027)	(0.034)	(0.026)	(0.025)	(0.046)	(0.035)	(0.077)	(0.068)	(0.032)	(0.043)	(0.028)	(0.036)	(0.053)
16	0.000	-0.006	0.101	0.193	0.066	0.267	0.283	0.126	0.167	0.132	0.231	0.101	0.107	-0.120	0.003
	(0.000)	(0.035)	(0.035)	(0.028)	(0.020)	(0.022)	(0.026)	(0.032)	(0.033)	(0.055)	(0.023)	(0.030)	(0.023)	(0.032)	(0.062)
17	0.000	-0.078	-0.018	0.134	0.105	0.233	0.179	0.057	0.065	0.089	0.265	0.167	-0.020	-0.042	0.049
	(0.000)	(0.026)	(0.025)	(0.024)	(0.019)	(0.018)	(0.030)	(0.033)	(0.041)	(0.043)	(0.019)	(0.022)	(0.021)	(0.022)	(0.043)
18	0.000	0.030	-0.011	0.132	0.119	0.263	0.166	0.069	0.110	0.157	0.227	0.345	0.165	-0.027	0.139
	(0.000)	(0.032)	(0.039)	(0.047)	(0.024)	(0.022)	(0.039)	(0.070)	(0.064)	(0.055)	(0.030)	(0.026)	(0.024)	(0.022)	(0.047)
19	0.000	-0.052	-0.003	0.228	0.161	0.388	0.192	0.217	0.241	0.264	0.393	0.241	0.177	-0.216	0.237
	(0.000)	(0.027)	(0.028)	(0.031)	(0.023)	(0.020)	(0.047)	(0.027)	(0.036)	(0.038)	(0.024)	(0.027)	(0.021)	(0.026)	(0.046)
20	0.000	-0.036	-0.079	-0.040	0.008	0.021	-0.036	-0.056	0.022	-0.078	0.022	0.058	0.061	0.007	0.060
	(0.000)	(0.016)	(0.019)	(0.021)	(0.022)	(0.017)	(0.043)	(0.033)	(0.045)	(0.062)	(0.030)	(0.028)	(0.015)	(0.013)	(0.029)
21	0.000	-0.257	-0.167	0.012	-0.014	0.113	0.078	0.056	0.067	0.054	0.036	0.153	0.051	-0.050	-0.007
	(0.000)	(0.046)	(0.053)	(0.053)	(0.042)	(0.033)	(0.060)	(0.047)	(0.065)	(0.101)	(0.047)	(0.043)	(0.040)	(0.036)	(0.049)
22	0.000	-0.049	0.217	0.278	0.017	0.170	0.267	0.282	0.264	0.167	0.304	0.308	-0.022	-0.010	0.350
	(0.000)	(0.063)	(0.055)	(0.052)	(0.043)	(0.051)	(0.119)	(0.041)	(0.065)	(0.123)	(0.038)	(0.038)	(0.040)	(0.042)	(0.049)
23	0.000	0.095	0.115	0.109	0.110	0.178	0.192	0.226	0.290	0.288	0.386	0.259	0.180	0.031	0.494
	(0.000)	(0.109)	(0.108)	(0.102)	(0.097)	(0.067)	(0.236)	(0.111)	(0.097)	(0.117)	(0.058)	(0.064)	(0.095)	(0.073)	(0.108)
24	0.000	0.059	0.149	0.147	0.244	0.114	0.338	0.015	0.025	0.177	0.178	0.089	0.144	-0.005	0.222
	(0.000)	(0.094)	(0.095)	(0.094)	(0.090)	(0.090)	(0.080)	(0.095)	(0.188)	(0.108)	(0.079)	(0.098)	(0.106)	(0.088)	(0.093)
25	0.000	0.083	0.120	0.062	0.111	0.179	0.160	0.216	0.224	0.042	0.347	0.167	0.139	-0.123	0.236
	(0.000)	(0.024)	(0.037)	(0.043)	(0.034)	(0.027)	(0.051)	(0.029)	(0.037)	(0.051)	(0.020)	(0.020)	(0.025)	(0.027)	(0.073)
26	0.000	0.111	0.148	0.129	0.172	0.162	0.124	0.086	0.193	0.022	0.362	0.213	0.113	0.024	0.173
	(0.000)	(0.038)	(0.044)	(0.050)	(0.032)	(0.026)	(0.056)	(0.044)	(0.045)	(0.090)	(0.023)	(0.030)	(0.032)	(0.032)	(0.065)
27	0.000	0.012	0.023	0.070	0.058	0.193	0.101	0.152	0.072	0.181	0.243	0.170	0.156	-0.026	0.096
	(0.000)	(0.042)	(0.041)	(0.039)	(0.035)	(0.027)	(0.053)	(0.045)	(0.066)	(0.071)	(0.027)	(0.028)	(0.040)	(0.030)	(0.076)
28	0.000	0.091	0.166	0.154	0.185	0.367	0.297	0.287	0.330	0.222	0.379	0.427	0.209	-0.054	0.232
	(0.000)	(0.032)	(0.032)	(0.042)	(0.028)	(0.028)	(0.056)	(0.042)	(0.053)	(0.073)	(0.028)	(0.026)	(0.032)	(0.025)	(0.072)
29	0.000	0.073	0.091	0.238	0.207	0.203	-0.070	0.108	0.135	0.096	0.289	0.374	0.068	-0.035	0.113
	(0.000)	(0.062)	(0.045)	(0.078)	(0.034)	(0.042)	(0.105)	(0.072)	(0.085)	(0.127)	(0.037)	(0.033)	(0.038)	(0.041)	(0.081)

Table 17
ϑ_{lm}

Field	Type			
	1	2	3	4
1	-0.500	0.073	0.908	-1.522
	(0.029)	(0.054)	(0.030)	(0.023)
2	-2.773	-2.463	1.340	0.988
	(0.020)	(0.009)	(0.018)	(0.036)
3	-4.536	-1.298	-0.243	0.904
	(0.034)	(0.034)	(0.019)	(0.026)
4	-3.564	0.116	0.263	1.531
	(0.037)	(0.040)	(0.025)	(0.026)
5	-5.163	-1.493	1.892	-0.717
	(0.028)	(0.040)	(0.035)	(0.030)
6	-6.318	-2.335	2.448	-0.087
	(0.020)	(0.017)	(0.040)	(0.023)
7	-4.190	0.591	4.852	3.421
	(0.020)	(0.057)	(0.043)	(0.042)
8	-4.086	1.984	4.779	0.734
	(0.017)	(0.109)	(0.061)	(0.033)
9	-4.440	-0.814	3.761	0.737
	(0.031)	(0.132)	(0.081)	(0.039)
10	-6.987	-3.427	4.167	0.936
	(0.004)	(0.005)	(0.031)	(0.035)
11	-4.596	-1.994	3.603	-0.105
	(0.010)	(0.027)	(0.036)	(0.026)
12	-4.162	-1.804	3.148	0.391
	(0.026)	(0.014)	(0.040)	(0.025)
13	-5.225	3.488	-0.894	-3.810
	(0.053)	(0.112)	(0.016)	(0.023)
14	-1.171	0.491	5.091	-3.424
	(0.034)	(0.079)	(0.021)	(0.029)
15	-4.827	3.711	3.178	-0.942
	(0.037)	(0.100)	(0.042)	(0.027)

Notes: Displays the type-field preference fixed effect.

Table 18
Additional Parameters

Parameter	Estimate
υ	1.363
	(0.008)

D Additional Counterfactuals

D.1 Maintaining Medicine Cutoffs

E Data Appendix

E.1 More on Denmark's Higher Education System and High School Tracks

Denmark's higher education system is structured to accommodate a diverse range of academic and professional aspirations. It includes two distinct pathways: Professional Bachelor's degrees and University Bachelor's degrees. While Professional Bachelor's programs focus on technical and marketable skills that cater to immediate employment, University Bachelor's programs are more academically rigorous and provide a foundation for further studies or research-based careers. This dual structure allows students to choose a path that best aligns with their goals, whether practical or theoretical.

Professional Bachelor's degrees typically last three years and include fields such as business, education, and engineering, alongside vocational training. These programs are characterized by their integration of hands-on learning and workplace training, often through internships, to ensure graduates are well-prepared for the labor market. University Bachelor's degrees, in contrast, follow a three-year curriculum rooted in theoretical frameworks and disciplinary knowledge, designed to prepare students for advanced academic pursuits or specialized professions.

The pathway to higher education begins with upper secondary schooling, which plays a critical role in determining students' higher education options. Danish upper secondary education consists of several tracks, including Stx (general academic), Hf (general and flexible), Hhx (business and economics), and Htx (technical and scientific). Each track prepares students for specific fields of higher education. For instance, Hhx emphasizes business-related disciplines, while Htx focuses on engineering and technology.

In 2008, reforms introduced greater flexibility to the upper secondary system, including supplemental summer courses (*Gymnasiale suppleringskurser*) to allow students to qualify for programs outside their original track. Despite these changes, the system retains a degree of rigidity, as the upper secondary specialization continues to influence the range of higher education programs accessible to students.

For those who complete University Bachelor's degrees, advancing to a Master's degree is a highly common choice, with the majority of cohorts pursuing this additional qualification. This continuation reflects the structured and sequential nature of Denmark's education system, which emphasizes both specialization and depth of knowledge. The typical trajectory for these students involves completing

a five-year combined program before entering the labor market in their mid-20s[38].

E.2 Classifying Field of Study

In this section we discuss the classification of fields of study and how we constructed our concordance. Stats DK provides a link between UDD codes and FAGOMRAADE codes, which can themselves be mapped to ISCED-15 and ISCED-F-27 *subject* codes. These concordances are publicly available and provided by Statistics Denmark. Below are the concordances:

- ISCED-F 2011: http://egracons.eu/sites/default/files/Isced%202013%20fields%20of%20education%20code%20list.pdf
- FAGOMRAADE-MELLEM: https://www.dst.dk/da/Statistik/dokumentation/nomenklaturer/disced15-udd#
  - Look at Subject Area variant.

UDD also provides a link to HOVEDOMRAADE, which can be mapped to ISCED codes *level* codes. We base our classification on how Denmark classifies its education programs in EURYDICE—the education research arm of the European Community.[39] This classification is below:

- HOVEDOMRAADE < 30 is level 1
- HOVEDOMRAADE $= 30, 40$ is level 2 (vocational + short-cycle) — ISCED 5
- HOVEDOMRAADE $= 50+$ is level 3 (bachelor's) — ISCED 6

Finally, due to data size constraints, we combine several ISCED-F components into aggregates. We also sub-divide ISCED-F components for health into those that are university based to those in the professional bachelor's system. This is because the former tend to be medicine/dentistry programs, and the latter nursing programs. Finally, we pool all defense/services sectors into vocational programs, even if offered at Professional Bachelor's universities. This affects a very small number of codes, and very few individuals. The exact changes we made are detailed below:

- Combine codes 51/52 (Biology + Environment)

[38] See https://eng.uvm.dk/upper-secondary-education/national-upper-secondary-education-programmes or https://www.uvm.dk/gymnasiale-uddannelser/uddannelser/gymnasial%20supplering/om-gymnasial-supplering for more background

[39] See https://eurydice.eacea.ec.europa.eu/national-education-systems/denmark/first-cycle-programmes.

- Combine codes 53/54 (Chem/Physics + Math/Stats)
- Split health into medium and high (based on 2d HOVEDOMRAADE)
- ** Combine all education and health into level 2
- Combine FAGOMRAADE 70/75/80 into level 2

E.2.1 Concording KOT to UDDA

KOT codes are codes that map to a particular program-institution-year combination. There is no direct concordance from KOT codes to fields of study. As such, we construct our own concordance. Our goal is to map KOT to UDDA and then to IDA, thus we construct a mapping that is based on the connection between the students who receive admission to a particular KOT code, and the UDD code that these students start according to UDDA. The UDD codes can be mapped to fields. We have supplemented this concordance with one that we did by hand using the names of the KOT codes, as a check against our procedure. There are a few things to note.

First, KOT codes within a span of time tend to be consistent. That is to say, the KOT code for “Psychology in Copenhagen University" will be consistent in each year. However, programs can change. If a program exits, KOT codes can be recycled. We generate synthetic KOT codes by hand based on the strings of the KOTs. These line-up with changes in region attached a KOT's institution, but can be more stringent. The reason we would like KOT codes to be consistent over time is because the field associated to a program can change if an administrator records a program. This can lead to spurious changes in the composition of fields. Thankfully this is a rare occurrence, affecting on the order of 10% of codes. Henceforth, a KOT code refers to a synthetic KOT code which we consider to be a *unique* and *time-consistent* program (field x institution).

Second, the same KOT code can map to different field codes over time. In this case we assign a time-consistent field to avoid spurious changes in the composition of field. When there is no time consistency in the field codes, we use the following set of rules to determine the case (where, below, code means an ISCED-F 4 digit field code):

1. There is a one time switch in the code, the first year of the program is before 2000, and there are at least 3 years where the initial program is observed
    - In this case we use the *initial* field. This is because we want to ensure backwards compatibility with people who have completed the program before 1996. It is rare that

the same program disappears for longer than 4 years, so that we assume if a program first appears in 2000 or later, it is a new program.

2. There is a one time switch in the code, the first year is before 2000, the *ISCED-F 2 Digit* code does not change, and there is less than 3 years that the initial program is observed
    - Assign the modal code.
3. There is a gap where a code switches and switches back
    - Assign the initial/final code
4. There is no change in the *ISCED-F 2 Digit* code and multiple switches
    - Assign modal code.
5. The initial year is 2000 or later (new program), there is a change in the program code that leads to a change in field, and it does not switch back
    - Assign modal code. Note that even if fields change, time consistency can be preferred because we update *everyone* who begins the program.
6. Remaining cases (of which there are less than 20 KOT codes out of ∼1400)
    - Decide case-by-case, and provide a note explaining the choice

After completing this concordance we have a unique synthetic KOT code that can be merged using true KOT-year combinations to a unique field of study. There are two remaining issues. First is that there are a small number of KOT codes that do not match to UDDA or a field directly. We assign these programs by hand based off of the text. Second, occasionally, there is a very poor match between the UDD and KOT. This occurs when students who gained admission to a KOT program end up assigned many different UDD codes. This can occur, for example, in cases where students do mixed-field-programs (e.g., “Chemistry and Physics"). For each KOT-year, we calculate the share of students in the modal UDD code for that KOT-year. If that is below 50%, we flag this and compare directly to the string in the text. We override the KOT-UDD match very rarely and note when and why.

There are two final issues.

First, throughout, language programs are often reclassified. This occurs, for example, when universities move language programs from business departments to communications departments.

It also happens within time-consistent KOT codes. Due to the tremendous variation in language code assignments, we hand-collected all language programs and assigned them to the most general super-set that many langugages were classified, which was Communications. Thus, we create a synthetic field code that essentially combines language programs with communications, and moves a handful of language programs that are not directly learning a foreign language (e.g., linguistics), and combine this with humanities.

Second, some KOT codes could not be matched to UDD. Obviously these codes will be for small programs, as it is implied that a very small number of people (<=5) matched to the modal UDD code in that KOT. For these KOT codes we use the text and impute the code by hand. On the codes that we did by hand, the match rate for ISCED-F 2 digit fields between the hand-code and the modal UDD match was 96% for non-language based fields. Thus, we think the hand-coding is not particularly subject to error. The final concordance contains a KOT code, year, synthetic KOT code, original UDD-based field, share of students in the KOT who have the UDD-based field, assigned-field, notes.

Merging Back to UDDA

With the concorded KOT codes, we can also map this back to UDDA. We find students who earned access to a particular KOT code. If their field of study begins with the dominant UDD code, or a UDD code that maps to a field in the same ISCED-F 1 digit grouping, we reassign them to the time-consistent, KOT-defined, field. Not including language codes, at the 4 digit level, this results in (possible) changes for 3.1% of code-years, while at the ISCED-F level, 2.0% of code-years are affected. In total, 1.36% of students—or about 13500 histories—need to be reassigned. Most individuals are not in KOT. This affects .11% of all person-years, and .65% of all person-years conditional on being college or more. We outline the procedure for updating below.

First, we define a variable for whether a student can be presumed to have started the field of study under their KOT. We do this based on three rules, all based on their first year of study in higher ed after the KOT admission year:

1. they are in a UDD code that has at least 5 people in it, and at least 10% of people start this UDD code
2. their UDD code is at least 50% of the total share of people
3. their UDD code is in the same 1-digit ISCED-F field as the field they will be reassigned to

After this, 92.5% of students are considered to have started their KOT code. We then use the following two rules to update their code to the concordance-assigned code:

1. The 1 digit ISCED field assigned by their UDD does not change over time. (using the 2 digit code is almost identical)

   (a) For example, if they start a program that is "natural science, n.e.s." according to UDD, and they are recorded as being in any natural science program *according* to UDD, we assume that their actual program has not changed and assign them the UDD based code.

2. If the UDD-based ISCED code switched to the KOT-Concordance based ISCED code at any point, we started using this UDD code going forward and reassign them to this code if it does not change. For example, if a certain program had been assigned as chemistry, and is reassigned to chemical engineering. Then we reassign people who are listed a chemistry to chemical engineering for as long as they are listed as having studied chemistry (and end their their history thereafter).

Ultimately, 69% of people have codes updated (recall that this is about a 1.36% subset of the full sample). Hence, a total of 63% of histories are updated. These people are considered to have started the program they applied for, finished it, and gone onto work using this field of study. What has changed is how we record the field of study relative to UDD. The remaining 37% are considered legitimate switches. That is to say, these are people who did not finish the degree they applied for. While this is much larger than the baseline reoptimization rate, it is worth pointing out that the overwhelming number of programs where there is a discrepancy are small programs—often interdisciplinary—where students are more likely to have been reshuffled. For individuals with an updated code, and no switch later, we consider them to have graduated with the updated KOT program. We merge them to the final UDDA file.

E.3 Defining Occupations

In order to define occupations we need to deal with three issues:

1. Between 2009 and 2010 there is a change in codes, as the ILO ISCO codes are changed to their 2010 version.

2. Many occupations are missing at various levels of digits (or have a “low quality" flag)

For example, there are codes that are 3000, even though there is no “30" 2-digit heading. Instead, this implies that only the first digit could be ascertained with confidence.

3. Defining a level of aggregation

In order to deal with the change in occupations, we rely on the concordance of Humlum (2021). This concords DISCO 6 digit codes across time. The difference between the ISCO's 4 digit occupations and the DISCO codes are in the last two digits, which are allowed to be more specific. Humlum's concordance is one-to-one. The match is imperfect, as not all occupations have the full 6-digit code. To create our occupations over time, we proceed in the following three major steps:

1. We use Humlum (2021)'s concordance. As there are multiple concordances and the number of digits in a code changes across years, we do so in the following:

    - First priority for a merge is given to exact 6-digit matches. We call these A merges. However, we do not allow switching to take place (or not take place) across ISCO-08 2 Digit groupings *if* there is not a concurrent ISCO-88 2 Digit grouping switch (or non-switch). That is to say, if there is a 2 digit change in the 88 codes, but not in the 08 codes, we do not accept the merge, if there is not a 2 digit change in the 88 codes, but there in the 08 codes, we do not accept the merge. We call this the switching-consistency-condition.
    - Second priority for a merge is given by matching the modal 6-digit string in a spell (connected set of plant-years). We call these C merges. Potential A and C matches agree 98% of the time at the 4 digit level. We again impose the switching-consistency-condition.
    - Final priority for a merge is given by matching 4 digit codes using Humlum (2021)'s 4 digit concordance. We call these B merges. Potential A and B merges agree at the 4 digit level 88% of the time, and at the 2 digit level 95% of the time. We again impose the switching-consistency-condition.
    - After the main merge is done, we allow for overriding of the switching-consistency-condition whenever there is a consistent 2-digit 08 code within a spell. There are not many violations of the condition in the first place, and this is a very conservative imputation that affects .4% of the observations.

After this merge is done, there is a consistent set of occupation codes over time. However, casual analysis of the data suggests significant spurious codes, especially around the break point 2009/2010.

To demonstrate this, we define a “Mass Event" as a 15% change in the absolute value of the share of an occupation in the economy year-on-year (using only high quality codes, as we do not use imputed codes later in the analysis). We also look at 10% changes in large occupations. The median change is on the order of 5% in a given year. Thus, such events ought to be rare and indeed they are, affecting some 5% of the sample.

Nevertheless, the events are unevenly distributed over time, affecting some 25% of the sample in 2009/2010 and 20% of the sample in each of 1997 and 2003, and 10% of the sample in 1999, while affecting less than 2% (and often less than .5%) of the sample in most other years. The reason for these breaks is two-fold: there is a change in the source of occupational codes over time. Some of this is dealt with by only retaining high quality occupation codes, and ignoring imputations. Because of this, a great many of the jumps—especially those before 2009/2010—are driven primarily by workers who had low quality codes or missing codes (e.g., only one digit code, like 500, or a 999) code, having their code attached. In 2009/2010, there are jumps for different several reasons:

- Occasionally it is obvious that several 08 codes were mistakenly given their 88 classification and the correction only occurred later. As an example, there is a spike in workers in nursing homes (code 513 in the 88 ISCO system) becoming bartenders (513 in the 08 ISCO system) and then returning (532 in the 08 ISCO system). There is documentation from Statistics Denmark suggesting that in cases where it was not clear which nomenclature was being used, as long as the 08 code existed, it was assigned as the occupation code. [40]

- Occasionally there is a true reclassification, likely stemming from gaps in the concordance (on account of trailing digits). These often “make sense" at first glance. For example, there is large reclassification of bank tellers into financial NES workers.

The issues are also concentrated in more occupational codes than others. In particular, ISCO codes that cover traditional “production" occupations (plant operation, machine work, etc.) suffer less from the problem because the occupation codes between 88 and 08 are almost the same for these occupations, making concording easier. The majority of issues are in occupation for higher skilled professions, as there were many changes in the coding for these occupations. For example, the creation of many IT and ICT occupations.

[40] Specifically see the document DISCO_koder_for_2010 available at https://www.dst.dk/da/TilSalg/Forskningsservice/Dokumentation/hoejkvalitetsvariable/personers-tilknytning-til-arbejdsmarkedet-set-over-hele-aaret--akm-/disco08-alle-indk.aspx.

Unfortunately, detecting these changes is hard. We documented all cases in the spread sheet XXX. We hand concord these cases. Since we are interested in occupational switching we take the following conservative stand: after visual inspection we back- or forward- fill occupation codes so that there is *no* switching occupation within a spell. This is similar to the assumption of Moscarini and Thomsson (2007).

Finally, after this is all done, we need to deal with low quality imputations, missing codes, or codes with only 1 leading digit. We impose the following rules:

1. asd

For codes that are not high quality, or affected by the above imputation scheme, we set the occupation codes to missing. In the estimation, we ignore missing data rather than integrate over it. This is not problematic (but perhaps inefficient) if missing data is random. In all of our choices we have aimed to be conservative, and assume less switching than actually occurs. Our final yearly switching rate at the 2 digit level is around 12%, reduced from 14%, which is in line with vom Lehn et al. The small change is small because, as previously discussed, only a handful of years are ever affected. To this end, the affect across years is more varied. After the procedure, the switching rate is more uniform across years—including 1997, 1999, and 2003. However, switching rates remain elevated between 2009 and 2010 . We do robustness by excluding this from the estimation of switching elasticities.

E.3.1 Tasks Construction

In addition to the Danish register data, we use the O*Net Database, which is maintained by the National Center for O*NET Development under sponsorship of the US Department of Labor. These data, which can be merged to the Danish occupation codes, are aggregated from survey responses pertaining to the importance of various tasks, skills, and types of knowledge to different occupations. (Add more here)

E.4 Constructing Education Histories

From UDG we can assign students a final exam score, as well as in which program they received the score. In our analysis we only consider students in the main 4 types of secondary program:

- Stx (split by speciality before 2008, mention grouping IB + Studentkursus here)

- Hhx
- Htx
- Hf

This covers XX% of KOT. XX% of applicants to KOT have a score that does not come from one of these programs. XX% are missing from the data. XX% are a small share covered primarily by the following programs: the Eux program (explain here—vocational thing) or prof. bachelors in teaching or nursing applying for another level of education.

In the KOT database, we know which type of exam is required by a program (`kot_eksvo` <– double check name). This exam highly correlates with the high school type assigned to students. We do not have information on conversion factors for students who apply from programs that do not match to the exam, so we treat their score as any other for the purposes of determining acceptance thresholds (described below).

E.5 Constructing Choice Sets

E.5.1 Determining Cutoffs

Students who apply to programs in Denmark can apply via Quota 1. This requires being allocated via one's exam score. In this case, we need to determine the set of schools that student's could *ex-post* have reasonably determined themselves eligible. As the published test score cutoffs have not been digitized (to the best of our knowledge), it is hard to determine the thresholds for admission. Moreover, even with this information, students can enter programs through alternative routes—for example through quota 2, standby, or other means. Thus we need a more consistent notion of cutoff rooted in the full set of pathways for admission to a program. We employ the following way to allocate a cutoff rule. We use this cutoff rule to then determine the number of programs in each field that each student could have entered. Specifically, we allow a 12/13 point (120/130 point scale) buffer.

1. Calculate the minimum score of a student admitted to a program with a quota 1 admission rule (otherwise assign a cutoff of 0) for each program, year, and grade scale (sometimes students overlap between the older and newer scales).
2. Calculate the median cutoff across years. If the a given year's cutoff is more than 20 points below the median, we assign the first percentile of admitted scores.

3. We update the median and recalculate. We then repeat replacing the first percentile with the 2.5th and the fifth.

4. If after this, there is a gap away from the median, we use the *largest* score of students who were applied on quota 1 and were rejected from the program (i.e., they recieved admission on a lower priority than the program in question). If the max is larger than the current cutoff, we use the 95th percentile.

5. For gaps far above the median cutoff, we use the score of the largest rejected applicant.

In general, there is a very high correlation between the minimum admit score and the largest reject score. The correlation between the minimum and the maximum is .62. This is high, but still reflects that there is a great deal of noise. This is especially the case in small programs, where one individual admitted via an alternative route (or perhaps misrecorded) can dramatically shift the cutoff. However, the correlation between the 5th percentile of admitted scores and 95th percentile of rejected scores has a correlation coefficient of .89, suggesting that this procedure is meaningfully capturing the difficulty of being admitted to some programs.

There is one additional issue: for a few years around 2008, there are students applying both under the old and new grading schema. Unfortunately, since we do not have the raw scores (only averages across classes), we cannot update the scales, and we do not know the conversion formula. However, for years when there is overlap, we take the cutoff for programs where people on both scales apply and regress one on the other. For example, we run,

$$\log GPA^{7}_{pt} = \beta_t + \beta \log GPA^{13}_{pt},$$

where GPA^{7}_{pt} is the cutoff under the 7 point system in year t for program p. These regressions in either direction produce R^2s over .7, suggesting a tight relationship between these cutoffs. We then impute the thresholds using this regression in years and programs where we do not observe people from one grade scale or the other applying. When this procedure is finished, we have a score cutoff for each GPA scale, for each program, in each year. There are also free programs, which have no threshold.

E.5.2 Determining Available Options

Students may not be able to apply to all programs despite a high score. This can occur because they lacked the requisite preparation. Before 2008, there was considerable tracking at the highschool

level in Denmark, with upper secondary programs being divided into a mathematics and linguistics track. Tracking remains in place afterward but was mildly loosened to encourage more flexibility. Nevertheless, programs often have minimum requirements. For example, to apply for Economics at the University of Copenhagen one *needs* to have taken A level mathematics. After 2008, the system was made to be more flexible. For example, summer programs have been expanded to allow students more flexibility. Nevertheless, rigidity remains. To determine which programs students can have entered we use their high school background. Before 2008 this is 5 different backgrounds—Hhx, Htx, Hf, Stx-Math, Stx-Linguistics—and after 2008 it is 4, as they combined the Stx programs. For each program, we look at all times this program was listed a priority for students from each background. We treat a program as eligible if it is allocated as a priority at least 5 times by students.

We know which exam students need to have taken to apply to a program, and this is *highly* correlated with educational background, so we believe that conditioning on high school background is an excellent proxy for programs which students can take. Of course, it is likely an overestimate given that we do not observe precise classes.

Once we have allocated the programs to which students with a background can apply for, we construct a choice set for each student based on a 12 (or 13, depending on scale), buffer around the calculated cutoff. Thus, for every year-high school-gpa-scale, we have a unique choice set that lists the programs available to each student within a field (including possibly 0). Our model implicitly assumes programs within a field are picked at random.

F Estimation Details

F.1 Labor Block Details

F.1.1 EM Algorithm Implementation Details

Inputs into the EM Algorithm are the following:

File Name	Description	Size
iDex	Key for *worker* for each i and t in panel	$NT \times 1$
tDex	Key for *calendar year* for each i and t in panel	$NT \times 1$
ptDex	Worker's observation within panel for each i and t	$NT \times 1$
field	Each worker's field	$N \times 1$
occupation	Each worker's occupation at time t	$NT \times 1$
logW	Each worker's log wage at time t	$NT \times 1$
zeta	Occupational characteristics	$\lvert\mathcal{O}\rvert \times Z$

These are the files read in by the Algorithm, however they are an incomplete list of what is needed for the Algorithm to run.

Step 1. Initialize the EM algorithm.

- Set types by k-means clustering on wages, ignoring occupations (as one possibility)
- Assign initial probabilities by,

$$q_{ik} = \begin{cases} 0.95 & \text{if clustered into } k \\ 0.05 & \text{if clustered into } k \end{cases}$$

- Construct the following sparse matrices *once*:
  - **XOT**: $NT \times |\mathcal{O}|$ matrix of occ-year dummies
  - **XOM**: $NT \times (M|\mathcal{O}| - M)$ matrix of occ-field dummies
  - For each k construct a matrix **XKZ** that is $NT \times (K-1) \times Z$ with all zeros for $k = 1$ (this is a normalization) and for $k \geq 2$:

$$\mathbf{XKZ}(k)(:, Z * (k-2) + Z) = \zeta_{o_{it},z}$$

  - Finally, construct the design matrix for each type:

$$\mathbf{XTEMP}(k) = [\mathbf{XOT}, \mathbf{XOM}, \mathbf{XKZ}(k)]$$

Note: In matlab this can be stored as a cell or multi-dimensional in array. In other languages, a multidimensional array. The notation here is Matlab's cell. This matrix does *not* change with iterations—only the weight matrix is updated.

- Estimate initial parameters for wages.
  (a) Preallocate the covariance matrix, **XtX**, to a $|\mathcal{O}| \times T + (M-1) \times |\mathcal{O}| + (K-1) \times Z$ square matrix of zeros.
  (b) Preallocate the covaraince matrix between Y and X, **XtY**, to a $|\mathcal{O}| \times T + (M-1) \times |\mathcal{O}| + (K-1) \times Z$ length vector.
  (c) For $k = 1, ..., K$...
    i. Generate the weight matrix,
    $$W = diag(q_{ik})$$
    ii. Update the RHS covariance matrix:
    $$\mathbf{XtX} = \mathbf{XtX} + \mathbf{XTEMP}(k)'W\mathbf{XTEMP}(k)/(NT)$$
    iii. Update the LHS covariance matrix
    $$\mathbf{XtY} = \mathbf{XtY} + \mathbf{XTEMP}(k)'W[\log w]/(NT)$$
  (d) Estimate the Mincer regression parameters,
  $$\beta^{Mincer} = (XtX)^{-1}XtY$$
  (e) Estimate residual variance
    – For each k,
    $$r_{itk} = \log w_{it} - XTEMP(k)_{it} \times \beta^{Mincer}$$
    – Estimate variance:
    $$\hat{\sigma}^2 = \frac{1}{NT}\sum_k \sum_i \sum_t r_{itk}^2 q_{ik}$$
- Estimate transitions:
$$\hat{\lambda}(o'|o,t,m,k) = \frac{\sum_{i:o_{it-1}=o \cap o_{it}=o' \cap i \in m} q_{ik}}{\sum_{i:o_{it-1}=o \cap i \in m} q_{ik}}$$

- Estimate initial conditions:

$$\hat{\lambda}^0(o|t,m,k) = \frac{\sum_{i:o_{it}=o \cap i \in m} q_{ik}\mathbf{1}(p_i^1 = t)}{\sum_{i:i\in m} q_{ik}\mathbf{1}(p_i^1 = t)},$$

  where p_i^1 is the first year the worker appears in the data.

Step 2. Initial value of EM objective, $J^{(0)} = 0$.

Begin E Step:

Step 3. For iteration g, update the EM objective, J^g. For each k calculate the log likelihood, LL_k

$$LL_k = \sum_{i=1}^{N} q_{ik} \left[\sum_{t=1}^{T} \left(-r_{itk}^2/(2\sigma^2) - \log(\sigma^2)/2 + \log \hat{\lambda}(o_{it}|o_{it-1}, m_i, k) \right) + \log(\hat{\lambda}^0(o_{io}|t_0, m_i, k) \right]$$

And define,

$$LL = \sum_k LL_k$$

Finally add cdf term:

$$J = \sum_k LL_k + \sum q_k \log(q_k),$$

where $q_k = \frac{1}{N}\sum_i q_{ik}$.

Step 4. Update weights,

$$q_{ik} = \frac{\exp(LL_{ik})q_k}{\sum_{k'} \exp(LL_{ik'})q_{k'}}$$

Begin M Step:

Step 5. Update unconditional weights,

$$q_k = \frac{1}{N}\sum_i q_{ik}$$

Step 6. Update Mincer Parameters

(a) For $k = 1, ..., K$...

i. Generate the weight matrix,

$$W = diag(q_{ik})$$

ii. Update the RHS covariance matrix:

$$\mathbf{XtX} = \mathbf{XtX} + \mathbf{XTEMP}(k)'W\mathbf{XTEMP}(k)/(NT)$$

iii. Update the LHS covariance matrix

$$\mathbf{XtY} = \mathbf{XtY} + \mathbf{XTEMP}(k)'W[\log w]/(NT)$$

(b) Estimate the Mincer regression parameters,

$$\beta^{Mincer} = (XtX)^{-1}XtY$$

(c) Update residual variance

- For each k,

$$r_{itk} = \log w_{it} - XTEMP(k)_{it} \times \beta^{Mincer}$$

- Estimate variance:

$$\hat{\sigma}^2 = \frac{1}{NT}\sum_k \sum_i \sum_t r_{itk}^2 q_{ik}$$

Step 7. Update transitions:

$$\hat{\lambda}(o'|o,t,m,k) = \frac{\sum_{i:o_{it-1}=o \cap o_{it}=o' \cap i \in m} q_{ik}}{\sum_{i:o_{it-1}=o \cap i \in m} q_{ik}}$$

Step 8. Update initial conditions:

$$\hat{\lambda}^0(o|t,m,k) = \frac{\sum_{i:o_{it}=o \cap i \in m} q_{ik}\mathbf{1}(p_i^1 = t)}{\sum_{i:i \in m} q_{ik}\mathbf{1}(p_i^1 = t)},$$

Step 9. Step 8: If $|J^g - J^{g-1}| < tol$, where g is iteration, stop. Else, go to Step 2.

Outputs to be recovered from β^{Mincer}:

File Name	Description	Size
$\hat{\mu}_{ot}$	Skill Prices	$\|\mathcal{O}\| \times T$
$\hat{\alpha}_{om}$	Occ-Field FE	$\|\mathcal{O}\| \times M$
$\hat{\alpha}_{kz}$	Type-Task Productivity	$K \times Z$

Additional outputs to save:

File Name	Description	Size
$\hat{\lambda}_{oo}$	Transition matrix estimate	$\|\mathcal{O}\| \times \|\mathcal{O}\| \times T \times M \times K$
$\hat{Q}_K$	Matrix of type weights	$NT \times K$

F.1.2 Stage 2 of Labor Supply

To derive (11),

$$\log\left[\frac{\lambda_t(o'|o,k,m)}{\lambda_t(o|o,k,m)}\right] = -\frac{1}{\nu}\Big(-C_{oo'km}+w_{o't}H_{o'}(k,m)-w_{ot}H_o(k,m)+\beta\big(V_{t+1}(k,m,o')-V_{t+1}(k,m,o)\big)\Big),$$

and

$$\begin{aligned}\log\left[\frac{\lambda_{t+1}(o'|o',k,m)}{\lambda_{t+1}(o'|o,k,m)}\right] =& -\frac{1}{\nu}\Bigg(C_{oo'km}+w_{o't+1}H_{o'}(k,m)-w_{o't+1}H_{o'}(k,m)\\ &+\beta(V_{t+2}(k,m,o')-V_{t+2}(k,m,o'))+(V_{t+1}(k,m,o')-V_{t+1}(k,m,o))\Bigg)\\ =& -\frac{C_{oo'km}+(V_{t+1}(k,m,o)-V_{t+1}(k,m,o'))}{\nu}\end{aligned}$$

Combining, we get,

$$\begin{aligned}\log\left[\frac{\lambda_t(o'|o,k,m)}{\lambda_t(o|o,k,m)}\right]+\beta\log\left[\frac{\lambda_{t+1}(o'|o't,k,m)}{\lambda_{t+1}(o'|o,k,m)}\right] =& -\frac{C_{oo'km}(1-\beta)}{\nu}+\\ &\frac{1}{\nu}\left(w_{o't}H_o(k,m)-w_{ot}H_o(k,m)\right)+\xi_{oo'kmt+1}\end{aligned},$$

In principle, to ensure positivity of the costs we apply an exponential transform. Specifically we estimate the following regression via unweighted NLLS:

$$\begin{aligned}\log\left(\frac{\hat\lambda_{oo'tmk}}{\hat\lambda_{ootmk}}\right)+\beta\log\left(\frac{\hat\lambda_{o'o't+1,mk}}{\hat\lambda_{oo't+1,mk}}\right) =& \exp\left(\gamma_{km}+\sum_z\gamma_z|\zeta_{o'z}-\zeta_{oz}|\right)\times(1-\beta)\tilde\nu-\\ &\tilde\nu\left(e^{\hat\mu_{o't}}\hat H(o',m,k)-e^{\hat\mu_{ot}}\hat H(o,m,k)\right)\end{aligned},$$

where $\tilde\nu=\nu^{-1}$ and $\hat H(o,m,k)=\exp(\alpha_{om}+\sum_z\alpha_{kz}\zeta_{zo})$.

F.1.3 Estimating Entry Costs

To derive (12),

$$\log\left[\frac{\lambda^e_t(o|k,m)}{\lambda^e_t(1|k,m)}\right] = -\frac{1}{\nu}\left(C^e_{okm}+w_{ot}H_o(k,m)-w_{1t}H_1(k,m)+\beta\Big(V_{t+1}(k,m,o)-V_{t+1}(k,m,1)\Big)\right)$$

Combining with $\log\left[\frac{\lambda_{t+1}(o'|o',k,m)}{\lambda_{t+1}(o'|o,k,m)}\right]$ and after some algebra, we obtain (12). In practice we take the mean across all occupations.

F.2 Details on Field Choice

F.2.1 Deriving the Recursive Representation

First we derive the transformation for the entry costs regression by expressing the entry value function in terms of future entry values and

$$\begin{aligned}
V_t^e(k,m) =& \nu \log\left(\sum_o \exp\left(\frac{-C_{okm}^e + w_{ot}H(o,k,m) + \beta V_{t+1}(o,k,m)}{\nu}\right)\right) \\
=& -C_{okm}^e + w_{ot}H(o,k,m) + \beta V_{t+1}(o,k,m) - \nu \log \lambda_t^e(o|k,m) \\
=& -C_{okm}^e + w_{ot}H(o,k,m) + \beta\left\{w_{o,t+1}H(o,k,m) + \beta V_{t+2}(o,k,m)\right\} - \nu \log \lambda_t^e(o|k,m) - \\
& \beta\nu \log \lambda_{t+1}(o|o,k,m) \\
=& -(1-\beta)C_{okm}^e + w_{ot}H(o,k,m) + \\
& \beta\left\{-C_{okm}^e + w_{o,t+1}H(o,k,m) + \beta V_{t+2}(o,k,m) - \nu \log \lambda_{t+1}^e(o|k,m)\right\} - \\
& \nu \log \lambda_t^e(o|k,m) - \beta\nu \log \frac{\lambda_{t+1}(o|o,k,m)}{\lambda_{t+1}^e(o|k,m)} \\
=& -(1-\beta)C_{okm}^e + w_{ot}H(o,k,m) + \beta V_{t+1}^e(k,m) - \nu \log \lambda_t^e(o,k,m) - \beta\nu \log \frac{\lambda_{t+1}(o|o,k,m)}{\lambda_{t+1}^e(o|k,m)}.
\end{aligned}$$

Group observable or estimated terms—wages, costs, and transition probabilties—into $V_t^F(k,m,s)$, where dependence on m and s is suppressed for convenience. Thus we have,

$$V_t^e(k,m) = V_t^F(k,m,s) + \beta V_{t+1}^e(k,m).$$

Averaging over k we have,

$$V_t^e(m,s) = \left[\sum_k \pi(k|m,s)V_t^F(k,m,s)\right] + \beta V_t^e(m,s)$$

Hence,

$$\begin{aligned}
\bar{U}_t(l,m,s) =& \vartheta_{lm} + \theta_l V_t^e(m,s) \\
=& (1-\beta)\vartheta_{lm} + \theta_l\left[\sum_k \pi(k|m,s)V_t^F(k,m,s)\right] + \beta\left[\vartheta_{lm} + \theta_l V_{t+1}^e(m,s)\right] \\
=& (1-\beta)\vartheta_{lm} + \theta_l\left[\sum_k \pi(k|m,s)V_t^F(k,m,s)\right] + \beta\bar{U}_{t+1}(l,m,s).
\end{aligned}$$

Iterating forward yields the expression for arbitrary t, and we fix $\bar{U}_T$ as the terminal value for all periods.

F.2.2 Likelihood Conditional on l

The FOCs are given by,

$$\frac{\partial LL}{\partial \vartheta_{lm}} = \sum_t (1-\beta^{T-6-t}) \sum_i q_{itl} \sum_r \left[\mathbf{1}(m_{r_i} = m) - \frac{N^{\upsilon}_{irm} \exp(U_{ml,t}(s_i))}{\sum_{m' \in \mathcal{C}_i} N^{\upsilon}_{irm'} \exp(U_{m'l,t}(s_i))} \right] \tag{25}$$

$$\frac{\partial LL}{\partial \theta_l} = \sum_t \sum_i q_{itl} \sum_r \left[XV^F_{t,T-6}(m_{r_i}, s_i) - \frac{\sum_{m \in \mathcal{C}_i} XV^F_{t,T-6}(m, s_i) \times N^{\upsilon}_{irm} \exp(U_{ml,t}(s_i))}{\sum_{m \in \mathcal{C}_i} N^{\upsilon}_{irm'} \exp(U_{ml,t}(s_i))} \right] \tag{26}$$

$$\frac{\partial LL}{\partial \upsilon} = \sum_t \sum_l \sum_i q_{itl} \sum_r \left[\log N_{irm} - \frac{\sum_{m \in \mathcal{C}_i} \log N_{irm} \times N^{\upsilon}_{irm} \exp(U_{ml,t}(s_i))}{\sum_{m \in \mathcal{C}_i} N^{\upsilon}_{irm'} \exp(U_{ml,t}(s_i))} \right]. \tag{27}$$

These first order conditions are intuitive.[41] The first says that ϑ_{lm} ensures that the time-averaged choice probabilities for each field by type l match the empirical shares. The second and third says that the sensitivity of individuals with respect to labor market fundamentals and the number of programs ensures that the model-implied average labor market outcome equals and log number of programs equals their empirical means. The discounting also has an intuitive appeal from a identification standpoint. The term $(1-\beta^{T-t-6})$ is larger the further away t is from the terminal period. Hence, the likelihood places more weight on earlier periods, where there is a longer history of observed discounted wage differentials, than on later periods. On a final note, the residuals in the $U_{lmt}(s)$ choice utilities will reflect measurement error and the expectational error that arise from using observed labor market outcomes in lieu of expectations.

F.2.3 EM Algorithm Implementation Details

Step 1. Initialize $\vartheta, \theta, \upsilon, q_{il}, U_T$.

- q_{il} initialized according to k-means clustering on test score and a matrix of priorities (0 for never chosen, 4 for first, 3 for second, etc.)
- U_T initialized by solving inner loop once with $\upsilon = .65$ (arbitrary intialization)
- Remaining parameters are estimated from type-specific multinomial regression on top priority.

Step 2. Initialize state payoffs:

$$U(m,l,s,t) = \vartheta_{lm}(1-\beta^{T-t}) + \theta_l XVF_t(m,s) + \beta^{T-t} U_T(m,l,s)$$

[41] In the implementation, to keep parameters bounded from 0, we parameterize them as $\theta = \exp(\vartheta)$ and search over ϑ. The first order conditions only change in that they are multiplied by the level of the current guess.

E Step:

Step 3. Calculate the current value of the complete data log likelihood:

$$LL = \sum_t \sum_i \sum_l q_{il} \left\{ \log q_l(s,f) + \sum_r \upsilon \log N_{m_{ri},r,i} + U(m_{ri},l,s_i,t) - \log \sum_{m' \in \mathcal{C}_i} N^{\upsilon}_{m',r,i} \exp\left(U(m',l,s_i,t)\right) \right\},$$

where $N_{m,r,i}$ is the number of programs in field m available to i at priority r.

Step 4. Update q_{il}:

$$q_{il} = \frac{LL_{il}}{LL_i}$$

M Step:

Step 5. Update the probability of being type l given scores, s, and gender, f, using sample mean:

$$q(l|s,f) = \frac{1}{N_{sf}} \sum_{i \in \{sf\}} q_{il}$$

Step 6. Inner Loop: Given the current value of U_t update model parameters.

1. Update U_T according to,

$$U'_T(m,l,s) = U_T(l,s) + \left[\sum_{t=1}^{T} \sum_{i \in \{s,t\}} q_{il} \beta^{T-t} \sum_{r=1}^{R} N_{i,r,m} \mathbf{1}(m_{i,r} = m) - \log \sum_{m' \in \mathcal{C}_\rangle, \nabla} N^{\upsilon}_{i,r,m'} \exp\left(U_t(m',l,s,t)\right) \right] / \sum_{i \in s} q_{il}.$$

2. Update remaining parameters using FOCs:

$$\begin{aligned}
\vartheta'_{lm} =& \vartheta_{lm} + \left[\sum_t (1-\beta^{T-6-t}) \sum_i q_{itl} \sum_r \left[\mathbf{1}(m_{r_i} = m) - \frac{N^{\upsilon}_{irm} \exp(U_{ml,t}(s_i))}{\sum_{m' \in \mathcal{C}_i} N^{\upsilon}_{irm'} \exp(U_{m'l,t}(s_i))} \right] \right] / \sum_i q_{il} \\
\theta'_l =& \theta_l + \left[\sum_t \sum_i q_{itl} \sum_r \left[XV^F_{t,T-6}(m_{r_i}, s_i) - \frac{\sum_{m \in \mathcal{C}_i} XV^F_{t,T-6}(m,s_i) \times N^{\upsilon}_{irm} \exp(U_{ml,t}(s_i))}{\sum_{m \in \mathcal{C}_i} N^{\upsilon}_{irm'} \exp(U_{ml,t}(s_i))} \right] \right] / \sum_i q_{il} \\
\upsilon' =& \upsilon + \left[\sum_t \sum_l \sum_i q_{itl} \sum_r \left[\log N_{irm} - \frac{\sum_{m \in \mathcal{C}_i} \log N_{irm} \times N^{\upsilon}_{irm} \exp(U_{ml,t}(s_i))}{\sum_{m \in \mathcal{C}_i} N^{\upsilon}_{irm'} \exp(U_{ml,t}(s_i))} \right] \right] / N.
\end{aligned}$$

3. Calculate the error as

$$error = ||(\upsilon'; \vartheta'; \theta'; U'_T) - (\upsilon; \vartheta; \theta; U_T)||_\infty$$

4. Terminate when the error is at tolerance, otherwise continue the loop.

Step 7. Terminate when $LL^{(g+1)} - LL^{(g)} < tol$. *Note: no absolute value is required as the EM algorithm should always increase. Failure to increase at an iteration suggests numerical error.*

F.3 Parameter Adjustments for Levels 1 + 2

After estimating the model for levels 1 and 2, we have estimates of μ_{ot}—in principle skill prices—for each group. Here we denote these by μ_{otg}, where g is for skill group. While it is the case that the average level of wages for each worker should be correctly matched because we normalize the *population* mean wages to 1, for each group we can only estimate a *relative* μ. That is to say, *something* must be normalized in the model. To see why, briefly, consider observing wages over time for different occupations and workers of different types (and ignore other heterogeneity), but condition on group. In this case, one has,

$$w_{otgk} = \mu_{otg} + \beta_{okg}. \tag{28}$$

This is essentially the decomposition we perform. This system is rank deficient. To see why is easy—one can add a constant to all the β and subtract this from all the μ and solve the system. To this end, we have imposed a normalization in estimation that sets $\beta_{11g} = 0 \forall g$. In reality, however, we do not want to do this because it is *not* the case that $\beta_{111} = \beta_{113}$. We need to renormalize. However, there is a further issue, which is that μ_{ot} are supposed to be [relative] skill prices, and we want these to match. I.e., we would like for $\mu_{ot1} = \mu_{ot2}$ for all groups.

We solve these issues with a simple renormalization:

$$\begin{aligned} w_{otgk} =& \mu_{otg} + \beta_{okg} \\ =& \mu_{ot3} + \underbrace{\left[\beta_{okg} + \mu_{otg} - \mu_{ot3}\right]}_{\tilde{\beta}_{okg}}. \end{aligned}$$

Another way to put this is that we want things measured in the same units. This normalization ensures that.

Unfortunately, this leads to time-varying β's. In order to avoid this, we only do the normalization for the year that we initialize from for counterfactuals. And so in practice we set,

$$\tilde{\alpha}_{omg} = \alpha_{omg} + \mu_{o0g} - \mu_{o03}$$

for all skill groups.

F.4 Bootstrapping

We calculate standard errors using the Bayesian block bootstrap, drawing weights for different stages. For the first stage of the labor block, the EM algorithm, by arguments in Arcidiacono and Miller (2011), is $\sqrt{N}$ consistent, so long as conditional on year, individuals remaining errors are independent. While Arcidiacono and Miller (2011) provide analytical covariance matrices, we use a bootstrap procedure so that in the second stage, where we bootstrap instead over year, we can capture first stage uncertainty. Specifically in the first stage we employ a Bayesian block bootstrap. For each bootstrap iteration, we draw from a Dirichlet distribution of length N with $\alpha_j = 1 \forall j$. We multiply this by N. This will be our vector of weights, w_i. Notice that this samples uniformly over individuals regardless of their total length in the sample, so that the total size of the sample (which we often abusively write as NT can change). Given the weights, we re-estimate the first stage where the likelihood is now weighted by w_i.

In the first stage, we estimate $\hat{\pi}$, which are the *true* transition rates. However, these transitions rates are based on *expectations* about the future. Thus, they encode information about $E_t V_{t+1}$. These expectations are not separately identified from other parameters, but as described in the main text, we plug in *future* transition rates and exploit finite dependence to estimate structural parameters in a linear regression. In *this* regression, that expectational errors may be correlated across units—now the occupation—within a year. Ignoring this clustering and re-estimating this stage using the bootstrapped values from the first stage would only capture uncertainty stemming from sampling error, but miss the within-year correlation. To this end, in this stage we also cluster by year. We now block bootstrap in the time dimension. Specifically, note that we may write the regression as depending on three periods $(t-1, t, t+1)$. We draw weights from a $T-2$ Dirichlet for the parameters in this estimation.

We employ the same Bayesian bootstrapping procedure (with newly drawn weights) for the Field Choice estimation. As a final caveat, we employ a minimum burn-in for the EM algorithm. We also employ slightly looser tolerances on interior optimizations for speed reasons. This leads to mild differences in where the bootstrap tends to settle. Nevertheless, the center is near the parameter estimates.

G Results Appendix

Approximating V^e

Given the Gumbel shocks, we derive above the following expression:

$$V_t^e(k,m) = -(1-\beta)C_{okm}^e + w_{ot}H(o,k,m) + \beta V_{t+1}^e(k,m) - \nu \log \lambda_t^e(o,k,m) - \beta\nu \log \frac{\lambda_{t+1}(o|o,k,m)}{\lambda_{t+1}^e(o|k,m)},$$

where λ are choice probabilities, and o is any choice of base occupation. Here we begin our approximation. Assuming that $V_t^e \approx V_{t+1}^e$ (i.e, that we are in a stationary environment),

$$V_t^e(k,m) \approx \frac{-(1-\beta)C_{okm}^e + w_{ot}H(o,k,m) - \nu \log \lambda_t^e(o,k,m) - \beta\nu \log \frac{\lambda_{t+1}(o|o,k,m)}{\lambda_{t+1}^e(o|k,m)}}{1-\beta}.$$

The issue with the approximation is that we cannot estimate C_{okm}^e in levels, only in differences. To this end, we need to take stand on this value. We assume that for retail, $C_{okm}^e \approx .25$, which is one quarter of the average annual income in the economy (which has been normalized to 1). Given this normalization, we can calculate the above for every value of o and every period t in the economy, and take the average. And so, calling the right hand side, $\tilde{V}_{ot}^e(k,m)$, where the $\tilde{V}$ refers to the approximated value evaluated at o, we defne,

$$V^{e,approx}(k,m) = \frac{1}{|\mathcal{O}| \times T} \sum_o \sum_t \tilde{V}_{ot}^e(k,m).$$

We then have,

$$V^{e,approx}(m,s) = \sum_k \pi(k|m,s) V^{e,approx}(k,m).$$

H Labor Demand Model

H.1 Model details

Given wages and prices of imported inputs, the profit maximizing choice of input quantity and the zero profit condition yield demand for occupation o in industry s:

$$H_{sot}^{P} = A_{st} \times \frac{\beta_{so}^{P}}{W_{ot}} \times A_{s}^{P} \times Q_{st} \times \left(P_{st}^{D}\right)^{\alpha} \times \left(P_{st}^{P}\right)^{\rho-\alpha} \times \left(W_{st}^{P}\right)^{-\rho}, \quad \forall o \in \mathcal{O}^{P}; \tag{29}$$

$$H_{sot}^{N} = A_{st} \times \frac{\beta_{so}^{N}}{W_{ot}} \times A_{s}^{N} \times Q_{st} \times \left(P_{st}^{D}\right)^{\alpha} \times \left(W_{st}^{N}\right)^{-\alpha}, \quad \forall o \in \mathcal{O}^{N}; \tag{30}$$

and demand for imported goods is

$$Q_{st}^{M} = A_{st} \times A_{s}^{M} \times Q_{st} \times \left(P_{st}^{D}\right)^{\alpha} \times \left(P_{st}^{P}\right)^{\rho-\alpha} \times \left(P_{st}^{M}\right)^{-\rho}, \tag{31}$$

where

$$Q_{st} = A_{s}^{F}\left(P_{st}^{D}\right)^{-\eta} + \mu_{s} E_{t} \times \left(\frac{P_{st}}{P_{t}}\right)^{1-\sigma} \times \frac{\left(P_{st}^{D}\right)^{-\eta}}{\left(P_{st}^{D}\right)^{1-\eta} + \left(P_{st}^{F}\right)^{1-\eta}}; \tag{32}$$

$$W_{st}^{f} = \prod_{o \in \mathcal{O}^{f}} w_{ot}^{\beta_{so}}, \quad f = \{P, N\}; \tag{33}$$

$$P_{st}^{P} = \left[A_{s}^{P}\left(W_{st}^{P}\right)^{1-\rho} + A_{s}^{M}\left(P_{st}^{M}\right)^{1-\rho}\right]^{\frac{1}{1-\rho}}; \tag{34}$$

$$P_{st}^{D} = \left[A_{s}^{N}\left(W_{st}^{N}\right)^{1-\alpha} + \left(P_{st}^{P}\right)^{1-\alpha}\right]^{\frac{1}{1-\alpha}} \tag{35}$$

$$P_{st} = A_{st}^{-1}\left[\left(P_{st}^{D}\right)^{1-\eta} + \left(P_{st}^{F}\right)^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{36}$$

$$P_{t} = \left[\sum_{s} \mu_{s} \left(P_{st}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{37}$$

$$E_{t} = \sum_{o} w_{ot}\left(\sum_{s} L_{sot}\right) \tag{38}$$

Finally, occupation wages, w_{ot}, adjust to clear the labor market, for $k = \{P, N\}$

$$H_{ot} = \sum_{s} H_{sot}^{k} \tag{39}$$

H.2 Demand Side Parameters Calibration

Table 19 summarizes the full set of parameters to be calibrated in order to close the model. All production coefficients are estimated as the time-averaged expenditure shares.

Table 19
Labor Demand Parameters and Variables

Parameters	Meaning	Value/Data Source
β_{so}^P, β_{so}^N	Labor inputs coefficients	IDA
P_{st}^M	Offshored input prices	UHDI, Broda and Weinstein (2006)
P_{st}^F	Imported good prices	UHDI, Broda and Weinstein (2006)
w_{ot}	Occupation wage	IDA
ρ, α	EOS in production	Doraszelski and Jaumandreu (2018)
σ	EOS in final good	Atalay (2017)
η	EOS btwn domestic and foreign	Simonovska and Waugh (2014)

Total wage bills in occupation-industry pairs identify the labor input coefficients. Danish customs data provides revenue and quantity level information on goods sold to Denmark by each trading partner, which can be fed into the procedure of Broda and Weinstein (2006) to calculate the price indices for offshored inputs and imported varieties.[42] We set the trade elasticity as 4, following Simonovska and Waugh (2014). We obtain raw wage data from the Integrated Database for Labour Market Research (IDA) at Statistics Denmark. Occupation wages are computed by first running a mincerian regression and then divide the residuals by the computed human capital using (3).

I Hat Algebra

I.1 Trade shocks

Consider shocks $\{P_{st}^{F,0}\} \to \{P_{st}^{F,1}\}$, $\{P_{st}^{M,0}\} \to \{P_{st}^{M,1}\}$. We use 0 superscripts to denote the initial steady state, and 1 superscripts to denote the final steady state.

Start from the initial steady state: $\{w_o^0\}$. We assume q_l and $\pi(k|m,s) = \frac{\exp(\delta_{km}+\beta_k s)}{\sum_{k'}\exp(\delta_{k'm}+\beta_{k'}s)}$ are constant. Denote relative changes in variable x by $\hat{x} = \frac{x^1}{x^0}$.

Step 1. Guess $\{w_o^1\}$

Step 2. Compute changes in prices

$$\hat{W}_s^f = \prod_{o\in\mathcal{O}^f} \hat{W}_o^{\beta_{so}}; \quad f = P, N$$

$$\hat{P}_s^P = \left(S_{st}^P \left(\hat{W}_s^P\right)^{1-\rho} + (1-S_{st}^P)\left(\hat{P}_s^M\right)^{1-\rho}\right)^{\frac{1}{1-\rho}}$$

$$\hat{P}_s^D = \frac{1}{\hat{A}_s}\left[S_{st}^N \left(\hat{W}_s^N\right)^{1-\alpha} + (1-S_{st}^N)\left(\hat{P}_s^P\right)^{\frac{1-\alpha}{1-\rho}}\right]^{\frac{1}{1-\alpha}}$$

[42] We follow the approach in Feenstra (1994) to correct for changes in product varieties across time.

$$\hat{P}_s = \left[(1 - S^I_{st}) \left(\hat{P}^D_s \right)^{1-\eta} + S^I_{st} \left(\hat{P}^F_s \right)^{1-\eta} \right]^{\frac{1}{1-\eta}}$$

$$\hat{P} = \left(\sum_s S_{st} \left(\hat{P}_s \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

$$P = \hat{P} P_0$$

Step 3. Compute Bellman Equations by iterating on the following until convergence:

$$V^1(k, m, o) = \nu \log \left(\sum_{o'} \exp \left(\frac{-C_{oo'km} + w^1_{o'} H_{o'}(k, m)/P + \beta V^1(k, m, o')}{\nu} \right) \right)$$

Step 4. Using $V^1(k, m, o)$, solve for the inclusive value of being type (k, m):

$$V^{e,1}(k, m) = \nu \log \sum_{o'} \exp \left(\frac{-C^e_{o,km} + w^1_{o'} H_{o'}(k, m)/P + \beta V^1(k, m, o')}{\nu} \right)$$

Step 5. Compute the NPV for each field m

$$V^1(m|s) = \sum_k V^{e,1}(k, m) \pi(k|m, s).$$

Step 6. Compute $\{L^1_{km}\}$ by Monte Carlo simulation

(a) **Outside the loop** Simulate N students. For each student i draw a test score s_i, gender f_i, high school background, x_i, and type, l_i from the joint distribution $Q(l, s, f, x)$.

(a) For each individual i, generate a test score s_i for i

(b) Generate a type for i from the conditional CDF of the type distribution given s_i:

$$l_i = 1 + \sum_l \mathbf{1}(e > Q_l(s_i)),$$

where e is a uniform random variable.

(c) For $r = 1, ..., R$ calculate the ranked list of choices for individual i:

- Calculate the conditional CDF of choices for field:

$$\tilde{Q}^F_{ir}(m) = \frac{\sum^m_{m'=1 \backslash m_{-r}} \exp \left(\vartheta_{mt} + \vartheta_{lm} + \theta_l V^1(m|s) \right)}{\sum^M_{m'=1 \backslash m_{-r}} \exp \left(\vartheta_{mt} + \vartheta_{lm} + \theta_l V^1(m|s) \right)}$$

- Generate R draws from this distribution:

$$m_{r_i} = 1 + \sum_{m \in m_{-r}} \mathbf{1}(e > \tilde{Q}^F_{ir}(m)),$$

for a uniform random variable.

(d) Assign the agent sequentially to a field in their choice set based on acceptance probabilities (conditional on s). If the agent is not admitted to the first $R-1$ choices, then they are automatically allocated to the R-th choice.

(e) Determine the agents type:

$$k_i = 1 + \sum_{m \in m_{-r}} \mathbf{1}(e > \tilde{Q}^K_i(m, s)),$$

where Q_k is type CDF for agent i with field m and score s, and e is a uniform random variable.

(f) Calculate the share of individuals belong to type (k, m), scale up to $\bar{L}$.

Step 7. Using $V^1(k, m, o)$ solve for λ^1_{km} and $\lambda^{e,1}(o|k, m)$

$$\lambda^1(o'|o, k, m) = \frac{\exp\left(\frac{-C_{oo'km} + w_{o't} H_{o'}(k,m)/P + \beta V^1(k,m,o')}{\nu}\right)}{\sum_{o''} \exp\left(\frac{-C_{oo''km} + w_{o''t} H_{o''}(k,m) + \beta V^1(k,m,o'')}{\nu}\right)}$$

$$\lambda^{e,1}(o|k, m) = \frac{\exp\left(\frac{-C_{o,km} + w_{ot} H_o(k,m)/P + \beta V^1(k,m,o)}{\nu}\right)}{\sum_{o'} \exp\left(\frac{-C_{o,km} + w_{o't} H_{o'}(k,m) + \beta V^1(k,m,o')}{\nu}\right)}$$

Step 8. Compute $\hat{H}^{supply}_o$

(a) Consider the flow equation:

$$L_{omk} = (1-\delta) \sum_{o'} \lambda^1(o|o', k, m) L_{o'mk} + \delta \lambda^{e,1}(o|k, m) L^e_{mk}$$

This implies,

$$L_{\cdot,mk} = \delta \left(I_{NOCC} - (1-\delta)\lambda^1(\cdot|o', k, m)'\right)^{-1} \times \left(\lambda^{e,1}(\cdot|k, m) L^e_{mk}\right)$$

(b) Construct H^1_{omk}:

$$H_{omk} = L_{omk} \times H(o, m, k)$$

(c) Compute changes in human capital supply across occupations

$$\hat{H}_o^{supply} = \sum_{k,m} H_{omk}^1 / H_o^0$$

Step 9. Calculate aggregate expenditure (and income)

$$\hat{E} = \sum_o S_{ot} \hat{H}_o \hat{W}_o$$

Step 10. Calculate changes in domestic output

$$\hat{Q}_s = S_{st}^x \left(\hat{P}_s^D\right)^{-\eta} + (1 - S_{st}^x) \hat{E} \hat{P}_s^{\eta-\sigma} \hat{P}^{\sigma-1} \left(\hat{P}_s^D\right)^{-\eta}$$

Step 11. Calculate a new guess for wages, by enforcing labor market clearing and output market clearing. Choose $\hat{W}_o$ such that

$$\begin{aligned}\hat{H}_o &= \sum_s \frac{H_{sot}}{H_{ot}} \hat{H}_{so} \\ &= \sum_s \frac{H_{sot}}{H_{ot}} \frac{1}{\hat{W}_o} \hat{Q}_s \left(\hat{P}_s^D\right)^{\alpha} \left(\hat{W}_s^N\right)^{-\alpha},\end{aligned}$$

for all $o \in \mathcal{O}^N$; and

$$\begin{aligned}\hat{H}_o &= \sum_s \frac{H_{sot}}{H_{ot}} \hat{H}_{so} \\ &= \sum_s \frac{H_{sot}}{H_{ot}} \frac{1}{\hat{W}_o} \hat{Q}_s \left(\hat{P}_s^D\right)^{\alpha} \left(\hat{P}_s^P\right)^{\rho-\alpha} \left(\hat{W}_s^P\right)^{-\rho},\end{aligned}$$

for all $o \in \mathcal{O}^P$.

Step 12. Update $\hat{W}_o$ and iterate from Step 1 again.

After the algorithm completes, one can update shares and quantities:

$$\begin{aligned}S_{s1}^P &= A^P (W_{s1}^P)^{1-\rho} / (P_1^P)^{1-\rho} = S_{s0}^P \times \left(\frac{\hat{W}_s^P}{\hat{P}^P}\right)^{1-\rho} \\ S_{s1}^N &= A^P (W_{s1}^N)^{1-\alpha} / (P_1^D)^{1-\alpha} = S_{s0}^N \times \left(\frac{\hat{W}_s^N}{\hat{P}^D}\right)^{1-\alpha} \\ S_{s1}^I &= 1 - A (P_{s1}^D)^{1-\eta} / (P_1^D)^{1-\eta} = 1 - (1 - S_{s0}^I) \times \left(\frac{\hat{P}_s^D}{\hat{P}_s}\right)^{1-\eta} \\ S_{s1}^S &= \mu_s P_{s1}^{1-\sigma} / P_0^{1-\sigma} = S_{s0}^S \times \left(\frac{\hat{P}_s}{\hat{P}}\right)^{1-\sigma}\end{aligned}$$

$$
\begin{aligned}
S^x_{s1} =& \frac{A^F_s (P^D_{s1})^{1-\eta}}{Q_{s1} P^D_{s1}} = S^x_{s0} \times \frac{(\hat{P}^D_s)^{-\eta}}{\hat{Q}_s} \\
S_{o1} =& \frac{w_{o1} H_{o1}}{E_1} = \frac{\hat{w}_o \hat{H}_o}{\hat{E}_h} \\
\Pi^P_{os1} =& \frac{(A^P_s)^{1-\rho} \beta_{so} Q_{s1} (P^D_{s1})^{\alpha} (P^P_{s1})^{\rho-\alpha} (W^P_{s1})^{1-\rho} / W_{o1}}{H_{o1}} = \frac{\hat{\beta}_{os} \hat{Q}_s (\hat{P}^D_s)^{\alpha} (\hat{P}^P_s)^{\rho-\alpha} (\hat{W}^P_s)^{-\rho} / \hat{W}_o}{\hat{H}_o} \Pi^P_{os0} \\
\Pi^N_{os1} =& \frac{(A^N_s)^{1-\alpha} \beta_{so} Q_{s1} (P^D_{s1})^{\alpha} (W^N_{s1})^{-\alpha} / W_{o1}}{H_{o1}} = \frac{\hat{\beta}_{os} \hat{Q}_s (\hat{P}^D_s)^{\alpha} (\hat{W}^N_s)^{-\alpha} / \hat{W}_o}{\hat{H}_o} \Pi^N_{os0}
\end{aligned}
$$

I.2 AI Shock

I.2.1 Defining the AI Shock

Suppose there is an external measure that predicts declines in *relative* labor demand. We do not model aggregate reductions in labor demand. Specifically we assume that there is a measure z_o such that within each sector we assume that at *original prices*,

$$
\frac{w_{so} H'_{so}}{w_s H'_s} \propto z_o.
$$

Let $\tilde{z}_{so}$ be a sector-occupation specific shock that ensures the wage share is left unchanged:

$$
\tilde{z}_{so} = z_o \times \frac{1}{\frac{w^N_s L^N_s}{w_S L_S} \left(\sum_{o \in N} \beta_{so} z_o \right) + \frac{w^P_s L^P_s}{w_S L_S} \left(\sum_{o \in P} \beta_{so} z_o \right)}
$$

We want to solve for $\tilde{\beta}$ and $\tilde{A}$ that matches this shock. We do this in two stages, first we look at non-production workers, then production workers. But the strategy is ultimately identical.

For non-production workers, the share of the wage bill is given by,

$$
\frac{w_o L_{so}}{w_s L_s} = \beta^N_{so} A_{sN} \left(\frac{w^N_s}{P^D_s} \right)^{1-\alpha}.
$$

So we need to solve the following:

$$
\tilde{\beta}^N_{so} \tilde{A}_{sN} \left(\frac{w^N_s}{P^D_s} \right)^{1-\alpha} = \tilde{z}_{so} \beta^N_{so} A_{sN} \left(\frac{w^N_s}{P^D_s} \right)^{1-\alpha},
$$

subject to the restriction that $\sum_{o \in N} \tilde{\beta}^N_{so} = 1$. Dividing out the price term and summing we have,

$$
\tilde{A}_{sN} = A_{sN} \sum_o \tilde{z}_{so} \beta^N_o.
$$

Hence,

$$
\hat{A}_{sN} = \sum_o \tilde{z}_{so} \beta^N_{so}.
$$

We also have,

$$\tilde{\beta}^N_{so} = \frac{\tilde{z}_{so}\beta^N_{so}}{\sum_{o'} \tilde{z}_{so'}\beta^N_{so'}}.$$

For production workers, there is a similar common price term that ultimately cancels out, so we have the same expressions:

$$\hat{A}_{sP} = \sum_o \tilde{z}_{so}\beta^P_{so}$$
$$\tilde{\beta}^P_{so} = \frac{\tilde{z}_{so}\beta^P_{so}}{\sum_{o'} \tilde{z}_{so'}\beta^P_{so'}}$$

I.2.2 Modified Algorithm

Consider shocks $\{P^{F,0}_{st}\} \to \{P^{F,1}_{st}\}$, $\{P^{M,0}_{st}\} \to \{P^{M,1}_{st}\}$ as well as $\beta_{so0} \to \beta_{so1}$. We use 0 superscripts to denote the initial steady state, and 1 superscripts to denote the final steady state.

Start from the initial steady state: $\{w^0_o\}$. We assume q_l and $\pi(k|m,s) = \frac{\exp(\delta_{km}+\beta_k s)}{\sum_{k'}\exp(\delta_{k'm}+\beta_{k'}s)}$ are constant. Denote relative changes in variable x by $\hat{x} = \frac{x^1}{x^0}$.

Step 1. Guess $\{w^1_o\}$

Step 2. Compute changes in prices

$$\hat{W}^f_s = \frac{\prod_{o\in\mathcal{O}^f} W^{\beta_{so1}}_{o1}}{\prod_{o\in\mathcal{O}^f} W^{\beta_{so0}}_{o0}}; \quad f = P, N$$

$$\hat{P}^P_s = \left(S^P_{st}\hat{A}_{sP}\left(\hat{W}^P_s\right)^{1-\rho} + (1-S^P_{st})\left(\hat{P}^M_s\right)^{1-\rho}\right)^{\frac{1}{1-\rho}}$$

$$\hat{P}^D_s = \frac{1}{\hat{A}_s}\left[S^N_{st}\hat{A}_{sN}\left(\hat{W}^N_s\right)^{1-\alpha} + (1-S^N_{st})\left(\hat{P}^P_s\right)^{\frac{1-\alpha}{1-\rho}}\right]^{\frac{1}{1-\alpha}}$$

$$\hat{P}_s = \left[(1-S^I_{st})\left(\hat{P}^D_s\right)^{1-\eta} + S^I_{st}\left(\hat{P}^F_s\right)^{1-\eta}\right]^{\frac{1}{1-\eta}}$$

$$\hat{P} = \left(\sum_s S_{st}\left(\hat{P}_s\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$$

$$P = \hat{P}P_0$$

Step 3. Compute Bellman Equations by iterating on the following until convergence:

$$V^1(k,m,o) = \nu \log\left(\sum_{o'}\exp\left(\frac{-C_{oo'km} + w^1_{o'}H_{o'}(k,m)/P + \beta V^1(k,m,o')}{\nu}\right)\right)$$

Step 4. Using $V^1(k,m,o)$, solve for the inclusive value of being type (k,m):

$$V^{e,1}(k,m) = \nu \log \sum_{o'} \exp\left(\frac{-C^e_{o,km} + w^1_{o'} H_{o'}(k,m)/P + \beta V^1(k,m,o')}{\nu}\right)$$

Step 5. Compute the NPV for each field m

$$V^1(m|s) = \sum_k V^{e,1}(k,m)\pi(k|m,s).$$

Step 6. Compute $\{L^1_{km}\}$ by Monte Carlo simulation

(a) For each individual i, generate a test score s_i for i

(b) Generate a type for i from the conditional CDF of the type distribution given s_i:

$$l_i = 1 + \sum_l \mathbf{1}(e > Q_l(s_i)),$$

where e is a uniform random variable.

(c) For $r = 1, ..., R$ calculate the ranked list of choices for individual i:

- Calculate the conditional CDF of choices for field:

$$\tilde{Q}^F_{ir}(m) = \frac{\sum_{m'=1\backslash m_{-r}}^{m} \exp\left(\vartheta_{mt} + \vartheta_{lm} + \theta_l V^1(m|s)\right)}{\sum_{m'=1\backslash m_{-r}}^{M} \exp\left(\vartheta_{mt} + \vartheta_{lm} + \theta_l V^1(m|s)\right)}$$

- Generate R draws from this distribution:

$$m_{r_i} = 1 + \sum_{m \in m_{-r}} \mathbf{1}(e > \tilde{Q}^F_{ir}(m)),$$

for a uniform random variable.

(d) Assign the agent sequentially to a field in their choice set based on acceptance probabilities (conditional on s). If the agent is not admitted to the first $R-1$ choices, then they are automatically allocated to the R-th choice.

(e) Determine the agents type:

$$k_i = 1 + \sum_{m \in m_{-r}} \mathbf{1}(e > \tilde{Q}^K_i(m,s)),$$

where Q_k is type CDF for agent i with field m and score s, and e is a uniform random variable.

(f) Calculate the share of individuals belong to type (k,m), scale up to $\bar{L}$.

Step 7. Using $V^1(k,m,o)$ solve for λ^1_{km} and $\lambda^{e,1}(o|k,m)$

$$\lambda^1(o'|o,k,m) = \frac{\exp\left(\frac{-C_{oo'km}+w_{o't}H_{o'}(k,m)/P+\beta V^1(k,m,o')}{\nu}\right)}{\sum_{o''}\exp\left(\frac{-C_{oo''km}+w_{o''t}H_{o''}(k,m)+\beta V^1(k,m,o'')}{\nu}\right)}$$

$$\lambda^{e,1}(o|k,m) = \frac{\exp\left(\frac{-C_{o,km}+w_{ot}H_o(k,m)/P+\beta V^1(k,m,o)}{\nu}\right)}{\sum_{o'}\exp\left(\frac{-C_{o,km}+w_{o't}H_{o'}(k,m)+\beta V^1(k,m,o')}{\nu}\right)}$$

Step 8. Compute $\hat{H}_o^{supply}$

(a) Consider the flow equation:

$$L_{omk} = (1-\delta)\sum_{o'}\lambda^1(o|o',k,m)L_{o'mk} + \delta\lambda^{e,1}(o|k,m)L^e_{mk}$$

This implies,

$$L_{\cdot,mk} = \delta\left(I_{NOCC} - (1-\delta)\lambda^1(\cdot|o',k,m)'\right)^{-1} \times \left(\lambda^{e,1}(\cdot|k,m)L^e_{mk}\right)$$

(b) Construct H^1_{omk}:

$$H_{omk} = L_{omk} \times H(o,m,k)$$

(c) Compute changes in human capital supply across occupations

$$\hat{H}_o^{supply} = \sum_{k,m} H^1_{omk}/H^0_o$$

Step 9. Calculate aggregate expenditure (and income)

$$\hat{E} = \sum_o S_{ot}\hat{H}_o\hat{W}_o$$

Step 10. Calculate changes in domestic output

$$\hat{Q}_s = S^x_{st}\left(\hat{P}^D_s\right)^{-\eta} + (1-S^x_{st})\hat{E}\hat{P}_s^{\eta-\sigma}\hat{P}^{\sigma-1}\left(\hat{P}^D_s\right)^{-\eta}$$

Step 11. Calculate a new guess for wages, by enforcing labor market clearing and output market clearing. Choose $\hat{W}_o$ such that

$$\begin{aligned}\hat{H}_o &= \sum_s \frac{H_{sot}}{H_{ot}}\hat{H}_{so} \\ &= \sum_s \frac{H_{sot}}{H_{ot}}\frac{1}{\hat{W}_o}\hat{Q}_s\left(\hat{P}^D_s\right)^{\alpha}\left(\hat{W}^N_s\right)^{-\alpha}\hat{\beta}^N_{so}\hat{A}_{sN},\end{aligned}$$

for all $o \in \mathcal{O}^N$; and

$$\begin{aligned}\hat{H}_o &= \sum_s \frac{H_{sot}}{H_{ot}} \hat{H}_{so} \\ &= \sum_s \frac{H_{sot}}{H_{ot}} \frac{1}{\hat{W}_o} \hat{Q}_s \left(\hat{P}_s^D\right)^{\alpha} \left(\hat{P}_s^P\right)^{\rho-\alpha} \left(\hat{W}_s^P\right)^{-\rho} \hat{\beta}_{so}^P \hat{A}_{sP},\end{aligned}$$

for all $o \in \mathcal{O}^P$.

Step 12. Update $\hat{W}_o$ and iterate from Step 1 again.

After the algorithm completes, one can update shares and quantities

$$\begin{aligned}
S_{s1}^P =& A^P (W_{s1}^P)^{1-\rho} / (P_1^P)^{1-\rho} = S_{s0}^P \times \left(\frac{\hat{W}_s^P}{\hat{P}^P}\right)^{1-\rho} \\
S_{s1}^N =& A^P (W_{s1}^N)^{1-\alpha} / (P_1^D)^{1-\alpha} = S_{s0}^N \times \left(\frac{\hat{W}_s^N}{\hat{P}^D}\right)^{1-\alpha} \\
S_{s1}^I =& 1 - A (P_{s1}^D)^{1-\eta} / (P_1^D)^{1-\eta} = 1 - (1 - S_{s0}^I) \times \left(\frac{\hat{P}_s^D}{\hat{P}_s}\right)^{1-\eta} \\
S_{s1}^S =& \mu_s P_{s1}^{1-\sigma} / P_0^{1-\sigma} = S_{s0}^S \times \left(\frac{\hat{P}_s}{\hat{P}}\right)^{1-\sigma} \\
S_{s1}^x =& \frac{A_s^F (P_{s1}^D)^{1-\eta}}{Q_{s1} P_{s1}^D} = S_{s0}^x \times \frac{(\hat{P}_s^D)^{-\eta}}{\hat{Q}_s} \\
S_{o1} =& \frac{w_{o1} H_{o1}}{E_1} = \frac{\hat{w}_o \hat{H}_o}{\hat{E}_h} \\
\Pi_{os1}^P =& \frac{(A_s^P)^{1-\rho} \beta_{so} Q_{s1} (P_{s1}^D)^{\alpha} (P_{s1}^P)^{\rho-\alpha} (W_{s1}^P)^{1-\rho} / W_{o1}}{H_{o1}} = \frac{\hat{\beta}_{os} \hat{Q}_s (\hat{P}_s^D)^{\alpha} (\hat{P}_s^P)^{\rho-\alpha} (\hat{W}_s^P)^{-\rho} / \hat{W}_o}{\hat{H}_o} \Pi_{os0}^P \\
\Pi_{os1}^N =& \frac{(A_s^N)^{1-\alpha} \beta_{so} Q_{s1} (P_{s1}^D)^{\alpha} (W_{s1}^N)^{-\alpha} / W_{o1}}{H_{o1}} = \frac{\hat{\beta}_{os} \hat{Q}_s (\hat{P}_s^D)^{\alpha} (\hat{W}_s^N)^{-\alpha} / \hat{W}_o}{\hat{H}_o} \Pi_{os0}^N
\end{aligned}$$

I.3 Transition Dynamics

Step 1. Calculate final steady state

Step 2. Guess a path of wages $\{w_{ot}\}_{o,t=1,T-1}$

Step 3. Solve for the path of prices using the hat algebra:

$$\hat{W}_{st}^{f} = \frac{\prod_{o \in \mathcal{O}^f} W_{ot}^{\beta_{sot}}}{\prod_{o \in \mathcal{O}^f} W_{o0}^{\beta_{so0}}}; \quad f = P, N$$

$$\hat{P}_{st}^{P} = \left(S_{s0}^{P} \hat{A}_{stP} \left(\hat{W}_{st}^{P} \right)^{1-\rho} + (1 - S_{s0}^{P}) \left(\hat{P}_{st}^{M} \right)^{1-\rho} \right)^{\frac{1}{1-\rho}}$$

$$\hat{P}_{st}^{D} = \frac{1}{\hat{A}_{st}} \left[S_{s0}^{N} \hat{A}_{stN} \left(\hat{W}_{st}^{N} \right)^{1-\alpha} + (1 - S_{s0}^{N}) \left(\hat{P}_{st}^{P} \right)^{\frac{1-\alpha}{1-\rho}} \right]^{\frac{1}{1-\alpha}}$$

$$\hat{P}_{st} = \left[(1 - S_{s0}^{I}) \left(\hat{P}_{st}^{D} \right)^{1-\eta} + S_{s0}^{I} \left(\hat{P}_{st}^{F} \right)^{1-\eta} \right]^{\frac{1}{1-\eta}}$$

$$\hat{P}_{t} = \left(\sum_{s} S_{s0} \left(\hat{P}_{st} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

$$P_t = \hat{P}_t P_0$$

Step 4. Solve the path of Bellman equations:

$$V_t(k, m, o) = \nu \log \left(\sum_{o'} \exp \left(\frac{-C_{oo'km} + w_{o't} H_{o'}(k, m)/P_t + \beta V_{t+1}(k, m, o')}{\nu} \right) \right)$$

Step 5. Solve for the path of inclusive values of being type (k, m):

$$V_t^e(k, m) = \nu \log \sum_{o'} \exp \left(\frac{-C_{o,km}^e + w_{o't} 1 H_{o'}(k, m)/P_t + \beta V_{t+1}(k, m, o')}{\nu} \right)$$

Step 6. Compute the NPV for each field m

$$V_t(m|s)_t = \sum_{k} V_{t+5}^e(k, m) \pi(k|m, s).$$

Notice that for $t = T-4, ..., T-1$, one will use the steady state value.

Step 7. Compute $\{L_{kmt}\}$ by Monte Carlo simulation

(a) For each individual i and period t, generate a test score s_i for i. We do not index by t because they will be copies each period.

(b) Generate a type for i from the conditional CDF of the type distribution given s_i:

$$l_i = 1 + \sum_l \mathbf{1}(e > Q_l(s_i)),$$

where e is a uniform random variable.

(c) For $r = 1, ..., R$ calculate the ranked list of choices for individual i:

- Calculate the conditional CDF of choices for field:

$$\tilde{Q}^F_{ir}(m) = \frac{\exp\left(\vartheta_{lm} + \theta_l V_t(m|s)\right)}{\sum_{m'=1\setminus m_{-r}}^{M} \exp\left(\vartheta_{lm'} + \theta_l V_t(m'|s)\right)}$$

Note: For cohorts $t = 1, ..., 5$ allocate students according to the steady state *values of $V_t(m|s)$. For remaining students, allocate according to $V_t(m|s)$.*

- Generate R draws from this distribution:

$$m_{r_i} = 1 + \sum_{m \in m_{-r}} \mathbf{1}(e > \tilde{Q}^F_{ir}(m)),$$

for a uniform random variable.

(d) Assign the agent sequentially to a field in their choice set based on acceptance probabilities (conditional on s). If the agent is not admitted to the first $R-1$ choices, then they are automatically allocated to the R-th choice.

(e) Determine the agents type:

$$k_i = 1 + \sum_{m \in m_{-r}} \mathbf{1}(e > \tilde{Q}^K_i(m, s)),$$

where Q_k is type CDF for agent i with field m and score s, and e is a uniform random variable.

(f) Calculate the share of individuals belong to type (k, m), scale up to $\bar{L}$.

Step 8. Using $V_t(k, m, o)$ solve for λ_{kmt} and $\lambda^e_t(o|k, m)$

$$\lambda_t(o'|o, k, m) = \frac{\exp\left(\frac{-C_{oo'km} + w_{o't} H_{o'}(k,m)/P_t + \beta V_{t+1}(k,m,o')}{\nu}\right)}{\sum_{o''} \exp\left(\frac{-C_{oo''km} + w_{o''t} H_{o''}(k,m)/P_t + \beta V_{t+1}(k,m,o'')}{\nu}\right)}$$

$$\lambda_t^e(o|k,m) = \frac{\exp\left(\frac{-C_{o,km}+w_{ot}H_o(k,m)/P_t+\beta V_{t+1}(k,m,o)}{\nu}\right)}{\sum_{o'}\exp\left(\frac{-C_{o,km}+w_{o't}H_{o'}(k,m)/P_t+\beta V_{t+1}(k,m,o')}{\nu}\right)}$$

Step 9. Compute $\hat{H}_o^{supply}$

(a) Use the flow equation:

$$L_{omkt} = (1-\delta)\sum_{o'}\lambda_t(o|o',k,m)L_{o'mkt-1} + \delta\lambda_t^e(o|k,m)L_{mkt}^e$$

(b) Construct H_{omkt}:

$$H_{omkt} = L_{omkt} \times H(o,m,k)$$

(c) Compute changes in human capital supply across occupations

$$\hat{H}_{ot}^{supply} = \sum_{k,m} H_{omkt}/H_{o0}$$

Step 10. Calculate aggregate expenditure (and income)

$$\hat{E}_t = \sum_o S_{o0}\hat{H}_{ot}\hat{W}_{ot}$$

Step 11. Calculate changes in domestic output

$$\hat{Q}_{st} = S_{s0}^x\left(\hat{P}_{st}^D\right)^{-\eta} + (1-S_{s0}^x)\hat{E}\hat{P}_{st}^{\eta-\sigma}\hat{P}_t^{\sigma-1}\left(\hat{P}_{st}^D\right)^{-\eta}$$

Step 12. Calculate a new guess for wages, by enforcing labor market clearing and output market clearing. Choose $\hat{W}_o$ such that

$$\begin{aligned}\hat{H}_{ot} &= \sum_s \frac{H_{so0}}{H_{o0}}\hat{H}_{so} \\ &= \sum_s \frac{H_{so0}}{H_{o0}}\frac{1}{\hat{W}_{ot}}\hat{Q}_{st}\left(\hat{P}_{st}^D\right)^{\alpha}\left(\hat{W}_{st}^N\right)^{-\alpha}\hat{\beta}_{so}^N\hat{A}_{sN},\end{aligned}$$

for all $o \in \mathcal{O}^N$; and

$$\begin{aligned}\hat{H}_{ot} &= \sum_s \frac{H_{so0}}{H_{o0}}\hat{H}_{sot} \\ &= \sum_s \frac{H_{so0}}{H_{o0}}\frac{1}{\hat{W}_{ot}}\hat{Q}_{st}\left(\hat{P}_{st}^D\right)^{\alpha}\left(\hat{P}_{st}^P\right)^{\rho-\alpha}\left(\hat{W}_{st}^P\right)^{-\rho}\hat{\beta}_{sot}^P\hat{A}_{sP},\end{aligned}$$

for all $o \in \mathcal{O}^P$.

Step 13. Update $\hat{W}_{ot}$ and iterate from Step 2 again.