

Field Choice, Skill Specificity, and Labor Market Disruptions*

Valerie Smeets
Aarhus University & CEPR

Lin Tian
INSEAD & CEPR

Sharon Traiberman
NYU & NBER

April 22, 2025

Abstract

We argue that college students' field-of-study choices significantly influence how economies respond to labor market disruptions. To do so, we develop and estimate a framework featuring forward-looking students who choose a field of study when entering college, and subsequently make decisions over occupations after graduating and entering the labor market. Different fields endow workers with distinct comparative advantages and varying costs associated with switching occupations. Simulating both a trade war and wide scale adoption of AI, we use our model to make three points. First, relative to models that ignore how new cohorts adjust their field-of-study choices, our framework predicts larger aggregate income responses and greater distributional differences. Second, policies that enhance flexibility in field-of-study decisions—such as relaxing capacity constraints in high-demand programs—raise aggregate output. Finally, these policies also lessen the adverse distributional consequences of shocks, by affording more opportunities to students with lower earnings potential.

Keywords: Skill Specificity, Human Capital Investment, Labor Market Disruptions

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, Christopher Flinn, Matt Notowidigdo, Esteban Rossi-Hansberg, Martin Rotemberg, Dan Treffer, Tim Van-Zandt, Jonathan Vogel, Dan Waldinger, Matt Wiswall, 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

Trade disruptions and rapid technological advancement are major labor market shocks that reshape occupational structures and skill demands worldwide. Understanding how societies adapt to such shifts is critical to answering many questions important to economists and policy makers: How effectively can workers mitigate adverse impacts? Who gains or loses, and how can policy influence these outcomes? We study these questions by examining how workers adapt their skills, particularly through college field-of-study choices of new cohorts, to evolving labor market conditions (Heckman et al., 1998; Adao et al., 2024). Fields of study differ substantially in earnings potential and adaptability to economic shocks; they shape individual career trajectories, match students to types of tasks and occupations, and determine aggregate skill distribution (Altonji et al., 2012a; Kinsler and Pavan, 2015; Kirkeboen et al., 2016). Thus, modeling the joint determination of education and career decisions is necessary for capturing how economies respond to shocks. Yet, despite their importance, few existing studies analyze college field of study choices alongside career paths.

We bridge this gap by developing and estimating a dynamic model that jointly captures educational choices by new entrants and subsequent occupational decisions by incumbent workers. Agents enter our model at high school graduation—having already decided whether to pursue a college degree—and select a field of study based on expected returns and idiosyncratic preferences. They subsequently make occupational decisions in each period, subject to switching costs. We use our model to perform several counterfactuals that answer the questions posed above. First, we examine how field-of-study decisions influence the aggregate and distributional consequences of labor market shocks, focusing on two specific scenarios: a global trade war that increases the cost of imported goods and reduces foreign demand for domestic exports, and a change in relative labor demand brought about by wide scale adoption of Large Language Models (LLMs). Second, we analyze how increased flexibility over students’ education choices, e.g., relaxing capacity constraints in high-demand programs, impacts long-term welfare and the economy’s adjustment to these shocks.

We leverage Danish administrative data that links workers to their college applications, including the ranked list of fields each student applied to. Denmark’s centralized, algorithmic college admissions system provides a useful setting for identifying student preferences separately from selection effects, allowing us to model how field choices respond to changes in expected earnings and to perform realistic counterfactual analyses.

Using this dataset, we establish three facts that motivate our model. First, field-of-study choices

strongly predict occupational specialization. For example, over 75% of Education graduates become teachers, while Social Science graduates are more uniformly spread out across occupations. Second, changes in the distribution of occupations over time are primarily driven by new workforce entrants rather than incumbent workers switching occupations. Third, this new-cohort driven adjustment reflects two distinct margins: students increasingly avoid fields tied to declining occupations (e.g., Engineering) in favor of growing ones (e.g., Law), while within fields, graduates increasingly divert to different career paths, e.g., STEM graduates shift from production-related occupations to financial services. These findings underscore the importance of modeling both educational choices and occupational dynamics to fully capture the mechanisms of labor market adjustments.

Having established these facts we turn to our model. At the education stage, students decide which field of study to pursue considering both expected labor market earnings and non-pecuniary preferences for different subjects.¹ Students differ on a wide set of dimensions, which in turn influences their preferences and abilities. Upon graduation they draw their labor market productivity from a distribution that depends on both their high school background and other observable characteristics, and their field of study. Labor market productivity is multidimensional, with workers having different comparative advantages across occupations. For example, students who have studied Social Sciences draw from a distribution that endows them with an ability at customer service, an ability at IT, and so on—on average these students are more likely to excel at customer service relative to IT. A student who studied Physics/Math will draw skills over the same set of occupations, but from a different distribution—their distribution may be such that they excel at IT, on average, relative to customer service.²

After drawing their labor market productivity, students choose an initial occupation. They choose occupations each period subject to switching costs à la [Artuç et al. \(2010\)](#). Students, even with the same educational background, ultimately make different career decisions and respond differently to changes in the labor market for two reasons: idiosyncratic switching cost shocks, and different comparative advantages across occupations.

We estimate our model in two stages. For the labor supply block of the model, we use a two-step approach following [Arcidiacono and Miller \(2011\)](#) and [Traiberman \(2019\)](#), first employing the

¹We do not model the extensive margin of entry into a Bachelor’s program. While this is an important adjustment mechanism ([Ferriere et al., 2021](#)), our focus is on the horizontal decision over different fields. An important path for future work would be to jointly model these decisions, taking account of changes in the composition of entrants to the college system.

²More precisely, each occupation consists of a set of tasks, and different tasks are differentially important across occupations. Each worker draws a vector of task-specific productivities, which—together with the task composition of each occupation—determines their comparative advantages. Section 4 provides more details.

Expectation-Maximization (EM) algorithm to recover comparative advantages and occupational transition rates, and then regressing these transition rates on wage differentials to obtain the remaining parameters as in [Artuç et al. \(2010\)](#). Intuitively, the panel structure of the data allows us to control for changes in labor supply over time, while the variation in wages induced by demand identifies labor supply elasticities. For the field choice block, we use Denmark’s centralized admissions data, exploiting students’ complete ranked choices to uncover latent heterogeneity using methods adapted from [Agarwal and Somaini \(2018\)](#), and [Fack et al. \(2019\)](#). By linking worker registry data and education data, we can use time variation in returns across fields to estimate students’ elasticity of field-choice with respect to expected income.

There are several key findings from our estimation. First, graduates from different fields differ significantly in how they respond to shocks. We measure this by the semi-elasticity of occupational exit rates with respect to wages. Workers from more generalized fields, such as Business Administration, are roughly six times more likely to exit their occupations following a given 1% decline in wages compared to those from highly specialized fields like Healthcare. Second, students’ field choices are themselves responsive to future income prospects, with a 1% increase in the net present value (NPV) of a field raising entry into that field by 0.92% (or 0.14 percentage points). We also identify strong clustering in field preferences, such as within STEM or healthcare disciplines, where students readily substitute across closely related fields but rarely switch to entirely different domains. Lastly, incorporating preference heterogeneity is critical for realistically capturing students’ substitution patterns. We demonstrate this by comparing our model’s predictions to those from a standard logit model (which is what can be identified without our ranked choice data). Our approach matches students’ observed preferences well, whereas the simpler model generates implausible substitution patterns that significantly deviate from observed behavior.

To quantify the general equilibrium (GE) effects of labor market shocks and education policies, we close the model by incorporating the labor demand side. Specifically, we extend a standard multi-sector framework in which each sector employs different mixes of occupations, allowing for heterogeneous substitutability across occupations. Using this framework, we first examine how flexibility in education policies affects labor market outcomes, by simulating two policy reforms: (1) eliminating program enrollment caps (e.g., capacity limits for high-demand fields) while retaining high-school tracking, which restricts students’ choices of programs based on their upper-secondary academic track; and (2) removing both tracking and enrollment caps, enabling students to freely choose fields aligned with their preferences and anticipated labor market returns. Removing

enrollment caps alone raises GDP by 0.3%, while completely removing all constraints generates a fourfold larger GDP gain of 1.3%. In both cases, the equilibrium distribution of field choices changes substantially, demonstrating that both constraints are binding for many students.

Finally, we evaluate how education policy flexibility shapes economic outcomes during two simulated labor market shocks: (1) a global trade war, modeled as increases in the costs of imported goods and a decline in foreign demand of domestic exports, and (2) an AI-driven technological shock, characterized by occupation-specific productivity shifts. Counterfactual simulations reveal two main findings. First, the flexibility in education policies allows individuals to more effectively respond to shifting labor demand through field choices, mitigating aggregate income losses. For example, relative to a model ignoring the field-choice-margin, even modest adjustments under the current set of capacity constraints, reverse an income loss due to the trade war to a small income gain. Relaxing constraints in line with the first counterfactual amplifies these gains, as students fully align choices with emerging opportunities.

Second, while more flexibility raises overall inequality, it improves lifetime earnings for the most vulnerable groups. When students can freely reoptimize their field choices, income gains shift toward initially lower-income individuals, while losses become concentrated among higher earners. Intuitively, constraints favor students with high earnings potential, because earnings potential and test scores are highly correlated. Relaxing these constraints allows students with lower scores to better re-optimize following a shock. Quantitatively, we find that it is more valuable to be a student with a low test score in a capped field, than to be in an unconstrained fields. These results underscore that field of study choices are a critical adjustment mechanism during structural disruptions, whereas flexibility in these choices enable more adaptive reallocation of skills, enhancing economic resilience.

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), labor market discrimination (Hsieh et al., 2019), recessions

(Grigsby, 2021), 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 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; Eckardt, 2024), a growing body of work examines variation in returns across fields of study (Hastings et al., 2013; Gemici and Wiswall, 2014; Kirkeboen et al., 2016; Bleemer and Mehta, 2022; Andrews et al., 2022) and the role of major specificity in shaping career trajectories (Kinsler and Pavan, 2015; Leighton and Speer, 2020).³ Our paper is most closely related to studies that estimate models of human capital accumulation alongside labor market processes (Heckman et al., 1998; Arcidiacono, 2004; Beffy et al., 2012; Ransom, 2021) and builds on the seminal work by Keane and Wolpin (1997), which integrates labor demand into college decisions.⁴ 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 and technological change—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

³Altonji et al. (2012b) and Altonji et al. (2016) provide extensive surveys, including discussions of structural modeling issues.

⁴While our paper focuses on higher education, a substantial literature highlights the importance of early human capital accumulation (e.g., Caucutt and Lochner, 2020).

(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 (Luflade, 2019; Kapor et al., 2024; Fack et al., 2019; Larroucau and Rios, 2022). There is also a growing literature on how particular institutional arrangements affect students’ decision making (Bordon and Fu, 2015; Bleemer and Mehta, 2024). Our approach brings these strands of the literature together, with a particular focus on occupational switching, and shocks that disrupt occupational structure.

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’ 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, International Standard Classification of Occupation (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 an 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 highest level of degree earned, the field of study if applicable, 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 college programs from the *Coordination Register* (KOT). The data are provided by the Coordinated Admissions System, which is run by the Danish Ministry of Higher Education and Science. For each student, we observe 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 track and GPA. We hand-link the education programs in KOT to UDDA and calculate admission cutoffs based on all applications data.⁵

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,⁶ and college and above. The last group is the primary focus of our analysis. The KOT data only includes information on students applying to college programs, and so we focus our attention on this group in the estimation, treating the decision to enter the KOT system (i.e., whether to pursue a college degree) as exogenous.⁷ Our definition of fields is based on the Danish variant of the ISCED classification of fields, i.e., ISCED-F. The ISCED-F classification breaks fields into 27 broad categories. We aggregate several codes—especially in the natural sciences—due to their small sizes. We also break apart the health code into programs at the Professional Bachelor's level and the University Bachelor's level. Appendix Table 9 contains our list of educational fields and their ISCED-F equivalent codes.

Occupations in Denmark are recorded according to the ISCO, developed by the International Labor Organization (ILO). Appendix C.3 contains a complete description of how we aggregate and

⁵We 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 C.2.1 contains a complete description of how we concurred KOT to UDDA and calculated admission cutoffs.

⁶Short-cycle programs, corresponding to ISCED Level 5, are practical programs designed to prepare students for specific occupations, similar to vocational programs. In Denmark, these include, for example, police academies, some health care programs, maritime programs, and some vocational tracks.

⁷There are a handful of programs below college administered through the KOT. However, this is a very small fraction of total programs. They tend to be highly specialized programs—such as military and maritime programs that are more outside the traditional scope of tertiary education.

concord occupations over time, as well as how we deal with a host of quality and measurement issues. Ultimately, we arrive at 22 time-consistent occupation codes. Appendix Table 10 contains our list of occupations and their ISCO-08 equivalent codes. Finally, Appendix C.3.1 discusses 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 college 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. An increasingly large share of Danish students pursue a degree from a Professional Bachelor's or University Bachelor's (henceforth "college"). Appendix Figure 11 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

⁸While 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.

⁹Our use of 30 year olds may 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>.

upper secondary education consists of several tracks, each influencing the types of programs students can pursue beyond high school (see Appendix C.1 for more information on college 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 college education. We also use high school tracks as a control when estimating students' returns and preferences over different fields of study. In our counterfactual analysis, we assess the economic implications of removing high school tracking from the college admissions process.

The Application System

Entry into college programs occurs through a centralized admissions process. Students can apply up to eight programs (each application being a combination of field and institution). Each program has its own high-school 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.

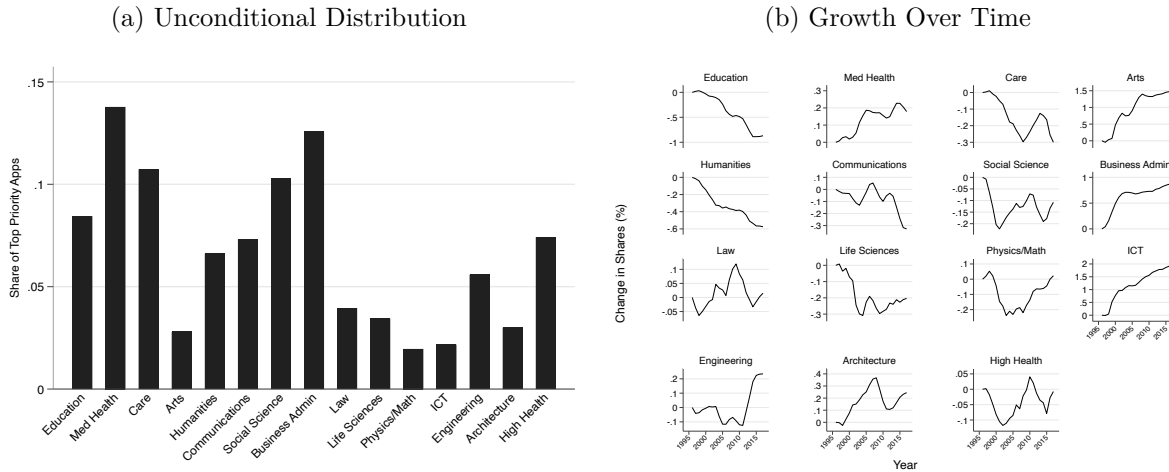
Students in Denmark are provided with significant resources to help them make informed education choices and improve their information about alternatives educational paths.¹⁰ Nevertheless, while students inevitably face some gaps in information and cannot fully predict the returns to different fields of study, significant efforts are made to provide them with comprehensive guidance and resources.¹¹

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). Quotas bind for many programs, leading most students to apply to multiple programs.¹²

¹⁰For example, the *UddannelsesGuiden* (Education Guide) website (<https://www.ug.dk/>) provides detailed information on program-specific prerequisites, historical GPA cutoffs, admissions numbers, labor market outcomes, demand trends across fields of study, and subject matter expectations. Study guidance centers (*Studievalg*) further support students through counseling services.

¹¹This is line with the fact that information plays a crucial role in educational choices, and that students often face

Figure 1
Summary Statistics on First Priority



Note: Panel (a) plots the share of first-choice fields (first priority) reported by high school graduates who apply to college education through the KOT system, regardless of final admission status. 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.

Figure 1 displays the unconditional distribution of first priorities for students across college fields of study, as well as the growth in fields over time. The most popular fields are in Care and Medium Health (e.g., nursing and health technicians), followed closely by Business, Social Sciences and Education. STEM fields, when including High Health (e.g., doctors and dentists), 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, whereas Education has experienced significant decline.

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.¹³ Second, preferences vary significantly by

significant information gaps when making education decisions (Wiswall and Zafar, 2015; Kapor et al., 2020).

¹²While the majority of admissions takes place through the scores-based assignment mechanism (known as "Quota 1"), a fraction of admissions take place based on supplementary exams or other criteria for students who were rejected from their Quota 1 application (known as "Quota 2"). In our analysis, we keep only Quota 1 admits. It is also not uncommon for students to wait between high school and college if they do not receive their top choice. This practice is discouraged, but continues. We have estimated the model pooling all applications, and only keeping applications the year students choose to enter school, and this has little impact on our final estimates.

¹³In our subsequent analysis, we do not distinguish between students choosing a similar program within the same institution (e.g., Physics vs. Mathematics at Aarhus) or applying to the same program at different institutions (e.g., Economics at Copenhagen vs. Aarhus). 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%

discipline. In STEM fields such as Life Sciences, Physics/Math, 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. These facts in turn motivate our model’s key ingredients: costly occupational switching for incumbents, and an ex ante choice over field of study which shapes one’s labor market comparative advantages.

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

We first demonstrate that fields of study are correlated with students’ career outcomes. In particular, we show 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, 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 by pooling across the all college workers. Figure 2 plots the HHI for each field of study, as well as the HHI calculated by pooling all fields (in grey).

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 0.2, while the average HHI conditional on field is nearly twice as high. Second, specialization varies significantly across fields. For example, over 75% of Education graduates become teachers, while Social Sciences graduates are more uniformly spread out across occupations with an HHI slightly below the pooled average. 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


of the time they apply for similar programs across institutions.

	First Priority														
Second Priority	Education	Care	Med Health	Arts	Humanities	Social Sciences	Communications	Business Admn	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 Sciences				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													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.

Figure 2
Field of Study-Occupation HHIs



Notes: Black bars represent the HHI for different fields of study, where fields are defined via the Danish variant of the ISCED classification (discussed in Section 2.2). The grey bar represents the HHI calculated by pooling workers from all fields. Data are pooled from 1996-2018, and conditioned on positive earnings.

greater challenges in adapting to shifting labor market conditions.

To better understand the role of occupational sorting in driving labor market outcomes, we document large heterogeneity in the returns across different fields of study in Appendix Figure 12. Panel (a) shows average earnings for workers by field relative to those a without college education, alongside the overall college premium, defined as the ratio of average earnings of all college graduates relative to non-college graduates. The variation in returns is enormous—with some fields yielding lower earnings than workers who never attended college. This speaks to possible selection, motivating our subsequent structural model and identification strategy. Panel (b) incorporates controls for observable characteristics.¹⁴ While controlling for these observables flatten the field premia, significant heterogeneity remains. Finally, connecting these to occupational sorting, Panel (c) additionally removes occupation fixed effects. Accounting for occupation fixed effects further reduces the variation in field premia (except notably for “Arts”) and aligns most fields closer to the overall college premium, suggesting that the matching of fields of study to specific occupations explains a large share of differences in returns across fields.

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

We now show that shift in the occupational structure over time occurs through two channels: incumbent workers switching occupations (or exiting the labor force) and new cohorts entering a

¹⁴We control for a polynomial in age, gender, and year fixed effects.

different mix of occupations.

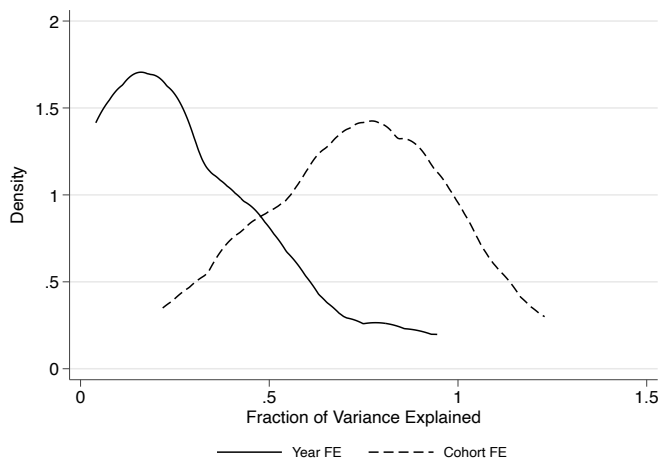
To quantify these channels, we adopt the approach of [Porzio et al. \(2022\)](#). Let $s_{o|c,t}$ be the share of workers in cohort c working in occupation o at time t . We project this share on cohort and time fixed effects for each occupation:

$$\log s_{o|c,t} = \varphi_{oc} + \psi_{ot} + \varepsilon_{oct}. \quad (1)$$

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.

To assess the relative importance of incumbents and new entrants, [Figure 3](#) reports the ratio of the variance in fixed effects to the variance of the fitted log shares: $\sigma_{\varphi_{oj}}^2 / \sigma_{\log s_{o|c,t}}^2, j \in \{c, t\}$.¹⁵ The distribution of these ratios suggests that cohort effects generally explain a larger share of the variation, underscoring the role of new cohort entry in occupational reallocation.¹⁶

Figure 3
Variance Explained by Cohort and Year Effects



Notes: Density distribution of the fraction of explained variance in log employment shares across occupations that is attributable to cohort and year fixed effects. For each fixed effect f , this plots $\hat{\sigma}_f^2 / \hat{\sigma}_{\log s_{o|ct}}^2$ across occupations o . Sample is all college workers from 1996-2018.

¹⁵The full set of cohort and year fixed effects are plotted in [Appendix Figure 13](#).

¹⁶This ratio can technically be larger than one due to the covariance of cohort and year fixed effects, though this term is generally negligible.

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 following decomposition:

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

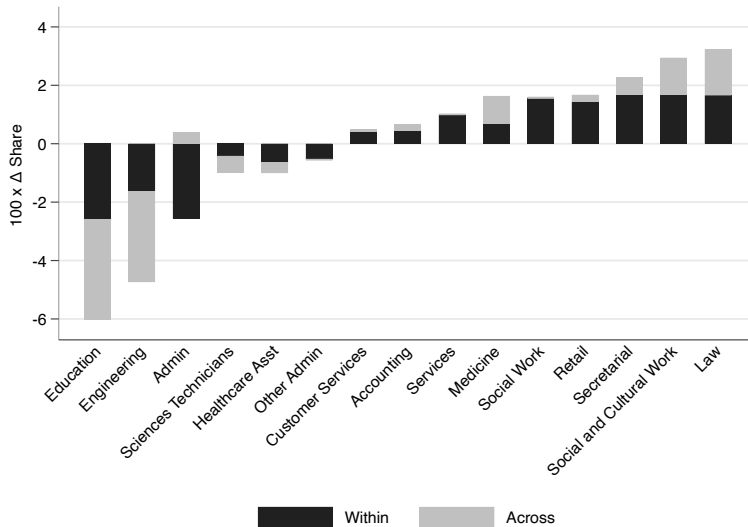
The first term is the within component of occupational reallocation: the amount of occupational reallocation one would expect holding fixed the initial distribution of fields, but changing the match between fields and occupations. The second term is the across component: the amount of reallocation one would expect if the matching between occupations and fields were held fixed in the base year, but new cohorts selected different majors. Finally, there is a covariance term which captures whether or not graduates in growing fields are also moving into growing occupations.

Figure 4 plots this decomposition for occupations with at least a 5% change in share.¹⁷ 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 social work and administration, where field choice appears less important.

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.

¹⁷The covariance term is generally small and is omitted here. The full decomposition is provided in Appendix Figure 14. Moreover, as mentioned, the number and types of students entering the centralized mechanism is growing over time. For this reason, some fields are declining in shares—but almost no fields are declining in levels.

Figure 4
Decomposing the New Cohort Entry Margin



Notes: This figure plots the decomposition of occupational reallocation into across- and within- field components according to the Kitagawa-Oaxaca-Blinder decomposition described in Equation (2). Occupations are sorted based on total growth over the sample period. Only occupations with at least 1% change are plotted. Appendix 14 contains second order terms.

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 heterogeneities in preferences over fields and in comparative advantages 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 $k \in \{1, \dots, K\}$. At this point, k is an index over latent types that captures any differences across workers in their productivity—for example, some workers may be better at all jobs than other workers (absolute advantage), and some workers may be relatively better at some occupations than others (comparative advantage)—as well as the costs associated with switching occupations. Importantly, it is known to the worker, but unobserved by the econometrician. In Section 5 we discuss how we identify these labor market types, and their dependence on field of study m .

Each worker supplies one unit of labor inelastically 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 observable characteristics X . The preference type l indexes heterogeneity in students' preferences over fields of study, as detailed later.

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 *labor supply 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 $\zeta_o = (\zeta_{1,o}, \dots, \zeta_{Z,o}) \in \mathbb{R}_{++}^Z$, 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. For example, $\zeta_{z,o}$ can be mathematical tasks, which will be more important in professions such as Engineering or Accounting, and less important in Law.

Each worker type has task-specific abilities $\alpha_{k,z}$, and their log output in an occupation is the dot product of α 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.¹⁸ Combining, we have:

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

A few comments are in order. First, the α_{om} terms allow workers in different fields of study to have distinct comparative advantages. For example, workers with an education background in STEM may have large α_{om} values in technical occupations, such as engineering. Second, even within a field, workers can have distinct comparative advantages across occupations due to differences in their labor market type k .¹⁹ Third, the conditional distribution of k depends on the chosen field of study m , a relationship we discuss in detail later. As a final note, we abstract from experience or occupation-specific tenure in order to focus on how educational fields shape skill development.²⁰

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,

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

where $V \equiv \int v(\cdot, \varepsilon) dG(\varepsilon)$ is the value function v integrated over shocks, and β is the effective discount factor, 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 function V simplifies to:

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

Moreover, the corresponding policy function, which determines the transition probabilities across

¹⁸Technically we use one index for all productivity differences across workers—absolute and comparative advantages. 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. Moreover, since α_{om} is flexible, no additional α_{mz} term could be identified. The purpose of the tasks is lower the dimensionality of the latent variables in estimation.

¹⁹As 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.

²⁰However, our framework can accommodate such an extension, (e.g., [Traiberman, 2019](#)), and understanding how incumbents can reshape knowledge more generally is likely an important avenue for future work ([Humlum et al., n.d.](#)).

occupations, is:

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

Equation (6) can be used to interpret ν as governing the elasticity of occupational switching with respect to wages, as opposed to idiosyncratic shocks. At the extremes, as $\nu \rightarrow 0$, the transition probabilities collapses to a deterministic outcome where all workers choose the occupation with the highest expected payoff. Conversely, as $\nu \rightarrow \infty$, 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 choose an initial occupation, according to the utility function,

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

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:

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

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

$$V_t^e(k, m) = \nu \log \sum_{o'} \exp\left(\frac{-C_{o,km}^e + w_{o't}H_{o'}(k, m) + \beta E_t V_{t+1}(k, m, o')}{\nu}\right). \quad (9)$$

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, $l \in 1, \dots, L$. They may also have additional covariates, such as their high school track or gender, which we collect in the vector X . For simplicity, we omit these covariates in the following discussion, though all outcomes are implicitly conditioned on them on these characteristics. We mention here that throughout the paper, we do not model the decision to enter college or not, this is treated as

exogenous. For those students who are not in the college system, we treat the distribution of fields as exogenous throughout.

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, given by 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 $\pi(k|m, s)$.²¹ 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:

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

where $V_{t+5}^e(k, m)$, given by (9), is the lifetime payoff to the student upon graduation with a labor market type k and a field m .

In addition to the expected labor market returns, students' field choices reflect non-pecuniary factors that vary by latent preference type l . 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 term $v \log(N_m)$, where N_m is the number of programs available in field m . For example, the "Physics/Math" field may encompass multiple mathematics and physics programs across institutions, which students may find attractive.²² Finally, students receive a conditionally independent Gumbel(0,1) preference shock $\varepsilon_m^{\mathcal{M}}$ over fields, introducing idiosyncratic variation in their choices.

²¹An 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.

²²As the focus of our paper is not on institutional quality or on within-field variation, we treat programs within a field of study as symmetric. Field-level differences in quality or popularity are partially reflected in the admission cut-offs described below.

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

$$\tilde{U}_{mst} = \vartheta_{lm} + \theta_l V_t(m|s) + v \log(N_m) + \varepsilon_m^{\mathcal{M}}. \quad (10)$$

Henceforth, we use \tilde{U} to denote the utility inclusive of the shock, $\varepsilon^{\mathcal{M}}$, and U to refer to the deterministic component of utility excluding the shock, i.e., $U_{mst} = \vartheta_{lm} + \theta_l V_t(m|s) + v \log(N_m)$.

Ranked List Determination

Upon receiving their preference shocks ε_m^m , students construct a ranked list of field choices 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.²³ Students' choice sets, $\mathcal{C} \subset \mathcal{F}$, may depend on their test scores, high school background, and other covariates. 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, \dots, m_R) is given by,

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

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.

5 Estimation

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

²³Since 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.

5.1 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. The estimation proceeds in two stages: first, we apply the EM algorithm to estimate parameters governing human capital supply, latent labor market types, and occupation transition probabilities; second, we estimate the remaining parameters— C , the switching costs and ν , the 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.²⁴ 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 (m, k) individual starts in occupation o at $t = 1$, while $\lambda_t(o_t|o_{t-1}, m, k)$ is the probability that the (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 is challenging because: (1) the likelihood is not log-linear; and (2) λ 's are outcomes of the structural model solution, which requires computation of expected future continuation values.

To address the first challenge, we employ the EM algorithm, which requires defining an auxiliary function—the mean of the likelihood conditional on the distribution of k , as if it were known. The second challenge, specific to our setting, is handled following [Arcidiacono and Miller \(2011\)](#), which propose estimating λ as a parameter directly and mapping it back to the structural parameters in the second stage. This approach allows consistent estimation of the type distribution and payoffs parameters by maximizing the following function (which we continue to call the log likelihood in a slight abuse of terminology):

$$\hat{\mathcal{L}} = \sum_{i=1}^N \sum_{k=1}^K \left[\sum_{t=1}^T \log f(w_{it}|k, o_{it}, m_i, t; \alpha) + \log \hat{\lambda}_t(o_{it}|o_{i,t-1}, m_i, k) \right] q_{ik}, \quad (11)$$

²⁴In [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.

where q_{ik} is the probability that individual i is type k , and the hats atop λ 's indicate that these are estimated directly from the data rather than derived from the model. Appendix D.1 provides implementation details. Briefly, the EM algorithm iterates on initial guesses of weights q_{ik} , updating them according to Bayes' Rule, and updates parameter estimates by maximizing (11) while treating q_{ik} as given. We also discuss several technical details including the treatment of initial conditions following Wooldridge (2005), the construction of $\hat{\lambda}$, and bootstrapping procedures for computing 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, $\lambda_t(o'|o, k, m)$.²⁵ As shown in Appendix D.1.2, (6) can be used to derive the following regression equation:

$$\log \left[\frac{\lambda_t(o'|o, k, m)}{\lambda_t(o|o, k, m)} \right] + \beta \log \left[\frac{\lambda_{t+1}(o'|o', k, m)}{\lambda_{t+1}(o'|o, k, m)} \right] = - \frac{C_{oo'km}(1 - \beta)}{\nu} + \frac{1}{\nu} (w_{o't}H_o(k, m) - w_{ot}H_o(k, m)) + \xi_{oo'kmt+1}, \quad (12)$$

where $\xi_{oo'kmt+1}$ is a time $t + 1$ expectational error. Unlike (6), this formulation differences out the current and next period's continuation values, addressing the remaining selection issue. Under rational expectations, $\xi_{oo'kmt+1}$ is orthogonal to time t variables, allowing (12) to be estimated via ordinary least squares.

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. As shown in Appendix D.1.2, this leads to the following estimating equation:

$$C_{o,km}^e = \left(w_{ot}H(o, k, m) - w_{1t}H(1, k, m) \right) + \beta C_{1,o,km} - \nu \left\{ \log \frac{\lambda_t^e(o|k, m)}{\lambda_t^e(1|k, m)} + \beta \log \frac{\lambda_{t+1}(o|o, k, m)}{\lambda_{t+1}(o|1, k, m)} \right\}. \quad (13)$$

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

²⁵Even though these are generated regressors, we will avoid the cumbersome use of hat notation. Instead, it should be understood that these are estimated.

5.2 Field Choice

To describe how we estimate field choice parameters, we first describe the likelihood function conditional on types, and then discuss how we use the EM algorithm to estimate the distribution of latent types.

We begin by deriving a recursive representation of payoffs, enabling us to leverage the panel structure of our application data. First, we project the previously-recovered q_{ik} , the probability that individual i is type k , onto student observables. This recovers the field production function, $\pi(k|m, s)$. With this in hand, the payoff to a field m at time t for type l individuals is,

$$U_{tlm}(s) = \vartheta_{lm} + \theta_l \left[\sum_k \pi(k|m, s) V_{t+5}^e(k, m) \right] + v \log(N_{mt}), \quad (14)$$

It is useful to define:

$$\bar{U}_{tlm}(s) \equiv U_{tlm}(s) - v \log(N_{mt}) = \vartheta_{lm} + \theta_l \left[\sum_k \pi(k|m, s) V_{t+5}^e(k, m) \right]. \quad (15)$$

The right hand side of (15) is *not* directly observable, because V^e is never directly estimated in the labor market estimation—the strategy relied explicitly on differencing out continuation values.²⁶ To resolve this issue, we treat the utility in the last period, i.e., $\bar{U}_{Tlm}(s)$ as a parameter to be estimated, where T refers to the last period for which we observe field choices. In Appendix D.1.2, we show that for any $t = T - \tau$, (15) can be iterated forward and rewritten with the following representation:

$$\bar{U}_{T-\tau, lm}(s) = \vartheta_{lm}(1 - \beta^{T-\tau}) + \theta_l \underbrace{\left[\sum_{\tau'=T-\tau}^{T-1} \beta^{\tau'-\tau} V_{T-\tau'+5}^F(m, s) \right]}_{XV_{\tau, T}^F(m, s)} + \beta^{T-\tau} \bar{U}_{Tlm}(s), \quad (16)$$

where V^F is an entry value term that *only* depends on observable data.²⁷ For ease of notation, we collect these observable terms into a single term and denote it as $XV_{\tau, T}^F(m, s) \equiv \sum_{\tau'=T-\tau}^{T-1} \beta^{\tau'-\tau} V_{T-\tau'+5}^F(m, s)$. 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.

²⁶The key challenge of estimating V^e directly is that we do not observe the final continuation value of V . Even differences across fields, $V^e(k, m) - V^e(k, m')$ cannot be directly observed either, as there is no simple representation of the difference that does not require knowledge of some terminal value.

²⁷Specifically,

$$V_{t+5}^F(m, s) = \sum_k \pi(k|m, s) \left\{ -(1 - \beta)C_{okm}^e + w_{o, t+5}H(1, k, m) - \nu \log \lambda_{t+5}^e(o|k, m) - \beta \nu \log \left(\frac{\lambda_{t+6}(o|o, k, m)}{\lambda_{t+6}^e(o|k, m)} \right) \right\}.$$

With (16) 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, high school track, and other observables), and let (m_1, m_2, \dots, 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:

$$L_{i,T-6,1,l} = \frac{N_{m_{r_i}}^v \exp(U_{m_{r_i}T-6}(s_i))}{\sum_{m' \in \mathcal{C}_i} N_{m'}^v \exp(\bar{U}_{m'l_iT-5}(s_i))}.$$

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:

$$L_{i,T-6}(m_1, m_2, \dots, m_R, s_i) = \prod_{r=1}^R \frac{N_{i,T-5,r,m_{r_i}}^v \exp(\bar{U}_{m_{r_i}l_iT-5}(s_i))}{\sum_{m' \in \mathcal{C}_i} N_{i,T-5,r,m'}^v \exp(\bar{U}_{m'l_iT-6}(s_i))}.$$

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

$$LL = \sum_t \sum_l \sum_i \sum_r q_{il} \left(v \log N_{irm_{r_i}} + \vartheta_{lm}(1 - \beta^{T-6-t}) + \theta_l X V_{t,T-6}^F(m, s) + \beta^{T-6-t} U_{T-6,lm}(s) - \log \left(\sum_{m' \in \mathcal{C}_i} N_{irm'}^v \exp(\vartheta_{lm}(1 - \beta^{T-6-t}) + \theta_l X V_{t,T-6}^F(m, s) + \beta^{T-6-t} U_{T-6,lm}(s)) \right) \right). \quad (17)$$

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 D.1.2](#).

We recover the distribution of latent preference types l using the EM algorithm. From Bayes' rule, we update the probability that individual i is of type l conditional on a guess of parameters:

$$q_{itl}^{(g+1)} = \frac{\exp(LL_{itl}^{(g)}) q_l^{(g)}}{\sum_{l'} \exp(LL_{itl'}^{(g)}) q_{l'}^{(g)}}. \quad (18)$$

We iterate on maximizing (17) and updating (18) until parameters cease changing to tolerance.

5.3 Identification

While ultimately all parameters are jointly determined through likelihood estimation, we provide some intuition for identification. First, wage parameters are identified from the panel structure of wage data, analogous to a fixed-effects strategy. Without occupations, workers would be clustered into labor market types based on long-run average wage quantiles.²⁸ With occupations, identification relies on occupational switchers' wage patterns, grouping them by revealed comparative advantages. If workers switch occupations yet remain at stable wage ranks, they possess absolute advantages in both occupations. If their rank declines post-switch, this suggests a comparative advantage in their initial occupation. These cases are illustrated in Appendix Figure 15.²⁹

Second, switching cost parameters are identified from the regression specification linking occupational switching to changes in wage differentials. In particular, the response of *net* flows to changes in wage *differentials* identifies ν . For example, if flows between two occupations remained unchanged after the wage gap grew, this would imply a small ν . The level of *gross* flows pins down the average level of switching costs, $C_{oo'km}$. Finally, 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, then $C_{oo'}$ is estimated to be smaller than $C_{oo''}$. See Artuç et al. (2010) for more discussion of these points.

Finally, intuition for the identification of the field choice parameters can be drawn by breaking apart the steps of the EM algorithm. Given a distribution of types, θ_l is identified from changes in field choices following shifts in the observed returns to that field, XV^F . Similarly, the fixed effects, ϑ_{lm} , reflect long-term sorting patterns under type-specific choice sets.

We identify the type distribution using information from students' ranked choice lists. To illustrate, without clustering, conditionally independent preference shocks imply that second choices are independent of first choices: $P(m_{r_2}|m_{r_1}) = P(m_{r_2})$. Latent types are determined in a way that restores conditional independence within a cluster.³⁰ Ranked lists are thus *crucial*: without them, single-choice observations could not separately identify type distributions from preference shocks.

²⁸Specifically, we would fit a Gaussian mixture model, assuming each worker's draw was fixed at career start and shifted earnings in all occupations by the same amount. This approach resembles Bonhomme et al. (2019), and indeed the EM algorithm underlies practical implementations of k-means clustering.

²⁹This intuition is not quite complete: workers with similar career are also grouped together. Seo and Oh (2024) discusses how observing switching at different horizons for the same cohort identifies persistent latent heterogeneity.

³⁰See 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

To maintain a short discussion, we only focus on discussing how workers of different educational backgrounds respond to changes in earnings. Appendix A.1 contains the full set of parameter estimates, with standard errors.

Table 2 displays estimates of the labor supply parameters that govern how workers reallocate across occupations. 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.³¹ The first column is estimated by ignoring worker heterogeneity, as in Artuç et al. (2010), while the second reflects our full model, which incorporates four labor market types ($K = 4$), test scores (when available), and field of study. Controlling for heterogeneity in k and education dramatically raises the measured responsiveness of occupation switching to wage differentials, and consequently leads to much smaller estimates of switching costs.

Table 2
Labor Supply Parameters

Parameter	(1)	(2)
\bar{w}/ν	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 D.4).

We use these estimates, in conjunction with how fields vary in their returns across occupations, to summarize how graduates from different fields respond to shocks. Specifically, we focus on the semi-elasticity of exiting occupation in response to a temporary shock to wages in that occupation:

$$\frac{d(1 - \pi_{oo})}{d \log w_o} = -\frac{w_o}{\nu} \pi_{oo}(1 - \pi_{oo}). \quad (19)$$

This semi-elasticity is the percentage point increase in exit from occupation o given a temporary 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 (19) for different fields of study. We define the

³¹For each individual, we take the q_{ik} weighted average across types.

mean semi-elasticity, \mathcal{E}_m , as the mean over all workers and years in a given field.³²

Table 3
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.

Table 3 shows the value of \mathcal{E}_m for each field of study. A striking feature of Table 3 is the substantial heterogeneity in semi-elasticities across fields of study. For example, in the most responsive field—Business—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 $w_o/(1 - \beta)$.³³ In this case, a 1% permanent decline in earnings would raise the exit rate by approximately 1.7 percentage points, illustrating the more substantial impact of persistent wage changes.

³²Formally, let

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

where q_{ik} is the probability a worker is type k , w_{imkt} is the wage, $\hat{\lambda}$ are transition probabilities, while the indexing $o(i)$ refers to i 's occupation. We define the mean semi-elasticity as,

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

³³Technically, 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.

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.³⁴ For our estimation, we set $L = 4$, which strikes a balance between computational efficiency and capturing meaningful heterogeneity in preferences and decision patterns.

Table 4
Field Choice Parameters

Parameter	Type 1	Type 2	Type 3	Type 4
θ	0.030 (0.001)	0.144 (0.022)	0.052 (0.004)	0.031 (0.001)
Q_l	0.183	0.197	0.301	0.320
\mathcal{E}^f	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 (14), 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, which is the mean of (20) across individuals.

Turning to the first set of results, Table 4 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 for each student i :

$$\mathcal{E}_{il}^f \equiv \frac{d\pi_{m(i)sl}}{d \log V_m^e} = \pi_{im(i)l}(1 - \pi_{im(i)l}), \quad (20)$$

where $m(i)$ is the students' field choice.³⁵ Since V^e is measured in differences rather than levels 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.³⁶ 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 . 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 4 highlights two key findings. First, students appear relatively responsive to shocks. The

³⁴The complete set of parameter estimates is provided in Appendix A.1.

³⁵This value also varies across observables like scores and high school track, which we omit to keep notation minimal.

³⁶Our 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 D.5 provides details on our approximation of V^e , including how we take averages over time.

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 fields 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 reallocation occurs within broad categories such as STEM.³⁷

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 suggests that the composition of affected types may influence how the economy adjusts to shocks, with potentially different implications depending on which groups of most exposed.

Beyond differences in sensitivity, we also find notable variation in preferences across types. [Table 5](#) 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.³⁸ Lastly, Type 4 corresponds to a humanities/social science profile, with a clear preference for those fields.

Finally, we examine aggregate substitution patterns across fields of study. To do this, we

³⁷Interestingly, the differences are not 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 obtain similar mean semi-elasticities, though substitution patterns differ. By contrast, accounting for choice sets matters greatly: assuming full access to all fields significantly lowers estimates. Ignoring all observable heterogeneity yields near-zero estimates.

³⁸There are two key reasons for this pattern. First, STEM is not a particularly popular choice overall, as reflected in the last column of [Table 5](#). 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 5
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.

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 logit model without heterogeneity, it is well known that these substitution patterns are independent of any parameters, and are instead proportional to initial conditional probabilities:

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

However, with agent heterogeneity, this relationship no longer holds. Instead, they 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 6 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

Substitute Field	Education	Medium Health	Care	Arts	Humanities	Communication	Social Sciences	Business Admin	Law	Life Sciences	Physics/Math	ICT	Engineering	Architecture	High Health
Education	12.11	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	63.93	17.40	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	1.24	1.04	2.43	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	5.58	2.27	4.43	20.17	7.73	4.24	7.36	1.96	3.18	0.28	0.49	1.83	0.65	1.32	0.13
Humanities	4.48	3.97	5.47	18.27	21.36	15.14	35.76	4.28	13.47	0.92	1.73	1.86	1.29	2.75	1.41
Communication	3.99	3.78	3.39	25.68	39.49	14.65	16.53	24.37	13.92	3.68	7.39	5.02	7.34	6.45	1.21
Social Science	1.36	2.42	1.73	11.02	9.76	33.95	10.64	11.68	25.09	2.58	3.66	5.66	3.36	7.22	6.57
Business Admin	0.76	1.82	0.91	6.02	8.69	7.97	12.80	12.19	23.27	9.31	25.28	21.25	26.19	22.26	1.50
Law	0.76	12.03	1.20	0.91	0.83	4.99	0.97	6.66	3.35	2.61	3.22	5.05	5.18	6.38	2.19
Life Sciences	0.40	1.17	0.63	0.58	0.92	2.23	1.42	7.09	2.01	5.00	11.30	8.59	14.52	9.14	29.61
Physics/Math	0.10	0.11	0.21	1.47	0.71	1.35	0.70	5.43	1.41	1.59	4.92	6.49	11.46	6.28	4.05
ICT	0.55	0.83	0.71	0.83	0.74	2.16	1.17	11.86	3.47	10.52	26.76	26.85	9.03	7.69	0.22
Engineering	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	25.21	2.91
Architecture	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	0.76
High Health															

Table 6
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.

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. This suggests that the model effectively captures realistic clustering patterns.³⁹ Appendix Table 11 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 6 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

In this section we close the model by specifying the labor demand side, and then describe the two counterfactual simulations. First, we assess the steady state consequences of relaxing capacity constraints in higher education, examining how expanded access influences labor market outcomes. Second, we simulate different labor market shocks, assess the role of field-of-study in the effects of these shocks, and evaluate the impacts across different education policies.

7.1 Closing the Model

We first outline the key components of goods markets and labor demand. Full details, including equilibrium conditions, are provided in Appendix E.1. In what follows, let $s \in \mathcal{S}$ index industries, and denote Denmark and Foreign by D and F , respectively.

³⁹Relatedly, 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.

The economy features a final non-traded consumption good. This final good is a nested CES aggregate of industrial outputs, combining both domestic and foreign varieties. At the top level, the final good aggregator, Y_t , is defined as:

$$Y_t = \left(\sum_s \mu_s^{\frac{1}{\sigma}} Y_{st}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (21)$$

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:

$$Y_{st} = \left((Y_{st}^D)^{\frac{\eta-1}{\eta}} + (Y_{st}^F)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \quad (22)$$

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:

$$X_{st} = A_{st}^F \left(P_{st}^D \right)^{-\eta}, \quad (23)$$

where A_{st}^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:

$$Q_{st} = A_{st} \left[\left(A_s^N \right)^{\frac{1}{\alpha}} \left(L_{st}^N \right)^{\frac{\alpha-1}{\alpha}} + \left(\left(A_s^P \right)^{\frac{1}{\rho}} \left(L_{st}^P \right)^{\frac{\rho-1}{\rho}} + \left(A_s^M \right)^{\frac{1}{\rho}} \left(Q_{st}^M \right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1} \frac{\alpha-1}{\alpha}} \right]^{\frac{\alpha}{\alpha-1}}, \quad (24)$$

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:

$$L_{st}^f = \prod_{o \in \mathcal{O}^f} \left(\frac{L_{sot}^f}{\beta_{so}} \right)^{\beta_{so}}, \quad (25)$$

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 E.1.


7.2 Capacity Constraints in General Equilibrium

In our first counterfactual, we examine how Danish capacity constraints and high school tracking affect workers career choices and aggregate output. As discussed in Section 2.3, students’ choice sets are limited in two ways: they may not meet prerequisite requirements for applying to a program (proxied by high school tracks), and if their final score is below some programs’ thresholds (capacity constraints). Starting from a calibrated 2018 steady-state equilibrium, we explore two scenarios.⁴⁰ 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”), where students are free to enter any field. Both scenarios align with relevant policies—for instance, Denmark’s expanded supplemental coursework (*Gymnasiale suppleringskurser*) to increase student flexibility.⁴¹

⁴⁰Fully implementing the Deferred Acceptance mechanism is challenging in our setting—not because of the computational cost of the algorithm, but because we do not explicitly model program acceptance conditional on field choices. To address this in a tradable and empirically grounded manner, we instead model admissions probabilistically. That is to say, a student with a test score of s from background h has probability $P(m|h, s)$ of being admitted to field m . These probabilities are recovered from historical acceptance rates and held fixed in counterfactuals. Removing capacity constraints amounts to setting these probabilities to 1.

⁴¹Two 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 fields (Ahn et al., 2024; Arcidiacono et al., 2025)—albeit this concern is partially alleviated by controlling for high school performance. In Appendix E.3, we offer one further scenario—we maintain cutoffs and tracking specifically for medicine, a field that attracts many students but our model may be ill-suited for.

Figure 5
Field Reallocation Without Capacity Constraints



Notes: Plots (100x) changes in the log shares of each field, $\Delta \log s_f$, completed by new cohorts. Panel (a) plots these changes removing capacity constraints but maintaining high school tracking, while panel (b) plots these changes removing high school tracking.

Figure 5 shows significant aggregate shifts across fields. To assess the extent of reallocation, we simulate individuals in both steady state and under each counterfactual policy. We then compute the proportion of simulated individuals who switch field from steady state. Reallocation is substantial, with 9% of students and nearly 20% of students switching fields under Partial Flex and Full Flex, respectively. Appendix Table 12 further 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 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.

We acknowledge that policy makers designing educational systems face a potential tension between respecting student preferences and meeting national objectives. We are agnostic on the source of these objectives.⁴² Nevertheless, our model can evaluate the economic and distributional consequences of removing capacity constraints, important inputs into any welfare function.

Table 7 summarizes further findings. Removing only capacity constraints (Partial Flex) raises GDP by 0.31%, whereas removing tracking and constraints (Full Flex) boosts GDP by 1.30%. Although both effects are positive, the magnitudes are small. One explanation for these responses, despite substantial field-of-study reallocation, is the more limited occupational reallocation, as

⁴²As an example, the objective could be paternalistic if student preferences reflect incomplete information (Hastings et al., 2016). Alternatively, fiscal considerations can militate for pushing students to higher output fields.

Table 7
GE Consequences of Lifting Capacity Constraints

	Partial Flex	Full Flex
%Δ GDP	0.31	1.30
%Δ s.d. log Income	1.07	1.62
Field Reallocation	8.66	19.77
Occupational Reallocation	6.90	16.33

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 the proportion of simulated agents switching between counterfactual and steady state.

shown in the fourth row of Table 7.⁴³ This suggests that when students switch fields, they tend to move into fields that offer similar career outcomes.

To examine distributional outcomes, we simulate a large cohort of 200,000 individuals from an initial choice of field through their career. For each individual we compute discounted lifetime earnings as,

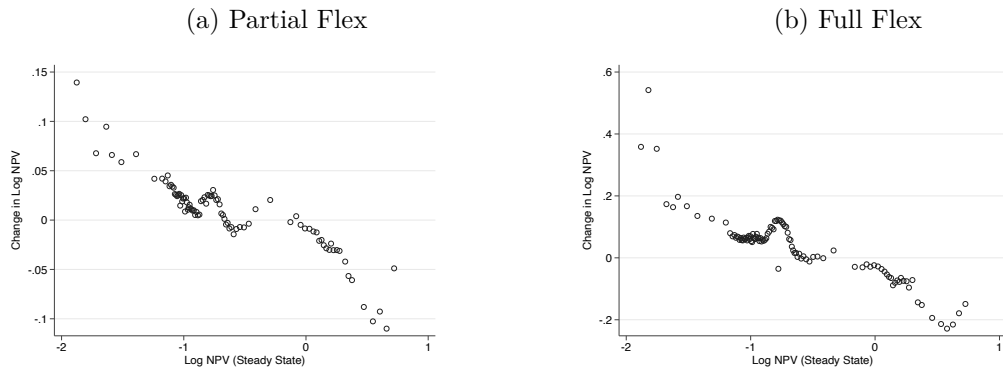
$$NPV_i = (1 - \beta) \times \sum_{t=0}^{50} \beta^t w_{it}. \quad (26)$$

As shown in the second row of Table 7, more flexibility increases inequality. To make sense of this result 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. In full GE, there are additional feedback effects caused by wage adjustments in response to reallocation.

Focusing on standard deviation may obscure richer distributional consequences. Figure 6 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 somewhat for top earners. This convergence effect is especially pronounced in the Full Flex scenario,

⁴³We calculate occupational reallocation analogously to field reallocation.

Figure 6
Removing Capacity Constraints and Inequality



Notes: For a simulation of $N = 200,000$ individuals, plots the change in steady state lifetime earnings, $(1 - \beta) \frac{1}{50} \sum_{t=1}^{50} \beta^t w_{it}$, across steady states, against the steady state lifetime earnings in the original steady state.

dominating overall distributional outcomes.⁴⁴

In sum, relaxing capacity constraints improves students' welfare, and has small but positive effects on aggregate output. Nevertheless, caution is warranted; if tracking reflects inherent student abilities, our Full Flex scenario would be too blithe. Moreover, relative to work such as Fu (2014), we do not have a model of the education supply side—achieving full flexibility in practice may incur considerable costs.

7.3 Labor Market Shocks

Finally, we examine two labor market shocks: an international trade shock that increases the cost of foreign goods and reduces export demand, and an AI-driven shock that alters demand for certain skills. We analyze their impacts on income and inequality across policy scenarios. We focus our analysis on college workers but briefly discuss aggregate outcomes at the end of this section.

Trade War Shock


In the first analysis, we examine the labor market effects of a simulated trade shock motivated by ongoing global trade tensions. Drawing on renewed protectionist policies and retaliatory measures, we assume a 30% increase in the cost of imported goods and offshored inputs, together with a 30% reduction in foreign demand for Danish exports. Starting from the calibrated 2018 steady-state equilibrium, we then assess the long-run impact of this dual shock on the domestic labor market.

⁴⁴The standard deviation of log income captures overall dispersion in ranks, not changes in relative rank. The standard deviation may still increase, despite the convergence across ranks, if the magnitude of gains and losses is large or highly variable across individuals.

Although the shock is applied uniformly, its effects vary widely across occupations. Production-related occupations, such as Science Technicians, Agriculture, and Engineering, and those linked to tradable services, including Administration, Accounting, and Business, face the highest exposure (Appendix Figure 17a).⁴⁵ At the education level, the most exposed fields include Physics/Math, Law, Life Sciences, and Business Admin (Appendix Figure 18a).⁴⁶

The trade war shock primarily affects labor demand, but the resulting equilibrium outcomes also depend critically on supply-side responses, particularly the ability for students to reallocate across fields of study. To assess the role of this flexibility, we contrast our baseline *Current System*, in which roughly 2.3% of students adjust fields within existing capacity constraints, with two alternative scenarios. Under the *No Reallocation* scenario, field choices remain fixed at initial steady-state shares, shutting down any educational adjustment. At the other extreme, the *Full Flex* scenario eliminates high school tracks and capacity constraints entirely, as discussed in the previous section, substantially increasing field switching to approximately 26%.

Figure 7
Field Reallocation under Trade War Shock



Notes: Plots (100x) changes in the log shares of each field, $\Delta \log s_f$, completed by new cohorts. Panel (a) plots these changes under the current system, while panel (b) plots these changes removing all capacity constraints.

Figure 7 highlights how both the scale and direction of student reallocation vary across scenarios.

Under the *Current System*, switching remains limited and tends to cluster within groups of fields

⁴⁵Occupational exposure is calculated by $Exposure_o = \sum_i \left(\frac{X_i}{WB_i} - \frac{M_i}{WB_o} \right) \frac{WB_{oi}}{WB_o}$, where X_i is the value of export, M_i is the combined value of the foreign inputs and imported goods, WB_i , WB_o and WB_{oi} are the wage bill at the sector, occupation, and sector-occupation level, respectively.

⁴⁶The variation in field-level exposures is a weighted average of occupation-level exposures with weights given by the share of students from each field working in each occupation. Physics/Math, for instance, sees high exposure due to a large share of graduates entering Science Technician and Engineering occupations, which are heavily impacted by the shock. Similarly, Law and Business Admin have high exposure because many of their graduates work in highly tradable service occupations such as Accounting and Administration. Life Sciences is affected both through scientific support roles and manufacturing-linked health sectors.

that reflect students’ estimated preference structures, such as movement within STEM or healthcare fields. In contrast, the *Full Flex* scenario generates broader reallocation patterns, which are also quantitatively larger. There are substantial outflows from fields like Physics/Math, ICT, and Engineering, and increased inflows into fields such as High Health, Arts, and Social Sciences. While these shifts partly reflect underlying exposure differences, they do not map cleanly onto exposure rankings. Instead, they combine both exposure considerations and students’ relative preference across fields, consistent with our model that field choices are shaped by a combination of economic incentives and non-pecuniary preferences.

Table 8
Results across Policy Scenarios: International Trade Shock vs. AI Shock

	Panel A: Trade War 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)
% Δ Income ($\Delta < 0$)	-0.73 (55.0%)	0.48 (56.7%)	9.20 (55.8%)	13.3 (40.3%)	15.4 (40.0%)	24.6 (41.7%)
Field Reallocation	0%	2.3%	26.1%	0%	3.5%	26.7%
Occupation Reallocation	12.1%	13.1%	25.2%	16.9%	18.3%	29.4%
% Δ s.d. log Income	18.3	18.5	20.4	15.4	15.7	17.7

Notes: Table reports the changes for row variables between steady state after a policy simulation and the initial steady state. % Δ Income is the percentage change in NPV between the counterfactual and baseline scenarios, with NPV calculated using (26). % Δ s.d. log Income is the percentage change in the standard deviations of log income for the 200,000 simulated individuals between the counterfactual and baseline scenarios. Occupation and Field reallocation reports the shares of the 200,000 simulated individuals who switch occupations and fields, respectively, between the counterfactual and baseline scenarios.


We next quantify how differences in educational flexibility shape economic outcomes under the trade war shock by simulating steady-state effects for 200,000 college-educated individuals. As shown in Panel A of Table 8, preventing educational adjustment leads to a decline in average lifetime income (-0.73%), while allowing even limited field switching under the *Current System* reverses this loss, resulting in a small gain of 0.48%. Removing all constraints in the *Full Flex* scenario substantially improves outcomes, with income rising by 9.2%.⁴⁷ Greater flexibility also increases occupational mobility significantly, with occupation switching rising from 12.1% under *No Reallocation* to 13.1% under the *Current System*, and to 25.2% under *Full Flex*. However, these gains come alongside modestly increased inequality. The standard deviation of log income rises slightly with greater flexibility, reaching 20.4% under *Full Flex*, compared to 18.3% under *No Reallocation* and 18.5% under the *Current System*. In addition, a larger share of individuals experience declines

⁴⁷All scenarios begin from the same initial field allocation; however, steady-state allocations may differ under greater flexibility due to endogenous re-sorting. We explore this in Section 7.2. Also, Appendix E.6 presents an alternative exercise in which the initial allocation is set to the steady state under the Full Flex scenario.

in lifetime income as flexibility increases.

While these summary statistics capture overall dispersion, they may mask important distributional shifts. Figure 8 plots lifetime income changes against baseline income levels. Under both *No Reallocation* and the *Current System* (Panels a and b), the most negative income growth is concentrated among individuals at the lower end of the initial income distribution, while higher-income individuals tend to experience moderate gains. In contrast, the pattern is reversed under *Full Flex* (Panel c): the largest positive gains are now concentrated among individuals with the lowest initial incomes, whereas those starting at the top of the distribution are more likely to experience income losses. Intuitively, capacity constraints tend to benefit students with high earnings potential, as it is strongly correlated with test scores. Relaxing these constraints enables lower-scoring students to re-optimize more effectively in response to shocks.

Figure 8
Trade War Shock and Inequality



Notes: For a simulation of $N = 200,000$ individuals, plots the change in steady state lifetime earnings, $(1 - \beta) \frac{1}{50} \sum_{t=1}^{50} \beta^t w_{it}$, across steady states, against the steady state lifetime earnings in the original steady state.

AI Shock

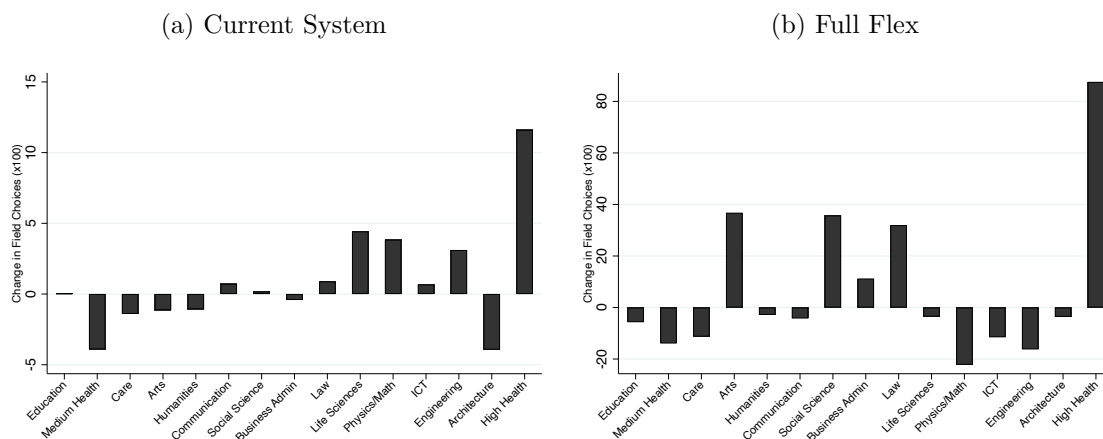
We now turn to the labor market implications of a technological shock driven by the adoption of artificial intelligence (AI), specifically large language models (LLMs). Starting from the same calibrated 2018 steady-state equilibrium, we simulate occupation-specific productivity changes that reflect differing levels of AI exposure.⁴⁸ While the AI shock differs fundamentally from the trade war shock, exposure similarly varies significantly across occupations: those involving routine cognitive tasks, such as Accounting, Secretarial Work, and Administration, face the highest exposure (Appendix Figure 17b). The exposure to AI shocks, as shown in Appendix Figure 18b, is relatively uniform across educational fields, with Business Administration and Law experiencing somewhat

⁴⁸Occupational exposure to AI is based on estimates from Eloundou et al. (2023).

higher exposure compared to other fields.

As with the trade shock, equilibrium outcomes under AI depend not only on shifts in labor demand but also on supply-side flexibility through occupational and educational adjustments. We again compare three policy scenarios: the baseline *Current System*, in which 3.5% of students change fields within existing capacity limits; the *No Reallocation* case, which fixes field choices at their initial steady-state levels; and the *Full Flex* scenario, where all constraints are removed, prompting reallocation by 26.7% of students (Panel B of Table 8). Under the Current System (Figure 9a), adjustments are again limited and primarily occur within closely related fields, such as among STEM disciplines or healthcare-related areas. The Full Flex scenario shown in Figure 9b exhibits quantitatively larger and more broad reallocation across fields. Fields such as Physics/Math and Engineering see substantial reductions in enrollment, while High Health, Arts, Social Sciences, and Law experience significant increases. Since the field-specific exposures are relatively uniform, most of the reallocation is driven by preferences.

Figure 9
Field Reallocation under AI Shocks




Notes: Plots (100x) changes in the log shares of each field, $\Delta \log s_f$, completed by new cohorts. Panel (a) plots these changes under the current system, while panel (b) plots these changes removing all capacity constraints.

Economic consequences across scenarios are summarized in Panel B of Table 8. Although the nature of the shock differs fundamentally from the trade war scenario, the patterns of adjustment and outcomes are strikingly similar. Greater educational flexibility again substantially enhances individual gains. Allowing limited field reallocation under the Current System increases lifetime income gains by roughly 16% (from 13.3% under No Reallocation to 15.4%), with gains reaching 24.6% under Full Flex. Occupational switching also rises markedly, from 16.9% without field adjustments to 18.3% and 29.4% under Current System and Full Flex scenarios, respectively.

Income inequality, measured by the standard deviation of log income, and the share of individuals experiencing negative income growth both increase somewhat under greater flexibility, particularly in the Full Flex scenario.

Finally, distributional impacts mirror those observed earlier under trade war shock. As shown in Figure 10, under both *No Reallocation* and the *Current System*, negative income growth is concentrated among individuals at the lower end of the initial income distribution, while middle- and higher-income individuals tend to see larger gains. This pattern reverses under *Full Flex* (Panel c), where the largest gains are concentrated among initially low-income individuals, and income losses shift toward the top of the distribution.

Figure 10
AI Shocks and Inequality



Notes: For a simulation of $N = 200,000$ individuals, plots the change in steady state lifetime earnings, $(1 - \beta) \frac{1}{50} \sum_{t=1}^{50} \beta^t w_{it}$, across steady states, against the steady state lifetime earnings in the original steady state.

Key takeaways: Greater flexibility in field choices consistently improves individuals’ ability to adapt to labor market shocks, whether from trade disruptions or technological change. Even modest reallocation yields better outcomes, while fully removing the constraints enables stronger responses to shifting demand. Although flexibility slightly increases overall income inequality, it also shifts the distribution of gains: income losses become more concentrated among initially high earners, while lower-income individuals experience larger improvements. These findings underscore the central role of field choice in reallocating talent during structural changes. Flexibility not only raises average income but also facilitates convergence in income trajectories, enhancing resilience both at the aggregate level and across the distribution.

One caveat is that these outcomes reported here apply specifically to college-educated individuals. In aggregate terms, differences in GDP across policy scenarios are much smaller.⁴⁹ The relatively

⁴⁹For the trade war shock, GDP changes by -14.0% (No Reallocation), -13.9% (Current System), and -13.4% (Full Flex). For AI Shocks, GDP increases by 28.2%, 28.6%, and 29.0%, respectively.

modest changes in GDP and the larger income changes for the college-educated imply offsetting losses for workers without college education. While our analysis emphasizes the value of flexibility within higher education, it also points to the importance of broader policies that support adaptability beyond the college-educated workforce.

8 Conclusion

In this paper, we model and estimate how students choose their field of study and eventual career paths. 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. Exploiting rich administrative data from Denmark, we are able to capture latent heterogeneity governing both preferences over fields-of-study, as well as unobserved heterogeneity in workers' comparative advantages across occupations.

Our key empirical findings are threefold: first, fields of study endow students with skills that differ in their generality across occupations; second, students are responsive to future income shocks that differentially impact fields, but nevertheless tend to cluster their decisions among narrow bands of related fields; and third, modeling heterogeneity is crucial to fitting realistic substitution patterns of students in response to shocks.

We run two sets of counterfactuals. The first finds that constraints faced by high school graduates are binding and important. We examine both capacity constraints on certain fields, as well as choice set limitations due to high school tracking. Interestingly, we find that offering students more flexibility may increase output, with muted effects on inequality. This is because constraints tend to benefit students with the highest earning potential by keeping out other students. Essentially, for lower earnings potential students, it is better to be at the bottom of high earning fields, than in lower earnings fields.

In the second, we use our model to understand how shocks unfold when students can shift their fields, while incumbents shift their occupations. We find greater flexibility in field-of-study choices leads to improved individual outcomes, higher occupational mobility, and a redistribution of gains toward lower-income individuals. Even modest reallocation under the current system enhances resilience, while removing all capacity constraints further amplifies these effects. In other words, flexibility consistently supports better adjustment to labor market disruptions and reduces the burden on more vulnerable groups.

Our findings have important policy implications. Reforms that improve students' responsiveness to changing labor market conditions may not only boost aggregate economic outcomes, but also promote a more dynamic and resilient workforce. Our paper motivates several avenues of future research. 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, as is the decision to enter college at all. We have also abstracted from institution quality, and other factors that might matter for students' decisions, such as their parents' backgrounds. A huge determinant of students' choices remain the field-specific preference parameters. Opening up this black box could shed light on promising policy interventions. Finally, we have completely abstracted from the education supply side—understanding the costs of more flexibility is crucial for proper cost-benefit analysis.

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.
- Ahn, Tom, Peter Arcidiacono, Amy Hopson, and James Thomas**, “Equilibrium Grading Policies With Implications for Female Interest in STEM Courses,” *Econometrica*, 2024, *92* (3), 849–880.
- 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.
- 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 –, “College Major Restrictions and Student Stratification,” Technical Report, CESifo Working Paper Series No. 11601 2024.
- Bonhomme, Stephane, Thibaut Lamadon, and Elena Manresa**, “A Distributional Framework for Matched Employer Employee Data,” *Econometrica*, 2019, 87 (3), 699–739.
- Bordon, Paola and Chao Fu**, “College-Major Choice to College-Then-Major Choice,” *The Review of Economic Studies*, 05 2015, 82 (4), 1247–1288.
- 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.
- 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.
- 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.
- Eckardt, Dita**, “Training Specificity and Occupational Mobility: Evidence from German Apprenticeships,” 2024.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock**, “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” 2023.
- 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.
- Ferriere, Alex, Ricardo Reyes-Heroles, and Gaston Navarro**, “Escaping the Losses from Trade: The Impact of Heterogeneity and Skill Acquisition,” Technical Report 2021.
- Fu, Chao**, “Equilibrium Tuition, Applications, Admissions, and Enrollment in the College Market,” *Journal of Political Economy*, 2014, 122 (2), 225–281.
- 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.
- , – , **Anely Ramirez, and Seth D. Zimmerman**, “(Un)informed college and major choice: Evidence from linked survey and administrative data,” *Economics of Education Review*, 2016, 51, 136–151. Access to Higher Education.

- 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.
- Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow**, “The Allocation of Talent and U.S. Economic Growth,” *Econometrica*, 2019, *87* (5), 1439–1474.
- Humlum, Anders**, “Robot Adoption and Labor Market Dynamics,” Working Paper, University of Chicago Booth School of Business 2021.
- , **Jakob R. Munch, and Pernille Plato**, “Changing Tracks: Human Capital Investment after Loss of Ability,” *American Economic Review*. Forthcoming.
- 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.
- Lufade, 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.
- Ransom, Tyler**, “Selective Migration, Occupational Choice, and the Wage Returns to College Majors,” *Annals of Economics and Statistics*, 2021, (142), 45–110.
- 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.
- Seo, Jaeun and Ryungha Oh**, “Sectoral Shocks and Labor Market Dynamics: A Sufficient Statistics Approach,” 2024. Working paper.
- 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.
- vom Lehn, Christian, Cache Ellsworth, and Zachary Kroff**, “Reconciling Occupational Mobility in the Current Population Survey,” *Journal of Labor Economics*, 2022, *40* (4), 1005–1051.
- 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 9
Field of Study Classification

Field	Field Name	ISCED-F Code	College	Notes
1	Education	11	✓	
2	Arts	21	✓	
3	Humanities	22, 23	✓	
4	Social Sciences	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 short-cycle education programs (ISCED 5). Checkmarks denote the fields of study for college education. See Section 2.2 for relevant discussions.

Table 10
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 Labor 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. See Section 2.2 for relevant discussions.

Table 11
 Substitution Matrix $\left(\frac{d\pi_{m'}/d\log V_m}{d(1-\pi_m)/d\log V_m}\right)$

Substitute Field	Initial Field														
	<i>Education</i>	<i>Medium Health</i>	<i>Care</i>	<i>Arts</i>	<i>Humanities</i>	<i>Communication</i>	<i>Social Sciences</i>	<i>Business Admin</i>	<i>Law</i>	<i>Life Sciences</i>	<i>Physics/Math</i>	<i>ICT</i>	<i>Engineering</i>	<i>Architecture</i>	<i>High Health</i>
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	.

Notes: This table 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. See Section 6.2 for relevant discussions.

Table 12
 Reallocation Table: Removing Capacity Constraints

Counterfactual Field	Initial Field														
	Education	Medium Health	Care	Arts	Humanities	Communication	Social Sciences	Business Admin	Law	Life Sciences	Physics/Math	ICT	Engineering	Architecture	High Health
Education		0.06	11.93	0.31		0.12	0.20			0.62	0.57	0.10	1.27		
Medium Health	6.33		3.71			0.36				7.86	1.51	0.30	1.10		37.04
Care	3.51	0.06							0.24	0.90	0.19	0.39	0.28	0.30	
Arts	13.01	0.12	4.61		6.56	3.55	5.12	4.20	10.58	2.27	4.15	3.55	4.31	5.07	3.70
Humanities	8.26	0.34	2.21	2.78		10.07	1.02	7.32	9.62	6.06	8.58	7.69	8.23	5.67	
Communication	2.38	0.06	4.00	0.93	0.15		0.41	0.29	1.68	2.00	2.07	1.68	2.65	0.60	3.70
Social Sciences	34.16	3.08	6.43	34.26	37.89	35.19		38.77	56.97	22.67	31.48	31.56	28.33	39.10	14.81
Business Admin	3.17	0.12	10.24	1.54	2.01	4.74	1.23		6.25	4.55	5.47	4.14	5.91	3.58	7.41
Law	13.12	1.37	15.48	52.16	46.22	34.24	82.58	41.50		25.02	29.03	36.69	24.57	37.31	29.63
Life Sciences	1.24	0.65	8.83	0.62	0.69	0.83	0.61	0.49	2.40		3.86	2.96	7.45	1.49	
Physics/Math	0.57	0.06	1.73	0.62	0.39	0.24	0.61	0.29	0.24	0.28		0.30	0.72		
ICT	0.57		5.73	1.23	0.93	1.18	0.41	0.88	1.44	1.45	2.54		2.15	0.30	
Engineering	0.34	0.03	1.60		0.15	0.12	0.41	0.29	0.72	1.31	1.60	0.79			3.70
Architecture	2.38	0.12	4.80	1.85	1.08	2.61	1.43	1.07	2.88	2.41	2.73	3.85	3.53		
High Health	10.97	93.90	18.71	3.70	3.94	6.75	5.94	4.88	6.97	22.61	6.22	6.02	9.50	6.57	

Notes: Each column is a base field. Each row is share of students, conditional on switching, who move to field f' from the base field, f , following a change in policy, $s_{ff'}/(1-s_{ff})$. Switchers are based on simulating individuals with the same characteristics and idiosyncratic shocks and computing optimal choices under each policy. See Section 7.2 for relevant discussions.

A.1 Complete Parameter Estimates

Table 13

α_k

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

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

Table 15
 ϑ_{lm}

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


Notes: Displays the type-field preference fixed effect.

Table 16
 Additional Parameters

Parameter	Estimate
v	1.363 (0.008)

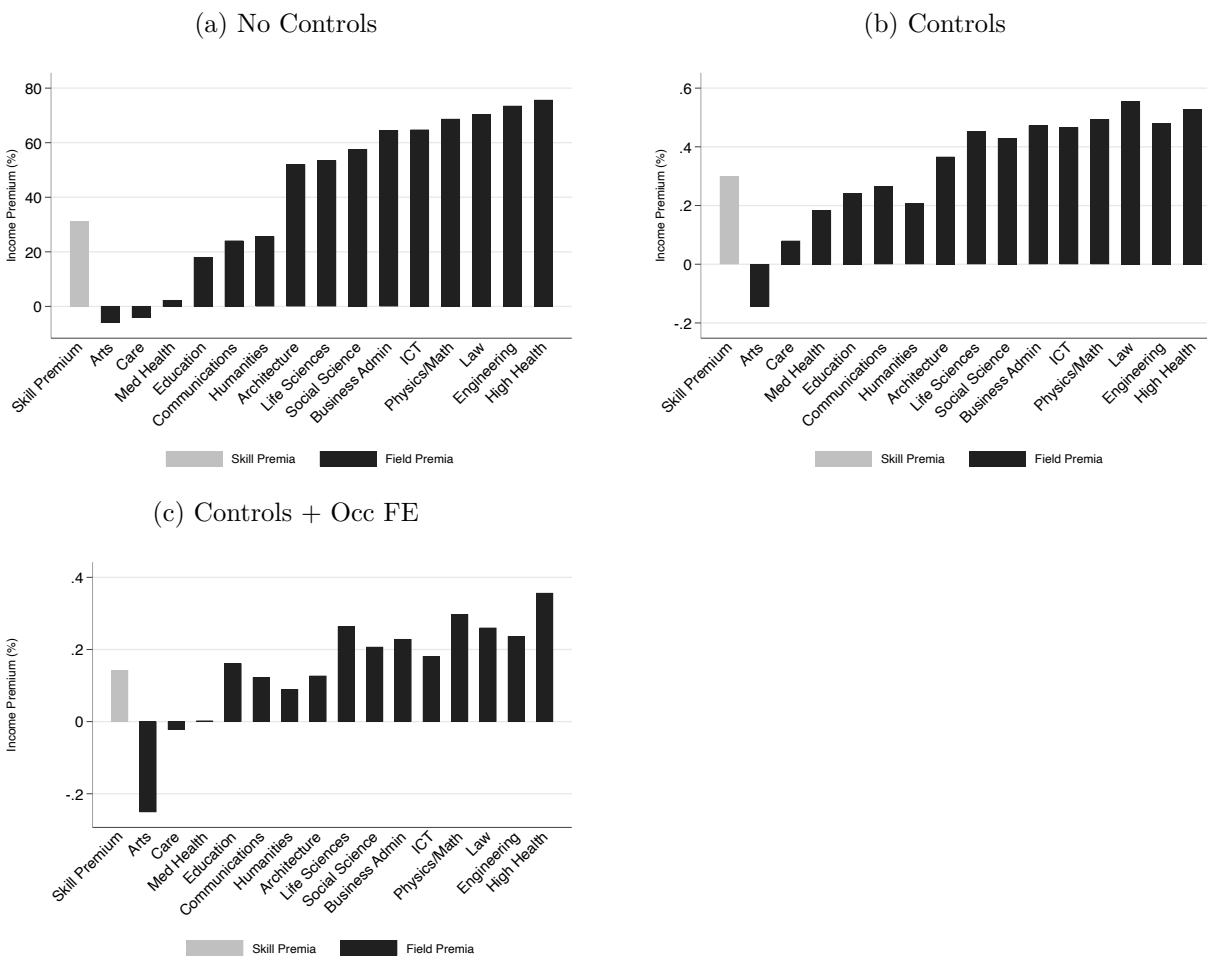
B Additional Figures

Figure 11
Education Level of New Cohorts at Labor Market Entry



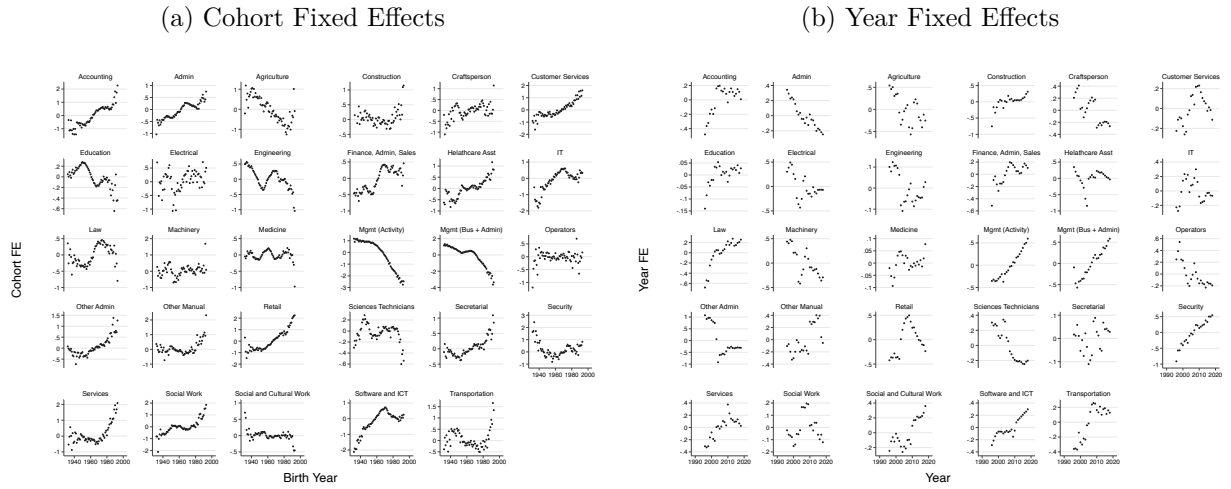
Notes: This figure plots the times series of the shares of highest attained level of education across workers in their first year on the labor market. Year of labor market entry is defined as the first year of employment after the worker achieves their highest level of education (i.e., work during schooling is not counted). See Section 2.3 for relevant discussions.

Figure 12
College Premia and Field Premia



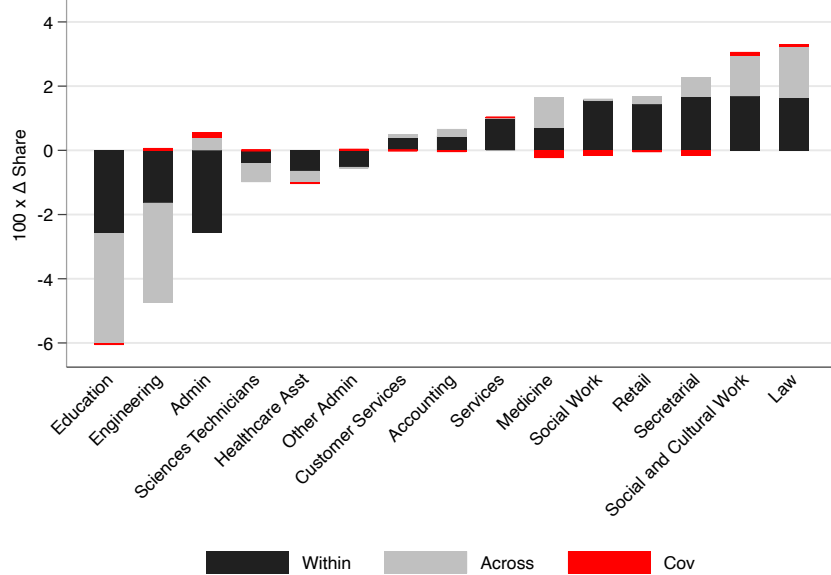
Notes: In panel (a), premia are calculated as average earnings college- or field-educated employed individuals in field f relative to those with no college, $\frac{\overline{Earn}_f}{\overline{Earn}_{No\ College}}$. In panels (b) and (c), college- or field- fixed effects are plotted from a regression of log earnings on a quadratic function in age, year fixed effects, and gender. In panel (c), occupation fixed effects are also included. In all panels, data is pooled from 1996–2018 and conditioned on being strictly positive. See Section 3 for relevant discussions.

Figure 13
Reallocation Decomposition



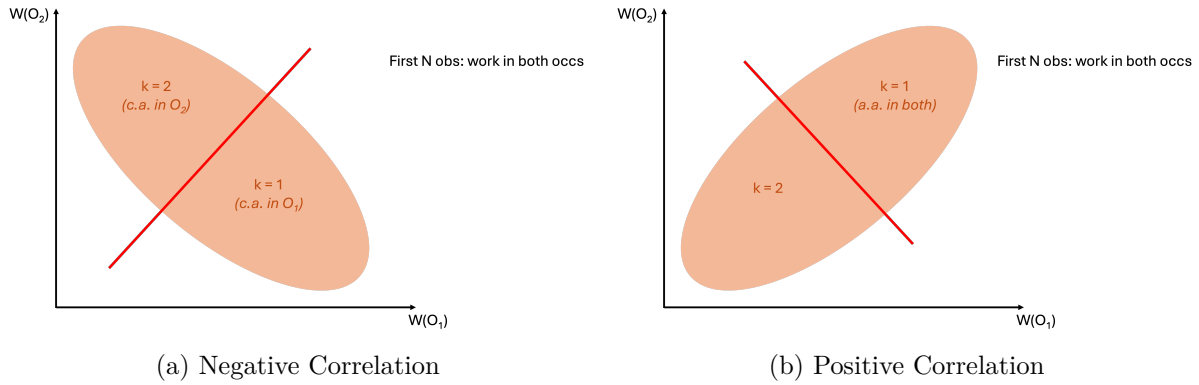
Notes: The figure plots the estimated cohort and year fixed effects from a projection of these variables onto log occupation shares (see Equation (1)). Plots are all MA3 smoothed series. The null hypothesis of perfect labor market fluidity is a horizontal line at 0 for all cohort plots. See Section 3 for relevant discussions.

Figure 14
Decomposition for All Occupations



Notes: This figure plots the decomposition of occupational reallocation into across- and within- field components, and their covariance according to the Kitagawa-Oaxaca-Blinder decomposition described in Equation (2). Occupations are sorted based on total growth over the sample period. See Section 3 for relevant discussions.

Figure 15
Intuition for the EM Algorithm



Notes: Panels (a) and (b) plot theoretical and appropriately time-discounted wages in occupation 1 (x-axis) against wages in occupations 2 (y-axis) for workers switching occupations. In panel (a), workers' rank changes with occupation switching, indicating likely differences in comparative advantages between the labor market types. In panel (b), workers switch occupations yet remain at stable wage ranks, which suggest one type possess absolute advantages in both occupations over the other type. See further discussions in 5.3.

C Data Appendix

C.1 More on Denmark's College System and High School Tracks

Denmark's college 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, including Law and Medicine.

The pathway to higher education begins with upper secondary schooling, which plays a critical role in determining students' higher education options. The 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 large 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-20s⁵⁰.

C.2 Classifying Field of Study

In this section we discuss the classification of fields of study and how we constructed our concordance. Statistics Denmark 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/discsed15-udd#>
 - Specifically see the Subject Area variant.

UDD also provides a link to HOVEDOMRAADE, which can be mapped to ISCED education *level* codes. We base our classification on how Denmark classifies its education programs in EURYDICE—the education research arm of the European Community.⁵¹ This classification is below:

- HOVEDOMRAADE < 30 is level 1 (high-school education)⁵²
- HOVEDOMRAADE = 30, 40 is level 2 (short-cycle education) — ISCED 5
- HOVEDOMRAADE = 50+ is level 3 (college and above) — 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 versus those in the professional bachelor’s system. This is because the former tends to be medicine/dentistry programs, and the latter nursing programs. Finally, we pool all defense/services sectors into

⁵⁰See <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

⁵¹See <https://eurydice.eacea.ec.europa.eu/national-education-systems/denmark/first-cycle-programmes>.

⁵²In Denmark, upper secondary education includes both general upper secondary education (Gymnasiale Uddannelser) and vocational basic courses (Erhvervsfaglige Grundforløb). The former is academically oriented, while the latter focuses on practical, job-specific training.

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)
- 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

C.2.1 ConCORDING KOT to UDDA

KOT codes 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 text 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.

At the end of the whole matching process, we still have two remaining issues to tackle. 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 languages 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, and 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 of these individuals are not in KOT. This affects 0.11% of all person-years, and 0.65% 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 college 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; or
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). 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.

C.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 concurs 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 (2021)’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 0.4% 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 0.5%) 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.⁵³
- 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 concordance 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.

⁵³Specifically see the document DISCO_koder_for_2010 available at <https://www.dst.dk/da/TilSalg/Forskningservice/Dokumentation/hoekvalitetsvariable/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 `concordanceNotes`.⁵⁴ 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 (which are hierarchal, so the subsequent rule only applies if there is no valid imputation before):

1. If a worker switchers firms from $t - 1$ to t , the code is missing in t but not in $t + 1$, and is assigned good in $t + 1$, we assign o_{t+1} .
2. If a worker switches from $t - 1$ to t , and the lag is good, impute the lag o_{t-1} .
3. Holes are filled. If o_t is missing and $o_{t-1} = o_{t+1}$ then o_{t-1} is assigned.
4. If the second digit is missing at t , and the leading digit between t and $t + 1$ matches, and o_{t+1} is good, then o_{t+1} assigned.
5. If the second digit is missing at t , and the leading digit between t and $t - 1$ matches, and o_{t-1} is good, then o_{t-1} assigned.

For codes that are not high quality, or not resolved 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. \(2022\)](#). 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 conduct robustness checks by excluding this from the estimation of switching elasticities.

⁵⁴This is available at Sharon Traiberman’s website, the servers on which the data can be accessed, and will be included in any future replication package.

C.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.

We conduct principal component analysis on the Knowledge, Abilities, and Skills modules. We match from SOC codes to ISCO codes and use census weights to collapse to the ISCO 4 digit level. We remove management tasks from all non-management occupations as they can badly bias estimates.⁵⁵ We collapse from the ISCO 4 to the ISCO 2 level using Danish population weights.

C.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 final analysis (which we call the high schooler sample), we only consider students who are within 5 years of high school graduation and who have applied at most twice. This also automatically drops students where this information is missing (for example, foreign students).⁵⁶ In our analysis we only consider students in the main 4 types of secondary program: Stx (split by speciality before 2008), Hhx, Htx, Hf. Ultimately, we keep 77% of KOT in our samples. For 10% of the sample, we do not have their high school track, or they went to a different track. For the remaining, we are missing other demographic information, or they do not meet the requirements to be included in the high schooler sample. A very small fraction of students are in specialized IB program or other programs (about 6200 students out of 1.4 million). We fold these into the Stx program. Similarly, there were in the past 1 and 2 year speciality Hhx programs, which affects about 18300 students. We combine these into a single Hhx program.

In the KOT database, we know which type of exam is required by a program (`kot_eksvo`). 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).

⁵⁵This occurs because occasionally production occupations will be matched with a “Supervisor” SOC code that is in turn given manager characteristics. Because of the structure of the match, where the census weight is all managers, including these occupation matches badly tilts all occupations towards managerial tasks.

⁵⁶For the purpose of constructing test score cutoffs, we actually pull all students for whom we have high school track data, as this can be used to discipline cutoffs.

C.5 Constructing Choice Sets

C.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 0.89, 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,

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

where GPA_{pt}^7 is the cutoff under the 7 point system in year t for program p . These regressions in either direction produce R^2 s over 0.7, 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 unconstrained programs, which have no threshold.

C.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 high school 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.

D Estimation Appendix

We begin by detailing the estimation of the labor supply block, including implementation of the EM algorithm. We then discuss how we adjust the estimated labor supply parameters for workers without a college education. Next, we describe the estimation procedure for the field choice block. Finally, we outline our bootstrap approach for computing standard errors.

D.1 Labor Market Block Details

D.1.1 EM Algorithm Implementation Details

Inputs into the EM Algorithm are the following:

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

$$q_{ik} = \begin{cases} 0.95 & \text{if clustered into } k \\ 0.05 & \text{if not clustered into } k \end{cases}$$

- 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 \mathbf{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.
 - Preallocate the covariance matrix, \mathbf{XtX} , to a $|\mathcal{O}| \times T + (M - 1) \times |\mathcal{O}| + (K - 1) \times Z$ square matrix of zeros.
 - Preallocate the covaraince matrix between Y and X , \mathbf{XtY} , to a $|\mathcal{O}| \times T + (M - 1) \times |\mathcal{O}| + (K - 1) \times Z$ length vector.
 - For $k = 1, \dots, K$...
 - Generate the weight matrix,

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

- Update the RHS covariance matrix:

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

- Update the LHS covariance matrix

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

- Estimate the Mincer regression parameters,

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

- 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

$$LL_k = \sum_{i=1}^N q_{ik} \left[\sum_{t=1}^T \left(-r_{itk}^2 / (2\sigma^2) - \log(\sigma^2) / 2 + \log \hat{\lambda}(o_{it} | o_{it-1}, m_i, k) \right) + \log(\hat{\lambda}^0(o_{i0} | t_0, m_i, k)) \right]$$

And define,

$$LL = \sum_k LL_k$$

Finally add cdf term:

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

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 $k = 1, \dots, K \dots$

i. Generate the weight matrix,

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

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	$NT \times K$

D.1.2 Stage 2 of Labor Supply

To derive (12),

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

and

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

Combining, we get,

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

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

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

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

D.1.3 Estimating Entry Costs

To derive (13),

$$\log \left[\frac{\lambda_t^e(o|k, m)}{\lambda_t^e(1|k, m)} \right] = -\frac{1}{\nu} \left(C_{o'km}^e + w_{ot}H_o(k, m) - w_{1t}H_1(k, m) + \beta(V_{t+1}(k, m, o) - V_{t+1}(k, m, 1)) \right)$$

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 (13). In practice we take the mean across all occupations.

D.2 Parameter Adjustments for Non-College Workers

After estimating the model for non-college workers, 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,

$$w_{otgk} = \mu_{otg} + \beta_{okg}. \quad (27)$$

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:

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

Another way to put this normalization ensures that things are measured in the same units.

Unfortunately, this leads to time-varying β 's. In order to avoid this, we only apply 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.

D.3 Details on Field Choice

D.3.1 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

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

Group observable or estimated terms—wages, costs, and transition probabilities—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,

$$V_t^e(m, s) = \left[\sum_k \pi(k|m, s) V_t^F(k, m, s) \right] + \beta V_{t+1}^e(m, s)$$

Hence,

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

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

D.3.2 Likelihood Conditional on l

The FOCs are given by,

$$\frac{\partial LL}{\partial \vartheta_{lm}} = \sum_t (1 - \beta^{T-6-t}) \sum_i q_{itl} \sum_r \left[\mathbf{1}(m_{r_i} = m) - \frac{N_{irm}^v \exp(U_{ml,t}(s_i))}{\sum_{m' \in \mathcal{C}_i} N_{irm'}^v \exp(U_{m'l,t}(s_i))} \right] \quad (28)$$

$$\frac{\partial LL}{\partial \theta_l} = \sum_t \sum_i q_{itl} \sum_r \left[XV_{t,T-6}^F(m_{r_i}, s_i) - \frac{\sum_{m \in \mathcal{C}_i} XV_{t,T-6}^F(m, s_i) \times N_{irm}^v \exp(U_{ml,t}(s_i))}{\sum_{m \in \mathcal{C}_i} N_{irm}^v \exp(U_{ml,t}(s_i))} \right] \quad (29)$$

$$\frac{\partial LL}{\partial v} = \sum_t \sum_l \sum_i q_{itl} \sum_r \left[\log N_{irm} - \frac{\sum_{m \in \mathcal{C}_i} \log N_{irm} \times N_{irm}^v \exp(U_{ml,t}(s_i))}{\sum_{m \in \mathcal{C}_i} N_{irm}^v \exp(U_{ml,t}(s_i))} \right]. \quad (30)$$

These first order conditions are intuitive.⁵⁷ 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.

D.3.3 EM Algorithm Implementation Details

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

- 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 $v = .65$ (arbitrary initialization)
- Remaining parameters are estimated from type-specific multinomial regression on top priority.

Step 2. Initialize state payoffs:

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

⁵⁷In 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:

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

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,

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

2. Update remaining parameters using FOCs:

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

3. Calculate the error as

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

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.

D.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.

D.5 Approximating V^e

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

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

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),

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

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 define,

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

We then have,

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

E Counterfactual Appendix

In this appendix, we provide more details for the counterfactual exercise. We start from presenting the labor demand in greater details, before showing the parameters adopted for the counterfactual exercise. We then describe the algorithm of counterfactual analysis. Finally, we present more results from the counterfactual experiments.

E.1 Labor Demand 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 :

$$H_{sot}^P = A_{st} \times \frac{\beta_{so}^P}{W_{ot}} \times A_s^P \times Q_{st} \times (P_{st}^D)^\alpha \times (P_{st}^P)^{\rho-\alpha} \times (W_{st}^P)^{-\rho}, \quad \forall o \in \mathcal{O}^P; \quad (31)$$

$$H_{sot}^N = A_{st} \times \frac{\beta_{so}^N}{W_{ot}} \times A_s^N \times Q_{st} \times (P_{st}^D)^\alpha \times (W_{st}^N)^{-\alpha}, \quad \forall o \in \mathcal{O}^N; \quad (32)$$

and demand for imported goods is

$$Q_{st}^M = A_{st} \times A_s^M \times Q_{st} \times (P_{st}^D)^\alpha \times (P_{st}^P)^{\rho-\alpha} \times (P_{st}^M)^{-\rho}, \quad (33)$$

where

$$Q_{st} = A_{st}^F (P_{st}^D)^{-\eta} + \mu_s E_t \times \left(\frac{P_{st}}{P_t} \right)^{1-\sigma} \times \frac{(P_{st}^D)^{-\eta}}{(P_{st}^D)^{1-\eta} + (P_{st}^F)^{1-\eta}}; \quad (34)$$

$$W_{st}^f = \prod_{o \in \mathcal{O}^f} w_{ot}^{\beta_{so}}, \quad f = \{P, N\}; \quad (35)$$

$$P_{st}^P = \left[A_s^P (W_{st}^P)^{1-\rho} + A_s^M (P_{st}^M)^{1-\rho} \right]^{\frac{1}{1-\rho}}; \quad (36)$$

$$P_{st}^D = \left[A_s^N (W_{st}^N)^{1-\alpha} + (P_{st}^P)^{1-\alpha} \right]^{\frac{1}{1-\alpha}} \quad (37)$$

$$P_{st} = A_{st}^{-1} \left[(P_{st}^D)^{1-\eta} + (P_{st}^F)^{1-\eta} \right]^{\frac{1}{1-\eta}} \quad (38)$$

$$P_t = \left[\sum_s \mu_s (P_{st})^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad (39)$$

$$E_t = \sum_o w_{ot} \left(\sum_s L_{sot} \right) \quad (40)$$

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

$$H_{ot} = \sum_s H_{sot}^k \tag{41}$$

E.2 Demand Side Parameters Calibration

Table 17 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 17
Labor Demand Parameters and Variables

Parameters	Meaning	Value/Data Source
$\beta_{so}^P, \beta_{so}^N$	Labor inputs coefficients	IDA
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. We obtain raw wage data from the Integrated Database for Labor 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). We set elasticity across sectors σ at 0.2, following Atalay (2017) and trade elasticity η at 4, following Simonovska and Waugh (2014).

E.3 Additional Results for Capacity Constraint Counterfactual

Maintaining Medicine Cutoffs

In this appendix, we analyze the “Partial Flex” counterfactual of Section 7.2, but we do not increase access to medical programs. Table 18 describes the key variables of interest we explored in the main

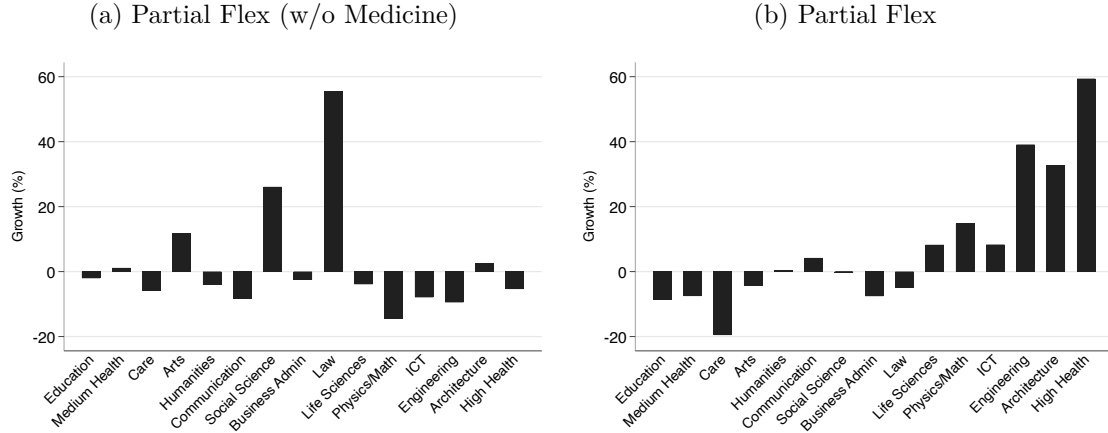
Table 18
GE Consequences of Lifting Capacity Constraints

	Partial Flex (w/o Medicine)	Partial Flex
% Δ GDP	0.02	0.31
% Δ s.d. log Income	0.49	1.07
Field Reallocation	6.67	8.66
Occupational Reallocation	5.61	6.90

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 calculated by simulating 200 thousand individuals in each counterfactual and steady state and determining the proportion of switchers.

text. The partial flex scenario is reprinted for ease of reading. Disallowing reallocation into medicine only partially dampens the amount of reallocation in the economy, but dramatically dampens the aggregate impacts on output on inequality. Nevertheless, directionally, the patterns are largely the same.

Figure 16
Field Reallocation Without Capacity Constraints



Notes: Plots (100x) changes in the log shares of each field, $\Delta \log s_f$, completed by new cohorts. Panel (a) removing capacity constraints except for medicine; Panel (b) removing capacity constraints.

To better understand the differences, 16 plots reallocation with and without increased access to medicine. In this case, there is actually a net *decline* in those pursuing medicine—reflecting those eligible reallocating away. Instead, there is much more entry into Law, and out of various STEM fields.

E.4 Construction of Labor Market Shocks

We describe how we construct the trade war and AI technological shocks for the second set of counterfactual exercises, in which we examine the economic outcomes in response to the labor market disruptions.

E.4.1 Trade War Shock

The trade war shock in our analysis aims to capture three key components of recent global trade tensions: increased costs of imported consumption goods, higher prices of offshored inputs, and reduced foreign demand for domestic exports. Specifically, we model a 30% increase in the price of imported goods (P_{st}^F) and offshored inputs (P_{st}^M), along with a 30% reduction in foreign demand for Danish exports (A_{st}^F). These values are not intended as precise forecasts but are chosen to capture

the possible scale of trade distortions and the economic costs associated with heightened uncertainty. The symmetric shock structure reflects disruptions to both production and consumption, as well as diminished access to export markets—offering a stylized representation of a sustained trade conflict.

E.4.2 AI Shock

Suppose there is an external measure that predicts declines in *relative* labor demand.⁵⁸ 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*,

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

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

$$\tilde{z}_{so} = z_o \times \frac{1}{\frac{w_s^N L_s^N}{w_s L_s} (\sum_{o \in N} \beta_{so} z_o) + \frac{w_s^P L_s^P}{w_s L_s} (\sum_{o \in P} \beta_{so} z_o)}$$

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,

$$\frac{w_o L_{so}}{w_s L_s} = \beta_{so}^N A_{sN} \left(\frac{w_s^N}{P_s^D} \right)^{1-\alpha}.$$

So we need to solve the following:

$$\tilde{\beta}_{so}^N \tilde{A}_{sN} \left(\frac{w_s^N}{P_s^D} \right)^{1-\alpha} = \tilde{z}_{so} \beta_{so}^N A_{sN} \left(\frac{w_s^N}{P_s^D} \right)^{1-\alpha},$$

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

$$\tilde{A}_{sN} = A_{sN} \sum_o \tilde{z}_{so} \beta_o^N.$$

Hence,

$$\hat{A}_{sN} = \sum_o \tilde{z}_{so} \beta_{so}^N.$$

We also have,

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

For production workers, there is a similar common price term that ultimately cancels out, so we

⁵⁸In the actual implementation, we adopt the predicted changes from [Eloundou et al. \(2023\)](#).

have the same expressions:

$$\hat{A}_{sP} = \sum_o \tilde{z}_{so} \beta_{so}^P$$

$$\tilde{\beta}_{so}^P = \frac{\tilde{z}_{so} \beta_{so}^P}{\sum_{o'} \tilde{z}_{so'} \beta_{so'}^P}$$

E.5 Algorithm

We describe the algorithm adopted to solve the counterfactual steady state. The algorithm for transition dynamics is available upon request.

Consider shocks $\{P_{st}^{F,0}\} \rightarrow \{P_{st}^{F,1}\}$, $\{P_{st}^{M,0}\} \rightarrow \{P_{st}^{M,1}\}$, $A_{s,0}^F \rightarrow A_{s,1}^F$, as well as $\beta_{so0} \rightarrow \beta_{so1}$. 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: $\{w_o^0\}$. 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

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

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

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

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

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

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

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

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

Note that the nominal wage $w_{o'}^1 H_{o'}(k, m)$ is deflated by the price index P to reflect real wage changes.

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

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

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 $\{L_{km}^1\}$ 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 $r = 1, \dots, R$ calculate the ranked list of choices for individual i :

- Calculate the conditional CDF of choices for field:

$$\tilde{Q}_{ir}^F(m) = \frac{\sum_{m'=1 \setminus m-r}^m \exp(\vartheta_{mt} + \vartheta_{lm} + \theta_l V^1(m|s))}{\sum_{m'=1 \setminus m-r}^M \exp(\vartheta_{mt} + \vartheta_{lm} + \theta_l V^1(m|s))}$$

- Generate R draws from this distribution:

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

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:

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

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_{km}^1 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:

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

This implies,

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

(b) Construct H_{omk}^1 :

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

(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

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

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

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

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

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

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:

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

E.6 Additional Labor Market Shocks Results

E.6.1 Alternative Base Steady State

In this appendix, we consider an alternative approach to modeling the initial steady state. In the baseline analysis, we project the 2018 economy forward under the existing college admissions system, allowing us to assess whether greater flexibility after the shock enhances resilience to economic shocks. Here, we instead begin from the *Full Flex* scenario introduced in Section 7.2, in which all constraints are removed. This exercise evaluates whether a more flexible educational system, which are already in place before the shock, better mitigates its long-run effects.

Table 19
Results across Policy Scenarios: Alternative Base Steady State

	Panel A: Trade War Shock			Panel B: AI Shock		
	Current Base to Current CF (1)	Current Base to Flex CF (2)	Flex Base to Flex CF (3)	Current Base to Current CF (1)	Current Base to Flex CF (2)	Flex Base to Flex CF (3)
% Δ Income ($\Delta < 0$)	0.48 (56.7%)	9.20 (55.0%)	0.04 (55.0%)	15.4 (40.0%)	24.6 (41.7%)	14.9 (40.2%)
Field Reallocation	2.3%	26.1%	2.1%	3.5%	26.7%	3.4%
Occ Reallocation	13.1%	25.2%	12.5%	18.3%	29.4%	17.7%
% Δ s.d. log Income	18.5	20.4	18.6	15.7	17.7	15.9%

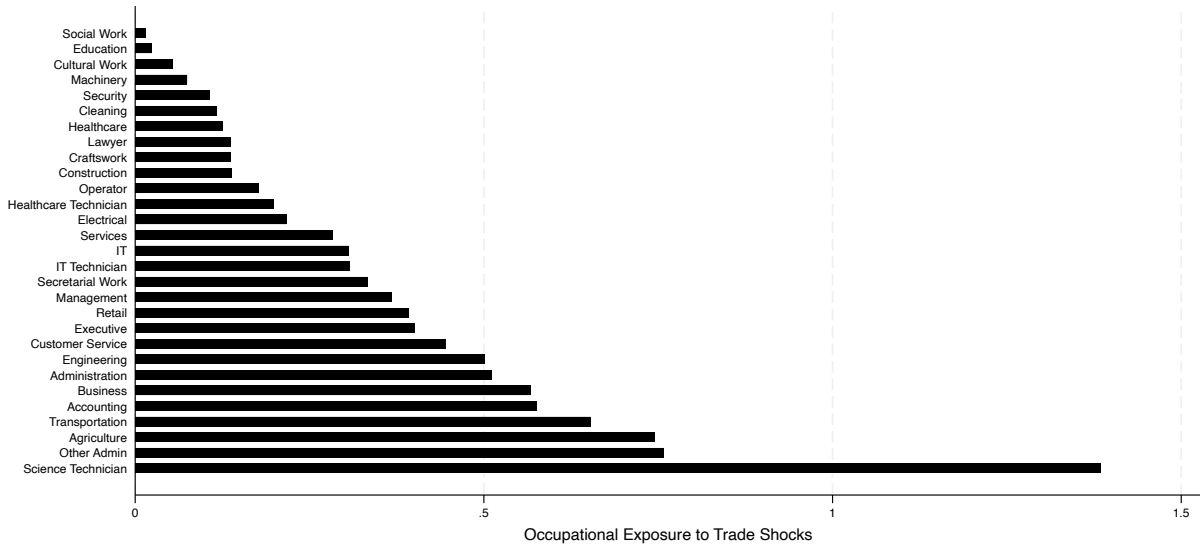
Notes: Table reports the changes for row variables between steady state after a policy simulation and the initial steady state. Current Base refers to an initial steady state in which the current constraints are maintained. Flex Base refers to the Fully Flex scenario discussed in Section 7.2, in which all capacity constraints are removed. Columns (1) and (2) for both panels are the same as the corresponding scenarios in Table 8. % Δ Income is the percentage change in NPV between the counterfactual and baseline scenarios, with NPV calculated using (26). % Δ s.d. log Income is the percentage change in the standard deviations of log income for the 200,000 simulated individuals between the counterfactual and baseline scenarios. Occupation and Field reallocation reports the shares of the 200,000 simulated individuals who switch occupations and fields, respectively, between the counterfactual and baseline scenarios.

Table 19 reports the findings. Columns (1) and (2) are reprinted from Table 8 (resp. Columns (2) and (3)) for easy comparison. Not surprisingly, starting from a more flexible educational steady state (Flex Base) leads to smaller gains from further flexibility, as much of the adjustment has already occurred. For both shocks, income gains under the Flex Base are small (0.04% and 14.9%, respectively) compared to scenarios starting from the Current Base. Field and occupational switching rates are also lower when flexibility is already embedded in the initial steady state. Interestingly, inequality remains more similar across specifications. These results suggest that economies with more flexible education systems in place prior to a shock are better positioned to absorb its long-run effects. Intuitively, key reallocations, both in educational choices and occupational sorting, have already occurred in anticipation of shifting demand, reducing the need for further costly adjustment.


E.6.2 Additional Figures from Counterfactual Analyses

Figure 17
Occupational Exposure to Labor Market Shocks

(a) Trade War Shock



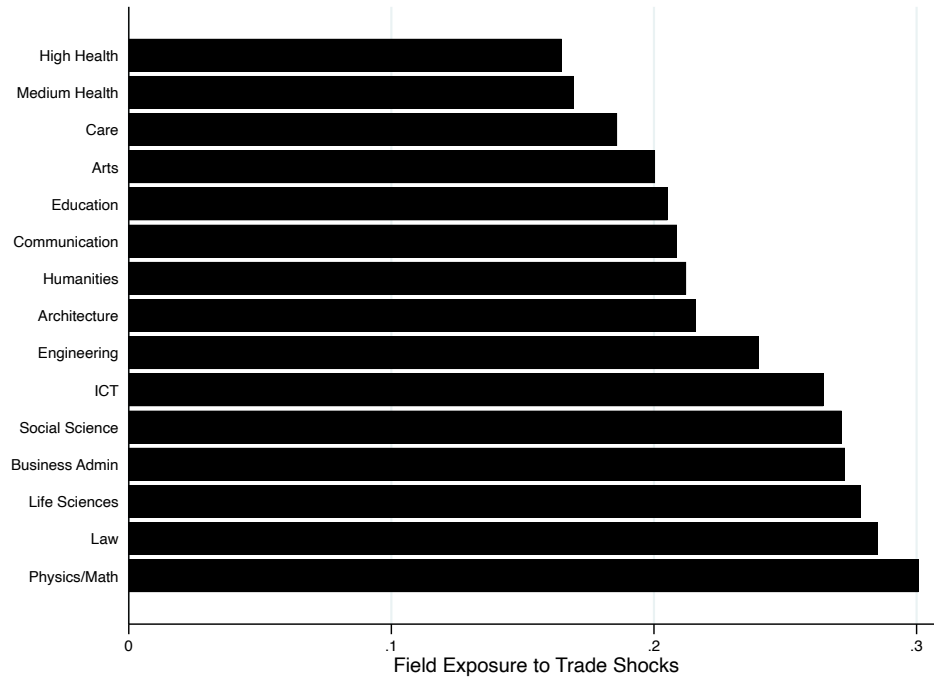
(b) AI Shock




Notes: We document the extent of exposure to the labor market shocks across occupations. In Panel (a), the simulated shock is a 30% price increase in imported goods and offshored inputs, and a 30% reduction in export demand. We calculate exposure using: $Exposure_o = \sum_i \left(\frac{X_i}{WB_i} - \frac{M_i}{WB_i} \right) \frac{WB_{oi}}{WB_o}$, where X_i is the value of export, M_i is the combined value of the foreign inputs and imported goods, WB_i , WB_o and WB_{oi} are the wage bill at the sector, occupation, and sector-occupation level. All are measured in 2008 values. The exposure metric in Panel (b) is obtained directly from [Eloundou et al. \(2023\)](#). See Section 7.3 for corresponding discussions.

Figure 18
Field Exposure to Labor Market Shocks

(a) Trade War Shock



(b) AI Shock



Notes: We document the extent of exposure to the labor market shocks across fields. In Panel (a), the simulated shock is a 30% price increase in imported goods and offshored inputs, and a 30% reduction in export demand. We calculate exposure using: $\sum_o Exposure_o \frac{L_{of}}{L_f}$, where $Exposure_o$ is the occupational exposure calculate in Figure 17, L_{of} is the 2018 allocation of workers in occupation o and field f , and L_f is the 2018 number of workers in field f . See Section 7.3 for corresponding discussions.