# Artificial Intelligence and College Majors

Anders Humlum* & Bjørn Bjørnsson Meyer†

July 29, 2020

**Abstract**

We study how Artificial Intelligence (AI) technology relates to the returns to college majors. Using micro data from Denmark, we first rank college majors according to whether their graduates work in AI firms. We show that AI cuts through the category of STEM degrees: while computer science and mathematics majors specialize in AI producer firms, we find that the laboratory sciences concentrate in firms that only use AI. We document that AI producer relevance correlates with higher wages, and that these earnings premiums are rising. In contrast, AI user majors do not earn wage premiums, nor have experienced differential earnings growth. Using a regression discontinuity design, we estimate the causal effects of admitting students to more AI relevant college majors. We find that admission cutoffs are effective in increasing the AI relevance of college graduates. The causal earnings gain from higher AI producer relevance is at least as large as the correlation.

*Princeton University, Department of Economics, humlum@princeton.edu.
†University of Copenhagen, Department of Economics, bjorn.bjornsson.meyer@econ.ku.dk.

# 1 Introduction

New technologies are a source of improved living standards, but may also pose risks of widening skill gaps and greater inequality. Today, Artificial Intelligence (AI) is one of the most salient technology arrivals, and its impact on labor markets is highly anticipated.

Artificial Intelligence, or Machine Learning, refers to algorithms that learn to complete tasks by identifying statistical patterns in data. AI is poised to be the next General Purpose Technology that will sweep through a wide range of industries (Brynjolfsson et al., 2018; Goldfarb et al., 2019). A distinctive feature of AI is that it may create winners and losers within the group of highly educated workers. While AI algorithms are being developed and implemented by specialized computer and data scientists, these same algorithms may be able to automate tasks done by other high-skilled workers, such as chemical engineers or optometrists (Webb, 2020).

In this paper, we document what types of college skills firms that use or produce AI employ. We view a college major as a set of skills that may be more or less relevant to work with AI. Using micro data from Denmark, we rank college majors according to their share of graduates that work in AI firms. Using these AI relevance rankings, we establish new facts about the college majors that work in AI firms, and we leverage an admission cutoff regression discontinuity (RD) design to estimate causal effects of pushing students toward more AI relevant degrees.

Our first key finding is that AI cuts through the group of Science, Technology, Engineering, and Mathematics (STEM) degrees. While computer science and mathematics majors specialize in AI producer firms, we find that biology, chemistry, and other laboratory sciences concentrate in firms that only use AI. We find that AI producer relevance is correlated with higher earnings in the labor market and that these earnings premiums have increased in recent years. In contrast, we find that AI user majors do not earn higher wages, nor have experienced differential earnings growth recently.

Our second contribution is to provide causal evidence on an often discussed policy response to AI, that is to adjust the supply of different college majors (President's Council of Advisors on Science and Technology, 2012; Trajtenberg, 2018). In Denmark, admission to college majors is based on high school GPA cutoffs. These admission cutoffs naturally give rise to a regression discontinuity (RD) design, comparing outcomes of applicants around cutoffs for AI relevant majors. The RD analysis delivers two key insights. First, we find that it indeed is possible to push marginal students to obtain more AI relevant skills. Many students apply to majors of widely differing AI relevance, and the cutoffs create discontinuous jumps in admission rates to these programs. We find that students comply with these discontinuities, such that an applicant just above a cutoff has a discontinuously higher chance of graduating in that major. Second, using these discontinuities in college major attainment rates, we find that earnings premiums from random assignment to higher AI producer relevance are as large as those suggested by the correlation. In contrast, we find that the returns to higher AI user relevance is a precisely estimated zero. These estimated causal effects are informative about policies that aim to increase the AI relevance of college students by adjusting the admission capacity of different college majors.

We make a methodological contribution in this paper by proposing new measures of the exposure of college majors to AI. A key advantage of our measures, which we microfound in task-based theories of AI, is that they rely on revealed and objective information only. We correlate our revealed measures with characteristics of the college majors to shed light on the types of skills that are relevant for AI. First, using linked patent-occupation data from Webb (2020), we find that application-specific breakthroughs in AI are a key driver of the AI user relevance of different college majors. Second, we find that a core in programming or mathematics is a necessary ingredient for work with AI production. We do not find evidence that interpersonal skills alone qualify for work with AI, as social science and the humanities rank as the least relevant majors for AI firms. Our measurement approach does, however, require very detailed micro data on both AI technology by firms and the college majors held by workers. We leverage here the unusually rich data available in the Danish setting, where we can link register data on college diplomas of workers with a novel survey of AI production and use by firms.

This paper combines insights from two strands of literature. The closest related work is a recent set of papers that provide evidence on how AI relates to the demand for worker tasks (Alekseeva et al., 2019; Goldfarb et al., 2019; Webb, 2020). We make two contributions to this line of work. First, we use new micro data from Denmark to rank college majors according to their relevance for AI firms. Second, we adopt an admission RD design to estimate the causal returns to AI relevant skills. Admissions cutoffs have been used successfully in the literature on the returns to college field choice (Hastings et al., 2013; Kirkeboen et al., 2016). We complement these papers by linking college majors to a specific technology using a continuous measure that goes beyond traditional divisions of fields of study.

The remainder of the paper is structured as follows. Section 2 explains how we rank college majors according to AI relevance. Section 3 uses the AI relevance rankings to document facts about the college majors employed in AI firms. Section 4 conducts the admission regression discontinuity analysis. Section 5 concludes.

# 2 Measuring the AI Relevance of College Majors

In this section, we rank college majors according to their share of graduates that work in AI firms, a measure we call "AI relevance". To measure AI relevance, we link a novel survey of AI production and use by firms (described in Section 2.1) to register data on college diplomas of workers (described in Section 2.2). In Section 2.3, we relate the AI relevance measures to task-based theories of AI.

## 2.1 Firm Survey on Artificial Intelligence

Our micro data on AI comes from the representative VITA firm survey conducted by Statistics Denmark in 2017. The survey questions were prepared in collaboration with the Danish Business Authority as a supplement to Eurostat's technology survey. The survey sampled 3,890 firms from the population of private non-agricultural, non-financial firms with more than 10 employees. The response rate was 97 percent, and the surveyed firms covered 27 percent of total employment in Denmark.

Figure 1 shows the questionnaire on Artificial Intelligence. AI is defined as "computer software that "thinks", analyzes, solves problems, and recognizes patterns in data", and examples include computer-generated annual reports, chatbots, or automated marketing. Firms were asked if they used AI in-house and whether they sold products or services containing AI technology.

Figure 1: Firm Survey on Artificial Intelligence

**Use of Machine Learning and Artificial Intelligence**

Machine learning and artificial intelligence is the use of computer software that "thinks", analyzes, solves problems, and recognizes patterns in data, e.g. image, sound or text data. This includes, among other things, computer generated annual reports, chatbots or automated marketing.

| | Yes | No |
|---|---|---|
| 18. **Does the firm utilize machine learning or artificial intelligence?** Incl. services from external providers containing these. | ○ | ○ |
| 19. **Does the firm sell products or services containing machine learning or artificial intelligence?** | ○ | ○ |

Among the surveyed firms, 2.2 percent answered 'yes' to selling AI-augmented products or services. AI producers employ 5.9 percent of workers, and 13.7 percent of college graduates. While most AI producers also use AI (70 percent), only a third of AI users sell AI-augmented products. For this reason, we will distinguish between firms that produce AI (henceforth, "AI producers") and those that use, but do not produce, AI (henceforth, "AI users"). Among the surveyed firms, 3.0 percent use, but do not produce AI. The AI users employ 10.3 percent of workers, and 18.9 percent of college graduates. Warzynski (2020) provides additional descriptives on the sample of AI using firms, studying how AI adoption correlates with firm industry, productivity, and workforce composition. In this paper, we use the AI survey to rank college majors according to their relevance for AI firms, which we turn to next.

## 2.2 Linking the AI Survey to College Diplomas

To calculate AI relevance, we first merge the AI survey to the employer-employee link file of the Integrated Database for Labor Market Research (IDAN). IDAN links the population of workers to the universe of firms in November of each year, allowing us to identify employees in AI firms. We then merge in worker diplomas from the Education Register of Statistics Denmark, allowing us to calculate the share of college majors that work in AI firms.

A bachelor" degree in Denmark is a three year major-specific program, but most students enroll in a two-year master's degree directly following their bachelor's studies. College education is rather specialized in Denmark with 226 distinct major (`AUDD`) categories. To increase the measurement accuracy of the AI relevance scores, we pool graduating cohorts of each major into 10-year intervals. Furthermore, to reduce noise from very small cells, we require that at least 100 graduates of a college major be present in the VITA survey before assigning the major an AI relevance score. For majors that do not meet this minimum criterion, we go up one level of aggregation up (75 different categories) and assign them the more aggregated AI relevance score of this category. In calculating the AI relevances, we use the survey weights provided by Statistics Denmark to ensure that the distribution of survey firms reflects the population of firms in Denmark.

### 2.3 A Micro Foundation for the AI Relevance Measures

How do our AI relevance measures relate to the task-based concepts of automation emphasized in the theoretical literature on AI? Using a task-based model of production, Acemoglu and Restrepo (2018) show that AI have three distinct effects on labor demand: (i) a negative *displacement effect* in the production tasks that AI automates, (ii) a positive *reinstatement effect* in the worker tasks that are needed to facilitate AI technology, and (iii) a positive *productivity effect* stemming from cost savings in the tasks that AI can do more efficiently than humans.

In Appendix B, we build on these insights to derive the following theoretical implications: (i) AI user relevance is the theory-consistent measure of a college major's exposure to AI's displacement effects, (ii) AI producer relevance is the theory-consistent measure of a college major's exposure to the reinstatement effects of AIs, and (iii) AI's productivity effects raise the demand for workers uniformly, and thus cancel out when we compare the relative wages for different skill groups.

## 3 Stylized Facts on Artificial Intelligence and College Majors

In this section, we use the AI relevance rankings from Section 2 to document a series of facts about the college majors that are employed in AI firms. Our first key finding is that AI cuts through the group of Science, Technology, Engineering, and Mathematics (STEM) majors. While computer science and mathematics majors specialize in AI producer firms, the laboratory sciences concentrate in firms that only use AI. Our second set of facts concern the wage premiums that accrue to AI relevant majors. We find that AI producer majors earn higher wages and that these earnings premiums have increased in recent years. In contrast, we find that AI user majors do not earn higher wages, nor have experienced differential earnings growth recently.

## Fact 1: Artificial Intelligence Cuts Through the STEM Degrees

Figure 2 shows a scatter plot of college majors according to their AI relevance. The figure highlights three key characteristics of AI relevant majors. First, computer science, mathematics, and physics are the most relevant majors for work with AI production. Second, biology, chemistry, and other laboratory sciences are highly concentrated in firms that only use AI. Reflecting this divide between the mathematical and laboratory sciences, engineering is also highly polarized with chemical engineers employed in AI using firms and robotics engineers employed with AI producers. Finally, the Arts and Humanities are the least relevant majors for both AI using and AI producing firms.

Figure 2: AI Relevance of College Majors

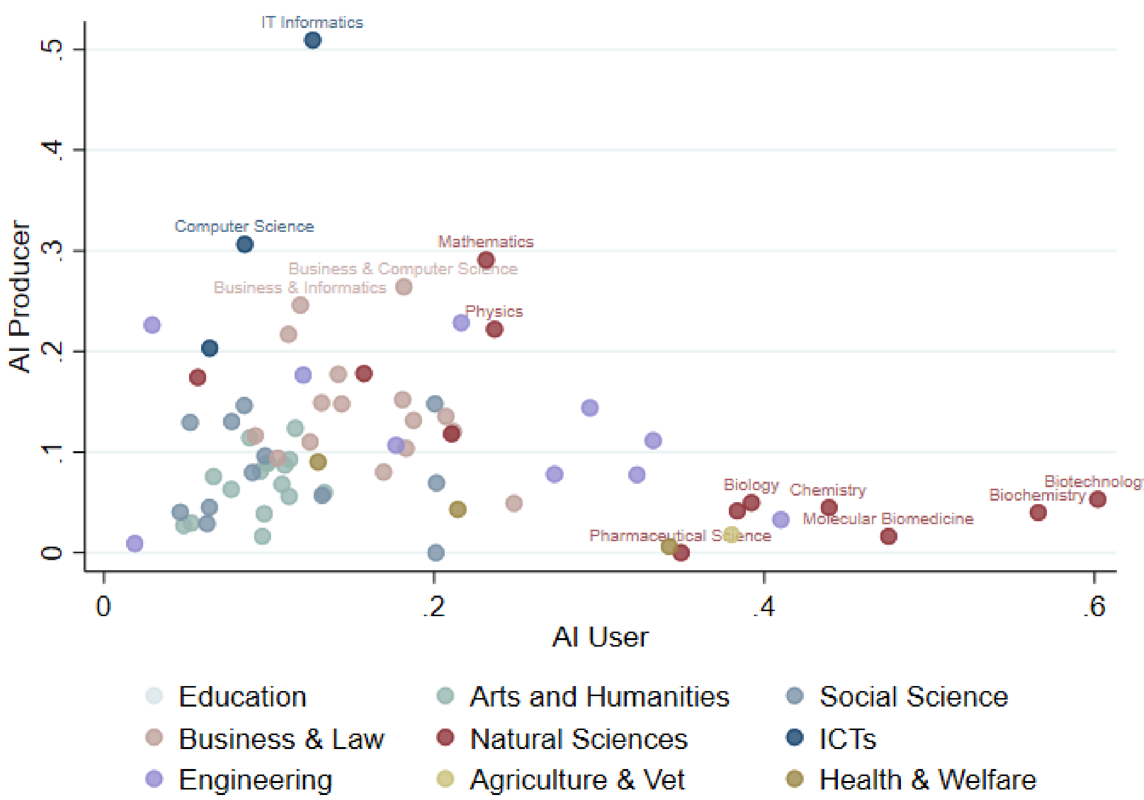


*Notes:* This figure plots the AI shares of college majors, defined as the fraction of college major graduates that work in AI user/producer firms. Majors are colored according to their one-digit International Standard Classification of Education (ISCED) code.

What drives the high relevance of biology and chemistry majors to AI user firms? One hypothesis is that AI has experienced technical breakthroughs in tasks that laboratory scientists traditionally carry out. Following on this idea, Webb (2020) links AI patents to job task descriptions and finds that important applications of AI relate to the automation of tasks by clinical laboratory technicians and chemical engineers. To shed light on this hypothesis, Figure 3 correlates the AI user relevance of a college major with its share of graduates that work in jobs with extreme exposure to AI applications as measured by Webb (2020). The strong positive correlation lends clear empirical support to the hypothesis that technical breakthroughs in AI applications are important determinants of the AI user relevance of different college majors. In comparison, Appendix Figure A.1 shows that AI producer relevance is not correlated with exposure to AI applications.

Figure 3: Share Employed in Jobs with Extreme Exposure to AI Applications (Webb, 2020)

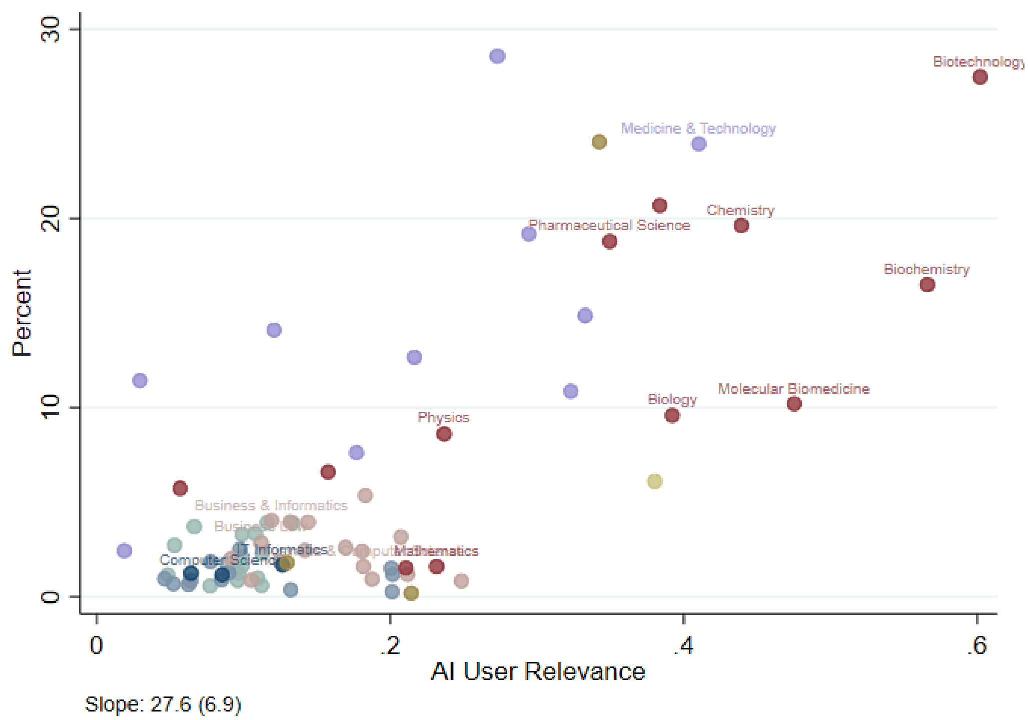


*Notes:* This figure shows the share of workers employed in occupations with extreme exposure to AI applications by AI user relevance of their college major. Extreme exposure to AI applications is defined as occupations (6-digit ISCO08) with AI exposure scores of 95/100 or above. See Webb (2020) for details about the AI exposure measure. Majors are colored according to their one-digit ISCED code; see Figure 2 for a color description. For the estimated slope coefficient, college majors are weighted by their number of graduates, and standard errors are clustered by college majors. Appendix Figure A.1 shows the corresponding figure by AI producer relevance.

What characterizes the college majors that are employed in AI producer firms? Inspection of Figure 2 suggests that a strong core in computer programming or mathematics may be a necessary ingredient for AI producer relevance. However, when combined with IT capabilities, business skills seems to be a complement in AI production. For example, majors that combine business with computer science or informatics rank in the top of AI producer relevance. Interestingly, among the top 10 for AI producer relevance, business law is the only major that does not include a strong quantitative component. This could suggest that AI production requires the support services of legal experts to navigate the unchartered legal territory that surrounds AI applications.

Finally, we do not find evidence that interpersonal skills alone qualify for work with AI, as otherwise often proclaimed in public debate (Trajtenberg, 2018; World Economic Forum, 2018). In fact, we find that social science and humanities are the least relevant majors for work in AI firms.

### Fact 2: AI Producer Majors Earn Higher Wages, AI Users Do Not

Figure 4.(a) shows that AI producer relevance is strongly correlated with higher starting salaries when entering the labor market. A standard deviation increase in AI producer relevance (6.4 pp) is associated with about $4,100 higher annual starting salary, and the relationship is linear throughout the distribution of AI producer relevances. In contrast, Figure 4.(b) shows that starting salary does not share a similar linear relationship to AI user relevance. While extremely low AI user relevance (as for the Arts and Humanities) is associated with lower earnings, extremely high AI user relevance (as for the laboratory sciences) are not associated with correspondingly higher earnings. In fact, once we drop the Arts and Humanities from Figure 4.(b), there remains no statistically significant correlation between AI user relevance and starting salary of college graduates.

Figure 4: Starting Salaries in 2016

(a) By AI Producer Relevance of College Major (b) By AI User Relevance of College Major

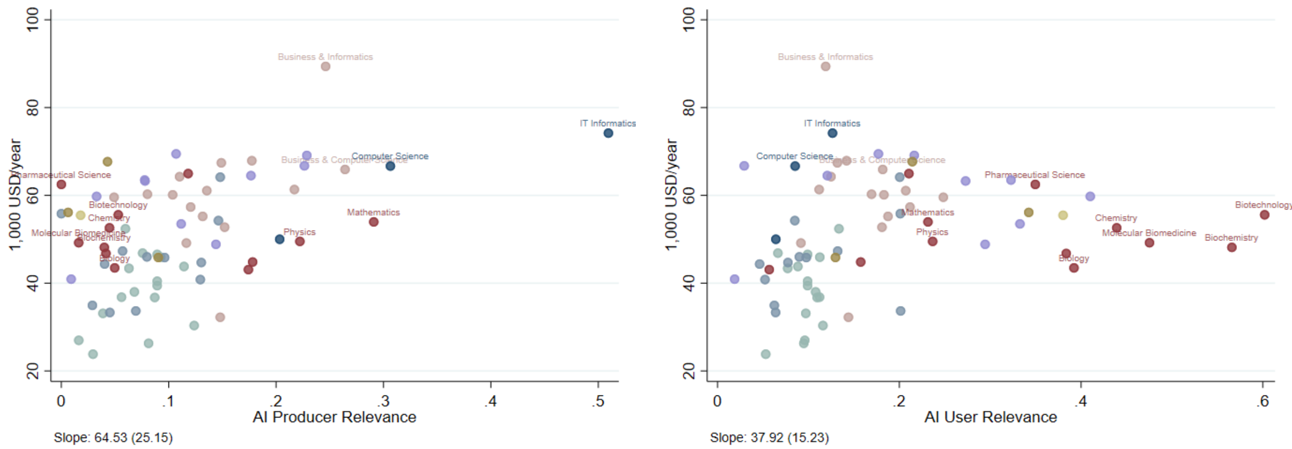


*Notes*: This figure shows the average starting salaries of college graduates by AI relevance of their major. Starting salary is defined as salary eight years after the start of a bachelor's degree. Salary includes wage compensation for employees and net profits for self-employed. Zero incomes are included. Salary is measured in 1,000 US dollars, using an exchange rate of 5.85 Danish Kroner per US dollar. For the estimated slope coefficients, college majors are weighted by their number of graduates, and standard errors are clustered by college majors. Majors are colored according to their one-digit ISCED code; see Figure 2 for a color description.

The earnings premium for AI producer skills is consistent with Alekseeva et al. (2019), who use job posting data from the United States to document large wage premiums for AI-specific programming skills such as proficiency in the machine learning packages TensorFlow and PyBrain. In Appendix Table A.1, we show that the earnings premiums of AI relevance are shared across workers with a given college major, and not confined to workers in AI firms only. This finding highlights the importance of studying AI earnings premiums throughout the economy, and not just in AI firms.

Finally, in Appendix Figure A.2, we document a sharp gender divide in the AI relevant majors, wherein men select into AI producer majors, and women select into AI user majors. For example, in the computer science degrees that are relevant for AI producers, less than ten percent of students are female. The gender divide persists within STEM degrees and suggests that an expansion of college majors according to AI relevance alone could exacerbate existing gender gaps in higher education and labor market outcomes.

### Fact 3. Wage Premiums Are Rising for AI Producer Majors, Not for AI Users

To probe the timing of AI technology, Figure 5 shows that the number of Google searches on "machine learning" increased by a six-fold from 2012 to 2017.

Figure 5: Google Searches on "Machine Learning"

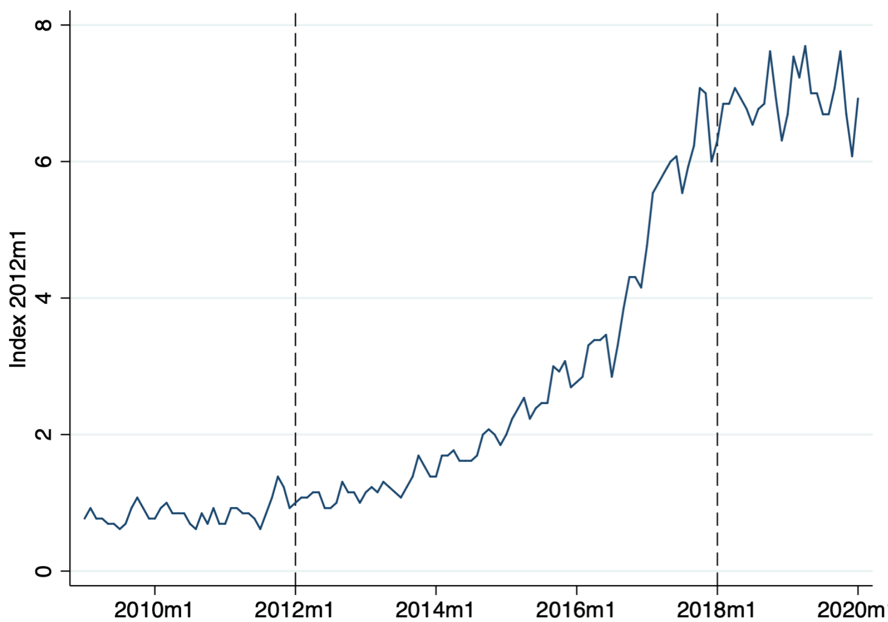


*Notes*: This figure shows the number of monthly Google searches on "machine learning" in the United States from January 2010 to January 2020. The units are normalized such that the maximum monthly search activity takes a value of 100.

Motivated by this recent rise of interest in AI, Figure 6 plots the growth in starting salaries of college major cohorts from 2012 to 2016. Salaries are deflated using the CPI with 2015 as the base year. Panel (a) shows that AI producer majors have experienced excess earnings growth since 2012. A standard deviation higher AI producer relevance (6.4 pp) is associated with 0.4 percent points higher annual growth in starting salaries. In contrast, Panel (b) shows that AI user majors have not experienced differential earnings growth in recent years.

Thus, if AI producer relevance captures exposure to AI's task reinstatement effects as implied by our theory in Appendix B, then Figure 6.(a) suggests that these positive impacts are already beginning to show in the starting salaries of college majors. In contrast, if AI user relevance captures exposure to AI's task displacement effects, then Figure 6.(b) shows that large negative displacement impacts of AI are yet to manifest themselves in starting salaries.

In Appendix Figure A.3, we replicate Figure 6 for an earlier period of years (2007-2012), showing that the differential earnings growth of AI producer majors is a recent phenomenon closely linked in time to the recent take-off in machine learning. This finding is consistent with Alekseeva et al. (2019), who document that the share of AI-skills in total job postings in the United States doubled from 2012 to 2016. Having said that, we want to emphasize that earning growth from 2012 to 2016 accounts for less than 20 percent of the observed level differences in earnings across AI producer relevance documented in Figure 4. Put differently, AI producer relevant skills have always been rewarded in the labor market, but the evidence presented here suggests that AI may have exacerbated these inequalities.

Figure 6: Annualized Growth in Starting Salaries from 2012 to 2016

(a) By AI Producer Relevance of College Major (b) By AI User Relevance of College Major

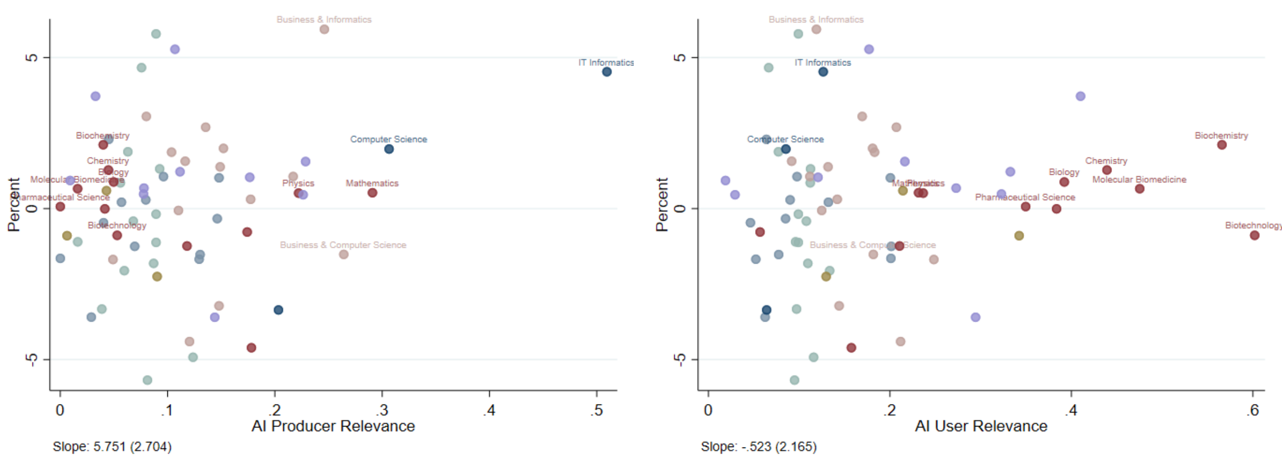


*Notes*: This figure shows the average annualized growth rates in starting salaries of college graduates from 2012 to 2016 by AI relevance of their major. Starting salary is defined as salary eight years after the start of a bachelor's degree. Salary includes wage compensation for employees as well as net profits for self-employed. Salaries are deflated using the CPI with 2015 as the base year. Zero incomes are included. For the estimated slope coefficients, college majors are weighted by their number of graduates, and standard errors are clustered by college majors. Majors are colored according to their one-digit ISCED code; see Figure 2 for a color description.

We view the evidence presented in Figure 6 as consistent with Brynjolfsson et al. (2018), who observe that firms currently are undertaking significant investments into AI, but hypothesize that the productivity gains from these investments may only arrive with a delay. The investment activities are already generating positive reinstatement effects for AI producers (Figure 6.(a)), but the productivity gains of AI – an important source of potential displacement of AI user majors – are yet to manifest themselves (Figure 6.(b)).

# 4 Causal Evidence on Increasing AI Relevant College Skills

In Section 3 above, we documented that college majors that are relevant for AI producers earn higher wages and that these earnings premiums are on the rise. On top of these expanding wage premiums, technologists and scholars expect that the diffusion of AI will accelerate in the years to come (Brynjolfsson et al., 2018; Goldfarb et al., 2019).

How can we prepare a future generation of college students for work with AI? An often discussed policy is to adjust the admission capacities of different college majors (President's Council of Advisors on Science and Technology, 2012; Trajtenberg, 2018). However, any policy that aims to increase the AI relevance of college students by adjusting admission criteria will face at least two challenges.

First, will students, on the margin of attending an AI relevant major, complete the degree if offered a seat? Second, if the students complete the degree, will they reap the positive returns in the labor market that are suggested by the correlations in Section 3? In this section, we study the Centralized Admission System in Denmark, where admission cutoffs allow us to provide causal evidence to each of these questions.

In Section 4.1, we first show that a substantial number of students stand on the margins of majors with widely differing AI relevances. Second, in Sections 4.3 and 4.4, we find that, if offered a spot at a more AI relevant major, these marginal students will not only complete the degree but also end up earning labor market returns that are at least as large as those suggested by the correlations.

## 4.1 Centralized College Admission in Denmark

The Danish university system consists of eight bachelor's granting institutions, all of which are public and use the Centralized Admission System (CAS).[1] A bachelor's degree is a three-year major-specific program, and switching between programs will, in most cases, mean reapplying through CAS. The admission capacities of different programs are centrally regulated by the Danish Ministry of Higher Education and the majority of program slots are allocated according to the high school GPA of applicants. Students can rank up to eight different major-

[1] All Danish university programs are free of tuition and offer the same small stipend to all students. Heinesen (2018) provides institutional details on the Centralized Admission System.

institution combinations on the application form.

We use CAS application data from 1996 to 2007. In these years, 52 percent of applicants listed more than one priority.[2] Among these students, Figure 7 shows the range of AI relevance for their listed priorities. The takeaway from the figure is that a substantial share of students applies to majors with widely differing relevance for AI firms. In Figure 7a, 27 percent of students apply to programs that differ in AI producer relevance by more than five percentage points (0.9 of a standard deviation). For AI user relevance in Figure 7b, the corresponding share of students that apply for programs that differ by more than five percentage points (0.6 of a standard deviation) is 32 percent. These differences indicate that many students could have been assigned to a program of very different AI relevance had admission capacities of different programs been different.

Figure 7: AI Relevances of Priorities on College Applications

(a) By AI Producer Relevance of College Major

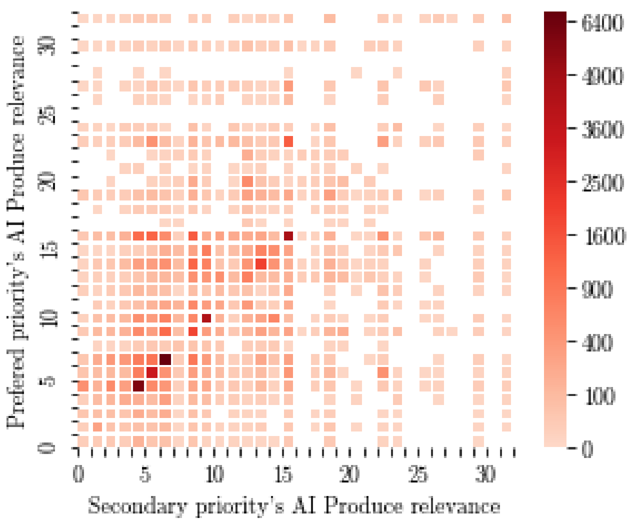


(b) By AI User Relevance of College Major

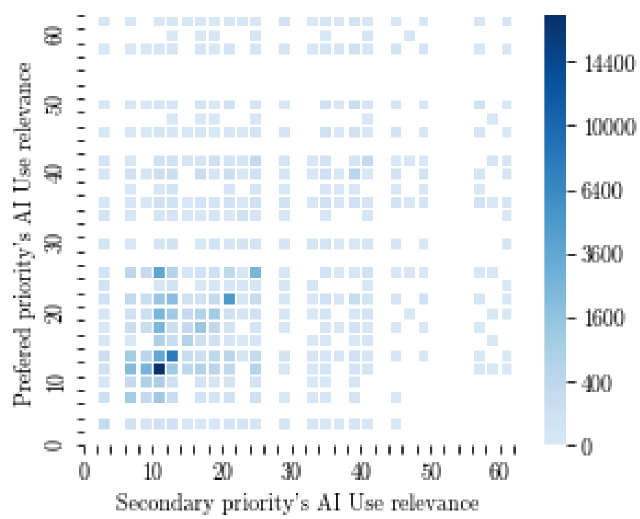


*Notes*: Color represents the total number of applications. The data is all applications from 1996-2007 with more than one priority and AI relevance measures. If there are more than two priorities, we select the two with the largest difference in AI relevance. In Panel (a), we exclude the outlier 'IT Informatics' (51 percent). Sample size is in Panel (a) is n=84,508 and in Panel (b) n=85,3687.

[2]Appendix Figure C.4a shows that the distributions of AI relevance in our analysis sample are representative of the population of applicants in CAS.

## 4.2 Regression Discontinuity Framework

This section explains how we exploit discontinuities generated by admission cutoffs to estimate causal effects of being offered a seat in a college major that is more or less AI relevant.

In the CAS, admission offers are made sequentially where the order is determined by the applicants' GPA: the highest-ranked applicant receives an offer for her preferred program; the second-highest applicant receives an offer for her highest-ranked program among the remaining slots, and so on. This is repeated until either slots run out, or applicants run out.[3] If slots run out, a cutoff is set equal to the GPA of the last applicant accepted to the program.[4]

As students are choosing between many different majors and prioritize them differently, it is not possible to assert whether crossing an admission threshold will increase or decrease AI relevance of the offered program without information on the next-best alternative program. Hence, we follow the approach in Kirkeboen et al. (2016) and take next-best choices into account.

We form pairs of priorities using the following rules: i) The preferred program must have excess applicants and thus a GPA cutoff; the second-best may have a cutoff. ii) The cutoff of the preferred program must be higher than the second-best. iii) An applicant must not be admitted to a higher priority than the preferred program in the pair. iv) The two majors must differ in terms of AI relevance.

Table 1a illustrates an application with three priorities. The three priorities result in two potential pairs: 1st/2nd and 2nd/3rd. The applicant has a GPA of 10.0 and is accepted to her first priority, program A.[5] The relevant next-best program is her second priority, program B. This is an example where the AI relevance is highest in the preferred program. Table 1b provides an example where the applicant is accepted in her second priority, program A. In this example, the more AI relevant major is ranked as next-best. The applicant is offered a seat in this AI relevant major because she does not meet the GPA cutoff of her preferred program, D.

[3] The allocation mechanism used in the CAS corresponds to a serial dictatorship, which is Pareto efficient and strategy proof. Serial dictatorship mechanisms are used in many college admission systems around the world, including Norway (Kirkeboen et al., 2016) and Chile (Hastings et al., 2013).

[4] The cutoffs are lax, since programs also are allowed to admit students on other criteria (*Quota Two*). In our sample, these are 18.5 percent of admissions.

[5] GPA is measured on a 10-point scale from 0 to 13 (00,03,5,6,7,8,9,10,11,13). The lowest passing GPA is 6, and 8 represents an 'average' performance.

Table 1: Application examples

(a) High AI Relevance of Preferred

| Applicant GPA = 10.0 | | | | |
| --- | --- | --- | --- | --- |
| Priority | Program | AI Relevance | Cutoff | |
| 1 | A | 0.17 | 9.9 | Preferred |
| 2 | B | 0.10 | 9.5 | Next-best |
| 3 | C | 0.15 | 8.0 | N/A |

(b) High AI Relevance of Second-best

| Applicant GPA = 10.0 | | | | |
| --- | --- | --- | --- | --- |
| Priority | Program | AI Relevance | Cutoff | |
| 1 | D | 0.08 | 10.5 | Preferred |
| 2 | A | 0.17 | 9.9 | Next-best |

The admission system creates discontinuities in the likelihood of being offered a seat around the GPA cutoffs; see Appendix Figure C.3. These exact cutoffs are unpredictable at the time of application as they depend on the number of seats, applicants, and GPAs. Our identifying assumptions are standard for RD designs. We assume that there are no other changes occurring at the thresholds that could confound our estimates; see Lee (2008) and Lee and Lemieux (2010) for a thorough treatment. Figure 8 represents robustness checks, showing no manipulation in the density of the running variable (GPA minus cutoff) and smooth behavior of predicted earnings near the threshold.

Figure 8: Manipulation Around the Cutoff

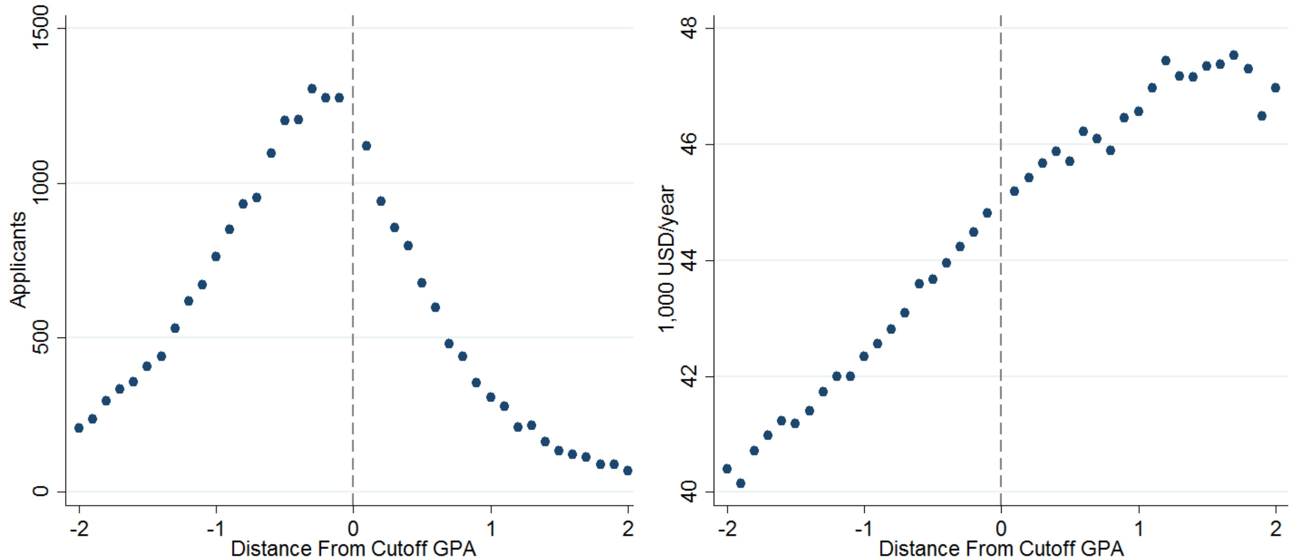


(a) Applicant Density

(b) Predicted Earnings

*Notes:* We pool all admission cutoffs and normalize so that zero on the x-axis represents the admission cutoff. The sample is all priorities of programs with cutoffs. (b) Variables used in prediction: age at the time of application, gender, and high school GPA.

Our main estimating equation is

$$y_i = \alpha + \beta Z_i + \alpha_1 d_i + \alpha_2 d_i^2 + (\alpha_3 d_i + \alpha_4 d_i^2) \times Z_i + \eta X_i + \mu_i \tag{1}$$

where $y_i$ represents an outcome of applicant $i$, such as degree attainment or earnings. $d_i$ denotes the distance from applicant GPA to the program cutoff. The indicator variable $Z_i$ equals one if the applicant's GPA is above the program cutoff and zero otherwise. Our parameter of interest is the causal effect of admission threshold crossing, $\beta$. We use a bandwidth of two grade points, and a second-order polynomial interacted with $Z$ to allow for different slopes on each side of the cutoff. We also include gender and age as a vector of predetermined control variables, $X$. In practice, these control variables have very little effect on our RD estimates, and mainly serve to improve precision.

## 4.3 AI Relevance of Completed College Major

Figure 9 shows how crossing an admission threshold impact the AI relevance of the completed major eight years later. The discontinuities around the thresholds correspond to the RD estimates of $\beta$ in Equation (1), which we also report in Table 2.(b).

Panels (a) and (b) show the impact on AI producer relevance of completed major. Crossing an admission threshold to a more (less) AI producer relevant program raises (lowers) the AI relevance of a student's completed degree by 1.1 (1.9) percentage points on average. As the discontinuities in offered programs are 2.6 and 2.9 percentage points (Panel (a) of Table 2), this implies a passthrough rate from offered seat to completed degree of around 40 to 60 percent. Panels (c) and (d) show the threshold crossing effects on AI user relevance, which imply similar passthrough rates of around 30 to 70 percent.

To be able to calculate an average realized AI relevance, Figure 9 focuses on applicants who have completed some bachelor's degree eight years later. Thus, a natural follow-up question is whether higher AI relevance comes the price of lower bachelor's completion rate. In Appendix Figure C.5, we show that there is no discontinuity in completion rates, eight years later, when crossing the admission threshold to a more AI relevant degree. However, we estimate an increase in completion rates of around five percentage points when crossing the threshold to a less AI relevant major. Put differently, it is only students who do not want a more AI relevant major whose completion rates suffer when pushed toward these AI relevant programs.

The estimated passthrough rates from offered seat to completed degree are centered around 50 percent,[6] which is sizable and suggests that adjusting admission cutoffs for different majors may be an effective way to increase the AI relevance of college graduates.[7] These findings are noteworthy, given that we in Section 3 documented that AI producer relevance requires hard skills that could make it difficult for a marginal student to complete a program.

[6] In the Danish admission system, several institutional features will tend to make the passthrough rate low. For example, applicants are free to reapply again next year or apply to their preferred program via the quota 2 system (see footnote 4)

[7] To evaluate any specific policy intervention that adjusts admission cutoffs, for example, by controlling the admission capacities of different majors, one would need to take into account where in the AI relevance distribution the marginal students are drawn in from. This would require a reweighting of the underlying marginal treatment effects relative to the local average treatment effects identified by our estimation strategy. See Mogstad et al. (2018) for a formal framework for using instrumental variables to draw inference about policy-relevant treatment parameters.

Figure 9: AI Relevance of Completed Degree and Threshold Crossing

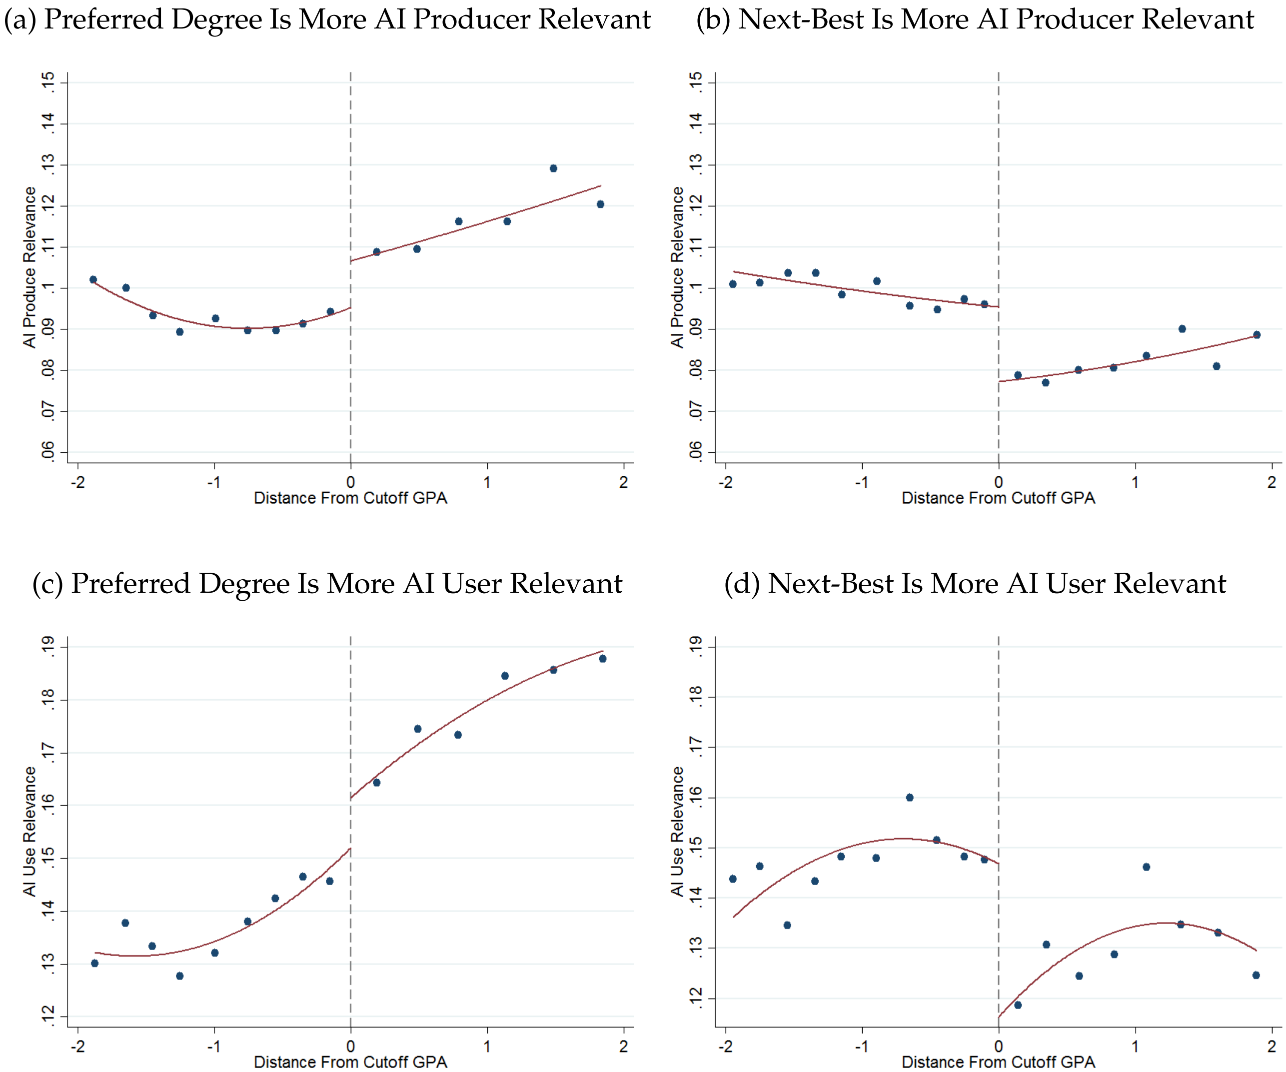


*Notes*: The figures plot the mean of AI relevance of completed degrees, eight years after application, within bins of GPA minus cutoff, and fit estimated quadratics using all the underlying data. Subsample of applicants who eight years later has completed a bachelor's degree categorized with the AI relevance measure. Excluding applicants exactly on the cutoff.

## 4.4 AI Relevance and Earnings

Figure 9 shows how crossing admission thresholds impact early-career salaries of college graduates. Panel (a) of the figure shows that crossing the threshold to a more AI producer relevant major leads to a statistically significant increase in yearly earnings of $3,770. In comparison, Panel (b) shows that crossing the threshold to a less AI producer relevant degree leads to an insignificant drop in yearly earnings of $1,692. In contrast, in Panels (c) and (d), we estimate that the causal returns to higher AI user relevance is a precisely estimated zero. Table 2.(c) reports the corresponding regression estimates.

Table 2 shows a noteworthy asymmetry in the effects of AI producer relevance, whereby students have higher bachelor's completion rates and receive larger earnings gains of AI producer relevance if they also prefer the AI relevant degree.[8] The asymmetry in completion rates can only partly explain the asymmetry in earnings, which persists even when we condition on students who complete their bachelor's studies. This estimated asymmetry in earnings effects of AI producer relevance is consistent with students sorting on individual comparative advantages, as documented in the Norwegian setting by Kirkeboen et al. (2016). If applicants sort according to comparative advantages, it can explain why students experience larger gains from AI producer relevance when they also rank these programs highest. From a policy perspective, these findings suggest that it may be more effective to expand AI relevant programs to students who prioritize them rather than to students who list them as an alternative.

To put the effects of AI producer relevance into perspective, in Appendix Table C.2, we compare them to the OLS earnings correlations estimated throughout the sample, and not just around the admission thresholds. For students who rank their AI relevant major as preferred, the causal effect of AI producer relevance is noticeable larger than the correlation. For students who rank the AI relevant major as next-best, the causal effects of AI producer relevance are similar to those suggested by the correlation. These findings suggest that the earnings premiums of AI producer relevance documented in Section 3 reflect causal effects of obtaining AI relevant college skills. The findings highlight that college education may be a powerful tool to prepare future generations for work with AI.

[8] Figure C.5 showed that the asymmetry in completion rate only applies to students who do not prefer the AI relevant degree

Figure 10: Earnings and Threshold Crossing

(a) Preferred Degree Is More AI Producer Relevant

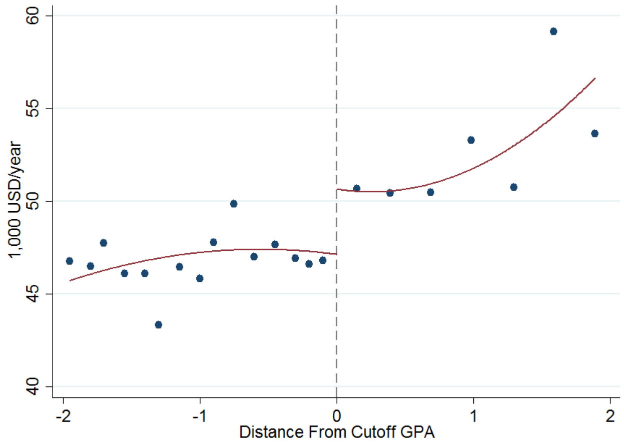


(b) Next-best Is More AI Producer Relevant

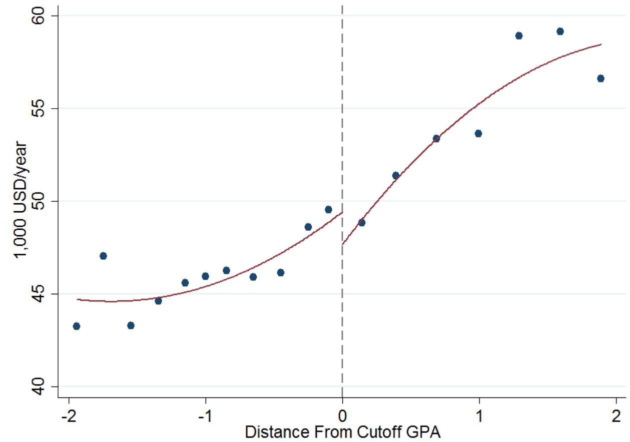


(c) Preferred Degree Is More AI User Relevant

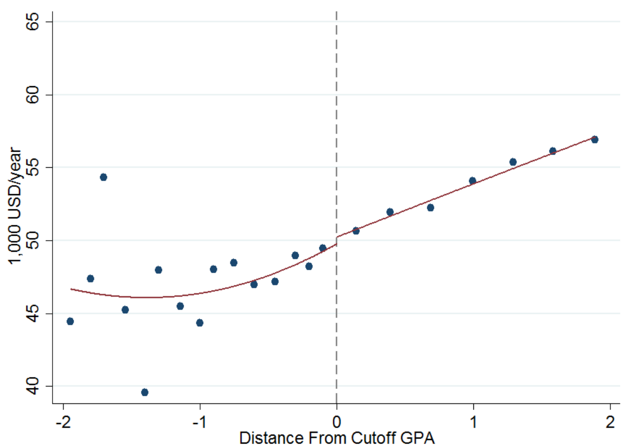


(d) Next-best Is More AI User Relevant

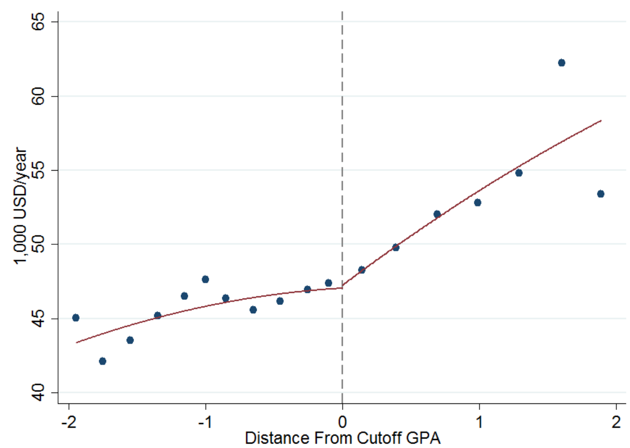


*Notes*: The figures plot the mean of earnings within bins of GPA minus cutoff, and fit estimated quadratics using all the underlying data. Excluding applicants exactly on the cutoff and restricting to earnings above zero. Earnings are the median of wage income plus self-employed business income in the years seven to nine years after application. Income variables are deflated using the CPI with 2015 as the base year.

Table 2: Threshold Crossing Effects

(a) AI Relevance of Offered Degree

| | High AI Preferred | | Low AI Preferred | |
|---|---|---|---|---|
| | Producer Relevance | User Relevance | Producer Relevance | User Relevance |
| | (1) | (2) | (3) | (4) |
| RD Estimate | 0.0261 | 0.0294 | -0.0325 | -0.0477 |
| | (0.0029) | (0.0042) | (0.0026) | (0.0050) |
| Mean below cutoff | 0.066 | 0.133 | 0.0936 | 0.133 |
| Observations | 8550 | 11665 | 12328 | 9221 |

(b) AI Relevance of Completed Degree

| | High AI Preferred | | Low AI Preferred | |
|---|---|---|---|---|
| | Producer Relevance | User Relevance | Producer Relevance | User Relevance |
| | (1) | (2) | (3) | (4) |
| RD Estimate | 0.0112 | 0.00915 | -0.0182 | -0.0321 |
| | (0.0035) | (0.0048) | (0.0030) | (0.0056) |
| Mean below cutoff | 0.0922 | 0.140 | 0.0981 | 0.149 |
| Observations | 5990 | 8100 | 8633 | 6530 |

(c) Earnings

| | High AI Preferred | | Low AI Preferred | |
|---|---|---|---|---|
| | Producer Relevance | User Relevance | Producer Relevance | User Relevance |
| | (1) | (2) | (3) | (4) |
| RD Estimate | 3770.4 | 609.84 | -1592.3 | 357.93 |
| | (1593.6) | (1378.9) | (1378.9) | (1494.2) |
| Mean below cutoff | 47151.6 | 47334.7 | 46517.6 | 46062.3 |
| Observations | 8550 | 11665 | 12328 | 9221 |

*Notes:* Estimation of equation (1) with gender and age as control variables. (b) is a subsample of applicants who, eight years later, has completed a bachelor's degree categorized with the AI relevance measure. Excluding applicants exactly on the cutoff. Bandwidth is +/- 2.0 grade points.

In summary, the findings in this section suggest that AI producer relevance may be a valuable target for college admission policies. However, how do the AI producer earnings effects compare to other targets that policymakers could follow when deciding which programs to allocate slots? A commonly used target in the Danish context is to rank programs according to their early-career average salaries (Kvalitetsudvalget, 2014). To compare our estimates to this benchmark, Appendix C.3 reruns our RD analysis, now ranking the majors according to their average starting salaries. Interestingly, we find that pushing students according to AI producer relevance yields larger causal earnings returns than simply pushing them toward degrees with higher average salaries. This suggests that it indeed is possible to acquire AI producer relevant skills in college, and that the earnings premiums related to these college skills are less driven by selection than the overall earnings differences between majors. This is a noteworthy finding, as AI producer relevance require quantitative skills that could be difficult to acquire late in life without prior training or inclination.

# 5 Concluding Remarks

This paper makes two contributions to our understanding of how artificial intelligence relates to the jobs of college majors. First, we use new micro data from Denmark to rank college majors according to whether their graduates work in AI firms. Second, we provide causal estimates on the effects of pushing students toward more AI relevant degrees.

A takeaway from our study is that traditional fields of study is an imperfect taxonomy for understanding AI relevance, and that AI cuts through the group of STEM degrees in particular. In further work, we welcome efforts to understand better what exact course contents are needed to prepare college students for work with AI.

An upshot of our regression discontinuity analysis is that adjusting admission cutoffs may be a powerful policy margin to push students toward more AI relevant skills. Our estimated causal effects inform policies that aim to increase the AI relevance of college students by adjusting admission capacities at different college majors, and we welcome further efforts to specify and evaluate such policy interventions.

Finally, while we estimated significant positive returns to the college majors that are relevant for AI producers, we also documented that men currently are heavily overrepresented in these same programs. To ensure that the fruits of AI are shared equally among men and women, we encourage further efforts to increase the female representation in the AI producer majors.

# References


Acemoglu, D. and P. Restrepo (2018). Artificial Intelligence, Automation and Work.

Acemoglu, D. and P. Restrepo (2020, jun). Robots and Jobs: Evidence From US Labor Markets. *Journal of Political Economy*.

Alekseeva, L., J. Azar, M. Gine, S. Samila, and B. Taska (2019). The Demand for AI Skills in the Labor Market. *SSRN Working Paper*, 1–38.

Brynjolfsson, E., D. Rock, and C. Syverson (2018). Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. *NBER Working Paper 24001*.

Goldfarb, A., B. Taska, and F. Teodoridis (2019). Could Machine Learning Be a General-Purpose Technology? Evidence from Online Job Postings.

Goldin, C. and L. F. Katz (1998). The Origins of Technology-Skill Complementarity. Technical Report 3.

Hastings, J. S., C. a. Nielson, and S. D. Zimmerman (2013). Are Some Degrees Worth More Than Others? Evidence From College Admission Cutoffs in Chile. *National Bureau of Economic Research Working Paper Series 19241*, 1–50.

Heinesen, E. (2018). Admission to Higher Education Programmes and Student Educational Outcomes and Earnings: Evidence From Denmark. *Economics of Education Review 63*, 1–19.

Kirkeboen, L., E. Leuven, and M. Mogstad (2016). Field of Study, Earnings and Self-selfselection. *The Quarterly Journal of Economics 131*(3), 1057–1111.

Kvalitetsudvalget (2014). Nye Veje Og Høje Mål. Technical report.

Lee, D. S. (2008). Randomized Experiments From Non-Random Selection in U.S. House Elections. *Journal of Econometrics 142*, 675–697.

Lee, D. S. and T. Lemieux (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature 48*(2), 281–355.

Mogstad, M., A. Santos, and A. Torgovitsky (2018). Using Instrumental Variables for Inference About Policy Relevant Treatment Parameters. *Econometrica 86*(5), 1589–1619.

President's Council of Advisors on Science and Technology (2012). Report To the President Engage To Excel: Producing One Million Additional College Gr Aduates W Ith Degrees in Science, Technology, Engineer Ing, and Mathematics. Technical report, Washington D.C.

Trajtenberg, M. (2018). AI As The Next GPT: A Political-Economy Perspective. *The Economics of Artificial Intelligence: An Agenda*, 175–186.

Warzynski, F. (2020). AI Adoption in Denmark.

Webb, M. (2020). The Impact of Artificial Intelligence on the Labor Market.

World Economic Forum (2018). 10 Skills You'll Need to Survive the Rise of Automation — World Economic Forum.

# A Stylized Facts on Artificial Intelligence and College Majors

Figure A.1: Share Employed in Jobs with Extreme Exposure to AI Adoption (Webb, 2020)

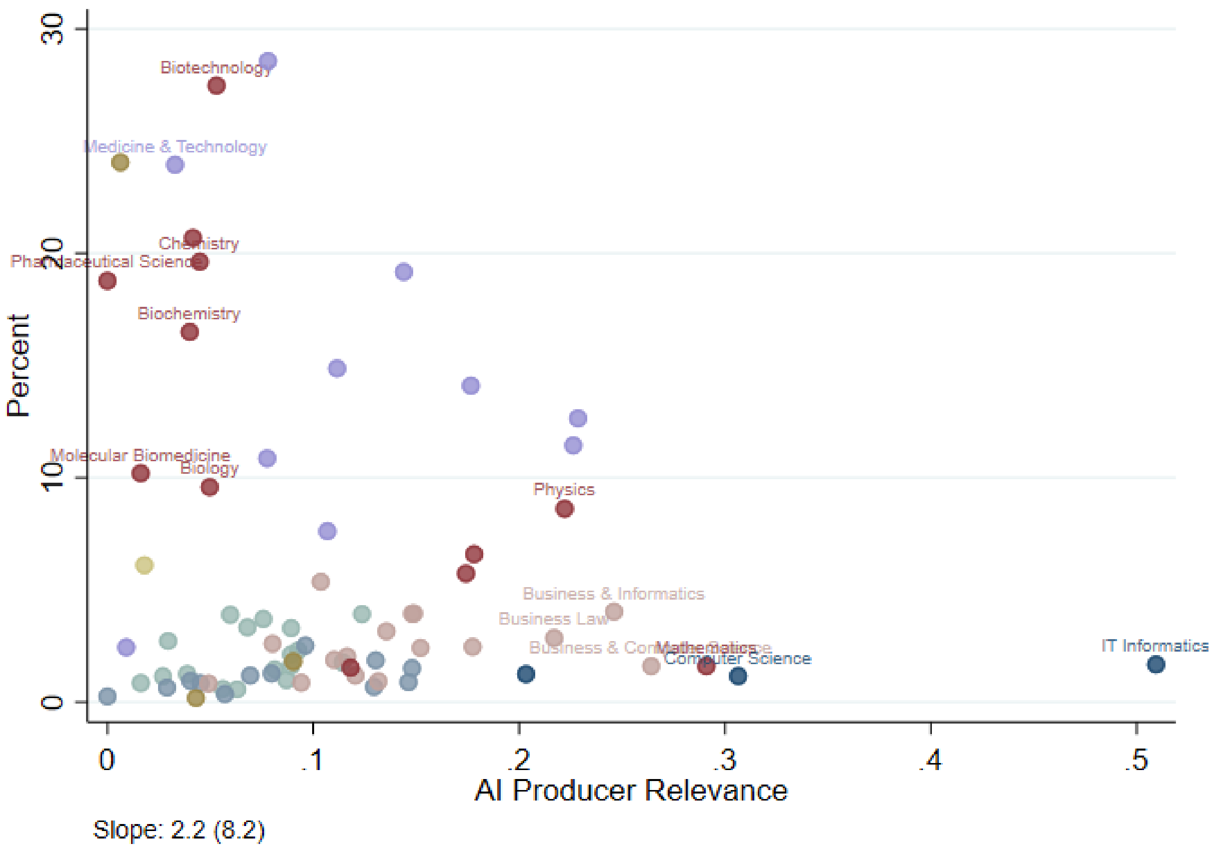


*Notes:* This figure shows the share of workers employed in occupations with extreme exposure to AI technologies by AI producer relevance of their college major. Extreme exposure to AI adoption is defined as occupations (4-digit ISCO88) with AI exposure scores of 95/100 or above. See Webb (2020) for details about the AI exposure measure. Figure 3 shows the corresponding figure by AI user relevance.

Table A.1: Earnings Premiums of AI Relevance and AI Firms

(a) Starting Salary in 2016 and AI Producer Relevance

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| AI Producer Relevance | 52.9 | | 44.7 | 47.3 |
| | (16.1) | | (15.1) | (18.6) |
| AI Producer Firm | | 9.2 | 6.6 | 8.4 |
| | | (1.7) | (1.6) | (4.7) |
| AI Producer Relevance × AI Producer Firm | | | | -11.4 |
| | | | | (23.0) |
| Observations | 2,424 | 2,424 | 2,424 | 2,424 |
| $R^2$ | 0.028 | 0.018 | 0.037 | 0.037 |

(b) Starting Salary in 2016 and AI User Relevance

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| AI User Relevance | 31.6 | | 18.6 | 23.2 |
| | (10.8) | | (10.3) | (11.3) |
| AI User Firm | | 11.7 | 10.1 | 12.7 |
| | | (2.0) | (1.8) | (3.2) |
| AI User Relevance × AI User Firm | | | | -12.9 |
| | | | | (12.5) |
| Observations | 2,424 | 2,424 | 2,424 | 2,424 |
| $R^2$ | 0.016 | 0.033 | 0.038 | 0.038 |

*Notes:* This table shows estimates from regressions of starting salaries on AI relevance of college majors (Row 1), an indicator for working in an AI firm (row 2), and the interaction of AI relevance and working in an AI firms (row 3). Starting salary is defined as salary eight years after the start of a bachelor's degree. Salary includes wage compensation for employees and net profits for self-employed. Zero incomes are included. Standard errors are clustered at the college major level.

Figure A.2: Male Share and AI Relevance of College Majors

(a) AI Producer Relevance (b) AI User Relevance

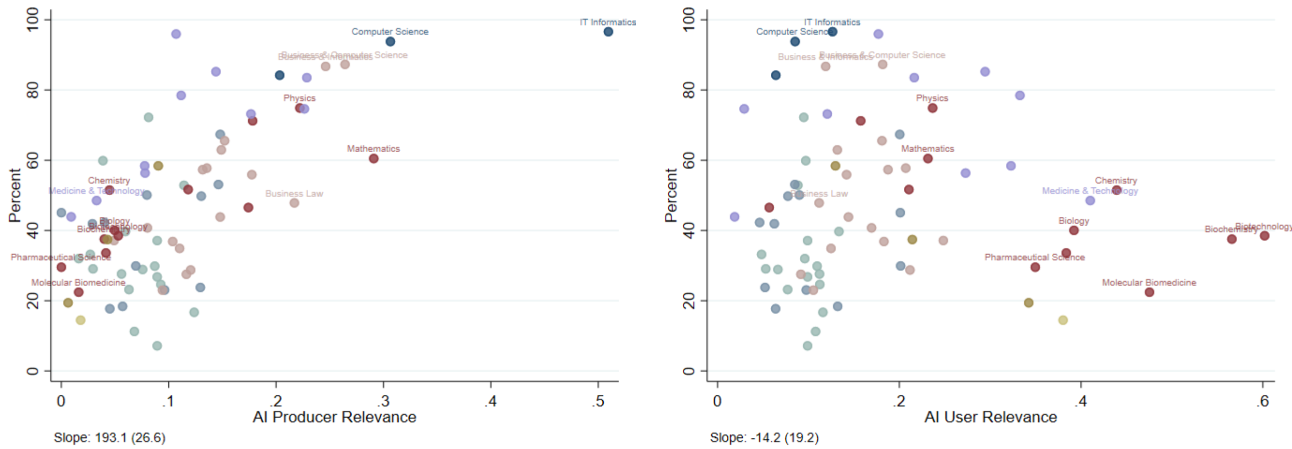


*Notes*: This figure shows the share of males by AI relevance of college majors. For the estimated slope coefficients, college majors are weighted by their number of graduates, and standard errors are clustered by college majors. Majors are colored according to their one-digit ISCED code; see Figure 2 for a color description.

Figure A.3: Annualized Growth in Starting Salaries from 2007 to 2012

(a) By AI Producer Relevance of College Major (b) By AI User Relevance of College Major

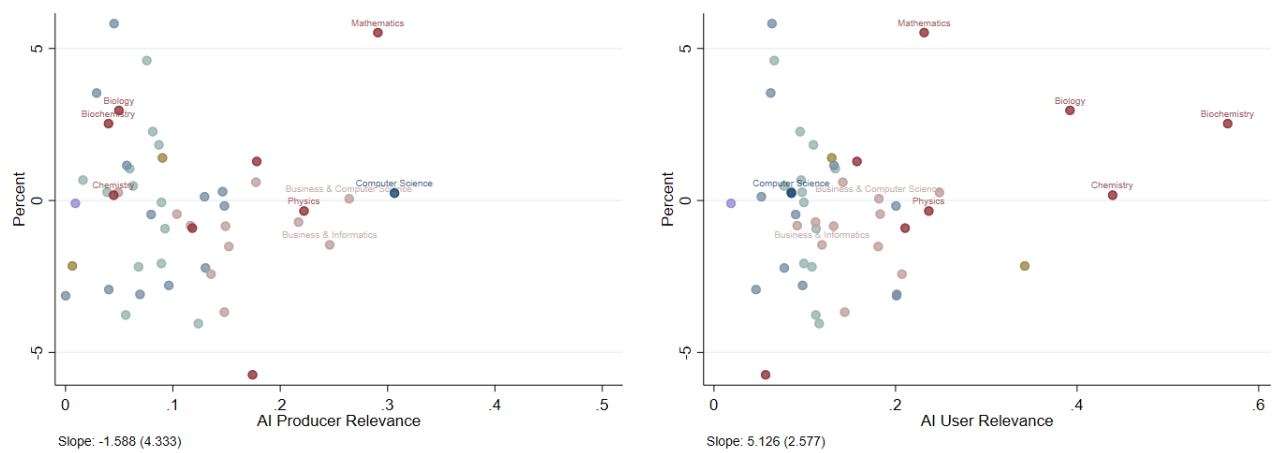


*Notes*: This figure shows the average annualized growth rates in starting salaries of college graduates from 2007 to 2012 by AI relevance of their major. Salary includes wage compensation for employees as well as net profits for self-employed. Zero incomes are included. Salaries are deflated using the CPI with 2015 as the base year. For the estimated slope coefficients, college majors are weighted by their number of graduates, and standard errors are clustered by college majors. Majors are colored according to their one-digit ISCED code; see Figure 2 for a color description.

# B A Micro Foundation for the AI Relevance Measures

This section provides a task-based micro foundation for the AI relevance measures proposed in Section 2. In our stylized model, the economy consists of three types of firms: AI users ($u$), AI producers ($p$), and non-AI firms ($n$).

AI users produce by combining a set of production tasks $i$

$$Y_u = \exp\left(\int_0^1 X_u(i)di\right) \tag{2}$$

Production tasks can be carried out by either labor or AI. We follow Acemoglu and Restrepo (2020) and assume that AI can automate a fraction $\theta$ of tasks

$$X_u(i) = \begin{cases} L_u(i) + \gamma(i) \times \mathrm{AI}(i) & \text{if } i \leq \theta \\ L_u(i) & \text{if } i > \theta \end{cases} \tag{3}$$

where $\gamma(i)$ denotes the relative productivity of AI in task $i$, and $L_u$ is the labor input bundle of AI user firms

$$L_u(i) = \prod_{s \in \mathcal{S}} L_{us}^{\alpha_{us}}(i) \tag{4}$$

AI is developed, implemented, and maintained by AI producer firms, who thus reinstate the demand for workers that are employed in these firms

$$\mathrm{AI}(i) = Y_p(i) = \prod_{s \in \mathcal{S}} L_{ps}^{\alpha_{ps}}(i) \tag{5}$$

Our incorporation of AI production is similar to Goldin and Katz (1998), who specify that the adoption of advanced machinery requires specialized inputs from skilled machinists.

Final output in the economy is a bundle of output from AI-user and non-AI firms

$$Y = Y_u^{\beta} Y_n^{1-\beta}$$

where non-AI firms operate a Cobb-Douglas production technology with skill shares $\alpha_{ns}$.

Product and labor markets are competitive with real wages $W_s$. Following Acemoglu and Restrepo (2020), we assume that AI is cost-effective in the tasks that are technologically feasible to automate.[9]

From cost minimization, the demand for workers of skill $s$ is

$$W_s L_s = Y\left((1-\beta)\alpha_{ns} + \beta((1-\theta)\alpha_{us} + \theta\alpha_{ps})\right)$$

Finally, also following Acemoglu and Restrepo (2020), we model technical progress in AI as a shift in $\theta$. The impact of AI on the demand for workers of skill $s$ is then

$$d\log(W_s L_s) = \underbrace{d\log(Y)}_{\substack{\text{Productivity}\\ \text{Effect}}} - \underbrace{R_{us} \times \frac{d\theta}{1-\theta}}_{\substack{\text{Displacement}\\ \text{Effect}}} + \underbrace{R_{ps} \times \frac{d\theta}{\theta}}_{\substack{\text{Reinstatement}\\ \text{Effect}}} \tag{7}$$

where $R_{us}$ and $R_{ps}$ denote the AI relevance measures proposed in Section 2. To recall, AI relevance is defined as the share of workers of skill $s$ who are employed in AI firms.[10]

Equation (7) shows that AI impacts the demand for worker skills through three channels: a positive productivity effect, a negative displacement effect, and a positive reinstatement effect. The equation delivers three key insights about how AI affects the demand for skills. First, the productivity effect, which is positive because AI is cost-effective at the margin, is shared uniformly across workers of different skills. This effect will lift all boats and cancels out when we compare the relative demands for different skill groups. Second, AI user relevance $R_{us}$ is the theory-consistent measure of a skill group's exposure to the negative displacement effects of AI automation, $-\frac{d\theta}{1-\theta}$. Finally, AI producer relevance $R_{ps}$ is the theory-consistent exposure measure to the positive reinstatement effects coming from AI production, $\frac{d\theta}{\theta}$.

[9]Formally, for AI to be cost-effective, we assume that

$$\prod_{s\in\mathcal{S}} \left(\frac{W_s}{\gamma(i)\alpha_{ps}}\right)^{\alpha_{ps}} \leq \prod_{s\in\mathcal{S}} \left(\frac{W_s}{\alpha_{us}}\right)^{\alpha_{us}}, \quad \forall i \leq \theta \tag{6}$$

[10]According to our theory, the AI relevance measures have the following structural expressions

$$R_{ps} = \frac{L_{ps}}{L_s} = \frac{\beta\theta\alpha_{ps}}{(1-\beta)\alpha_{ns} + \beta((1-\theta)\alpha_{us} + \theta\alpha_{ps})}$$
$$R_{us} = \frac{L_{us}}{L_s} = \frac{\beta(1-\theta)\alpha_{us}}{(1-\beta)\alpha_{ns} + \beta((1-\theta)\alpha_{us} + \theta\alpha_{ps})}$$

# C The Danish Centralized Admission System (CAS)

## C.1 CAS Data

We use CAS data from 1996-2007 merged with register data from Statistics Denmark (DST). Programs in nursing, teaching, child care, and social work are not offered by universities, and we do not consider these. In CAS, each program has a *kotnr*-code. These are not easily converted to the field codes of DST (*audd/udd*). Furthermore, *kotnr* is not stable over time and only unique within each application year. Hence we crosswalk from *kotnr* to *audd* with the following procedure. We compare the *kotnr* with the enrolled programs in DST's education register (*KOTRE*) for all accepted applicants each year. A student does not necessarily appear in the offered program in the DST register as it tracks actual enrollment in October. Not all students accept the offered seat, and some are able to enroll in other programs with empty slots. Hence, we crosswalk to the most prevalent *audd* code. We only crosswalk *kotnr*s when there is: i) At least ten applicants ii) At least 50 percent of the students are enrolled in the most prevalent *audd* code. Figure C.1 shows the distribution of match quality; the fraction of students with the most prevalent code.

Figure C.1: Crosswalk Match Quality Distribution

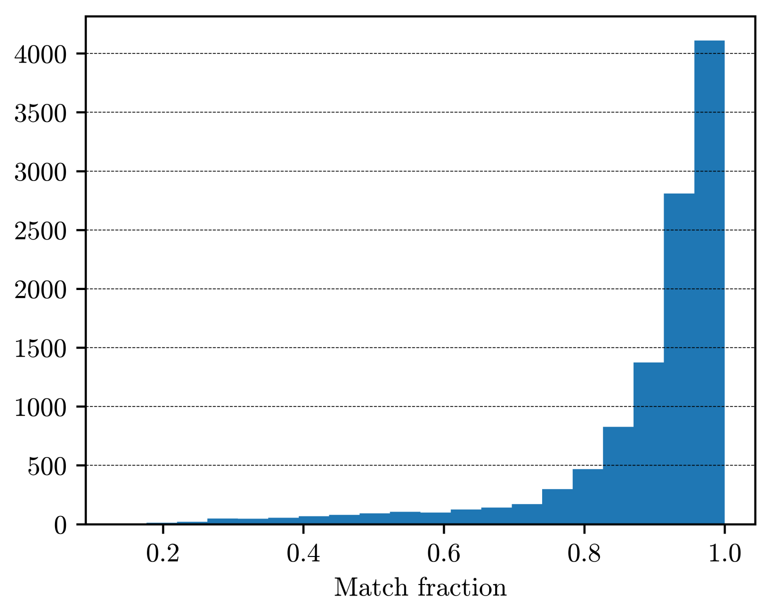


*Notes:* The figure shows a histogram fraction of students who are enrolled in the most typical field code (*audd*) by CAS program code (*kotnr*). We only crosswalk codes with a better match fraction than 0.5.

Table C.1: Applications in CAS

| | |
|---|---|
| Complete applications 1996-2007 | 274,067 |
| High School GPA available | 233,502 |
| Income data available | 229,811 |
| More than one priority | 119,715 |
| At least one priority with cutoff | 102,599 |
| *Forming pairs within applications:* | |
| Relevant pairs | 56,941 |
| Both priorities crosswalked and classified in VITA firm survey | 37,327 |
| *Final samples:* | |
| i) Priorities differ in AI produce relevance | 26,842 |
| ii) Priorities differ in AI use relevance | 26,850 |

*Notes:* Complete applications are applications with at least one priority crosswalked to a university program and does not include students with standby priority from the year before (as these have priority regardless of cutoffs). Only applicants with high school GPA and income data in DST's registers. This implies that we only use applicants with a Danish High School exam, who are living in the country seven to nine years later. Whenever a program has fewer applicants than slots; there is no cutoff. The amount of relevant pairs is the result of the first three rules described in section 4. The loss of observations in the crosswalk to AI relevance step is mainly due to certain *audd*-codes not being well represented in the VITA firm survey.

## C.2 Figures and Tables

Figure C.2: Offered Program and Threshold Crossing

(a) Preferred Degree Is More AI Producer Relevant

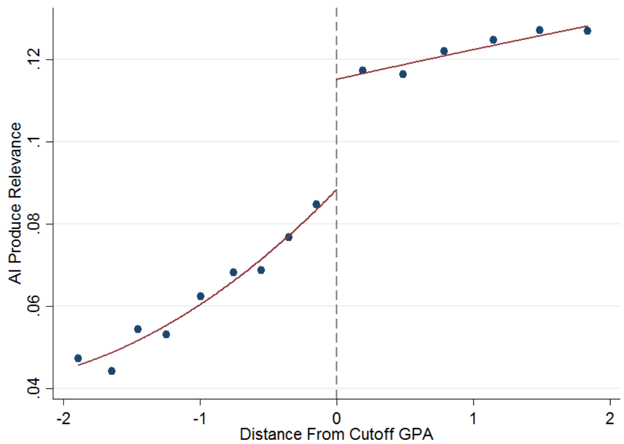


(b) Next-Best Is More AI Producer Relevant

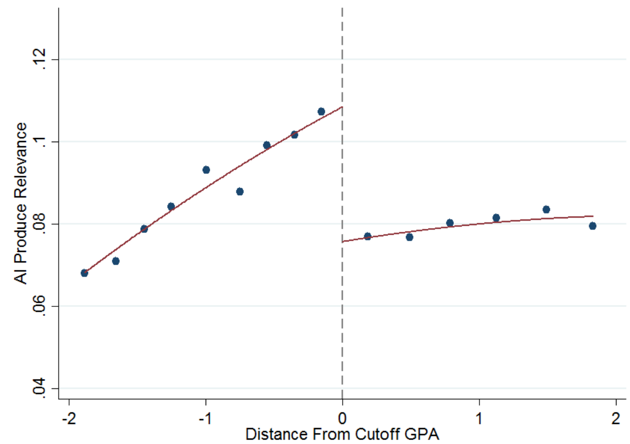


(c) Preferred Degree Is More AI User Relevant

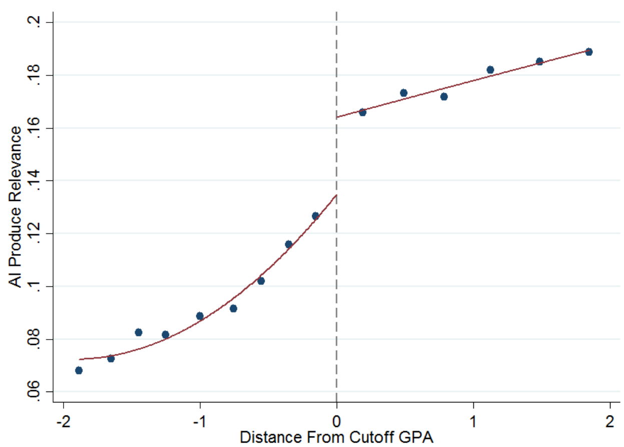


(d) Next-Best Is More AI User Relevant

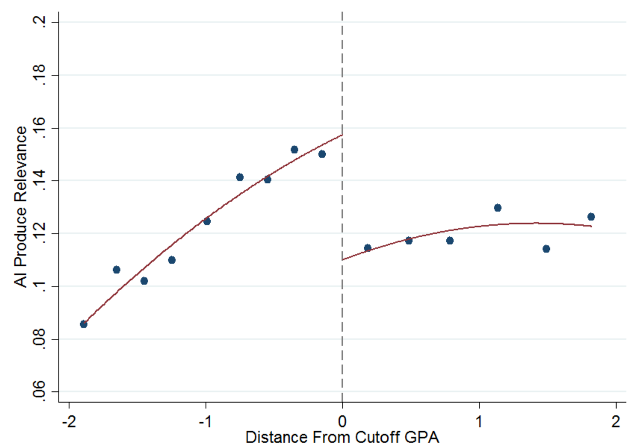


*Notes*: The figures plot the mean of AI relevance of offered degrees from the CAS system within bins of GPA minus cutoff, and fit estimated quadratics using all the underlying data. Excluding applicants exactly on the cutoff.

Figure C.3: Distance to Cutoff and Acceptance

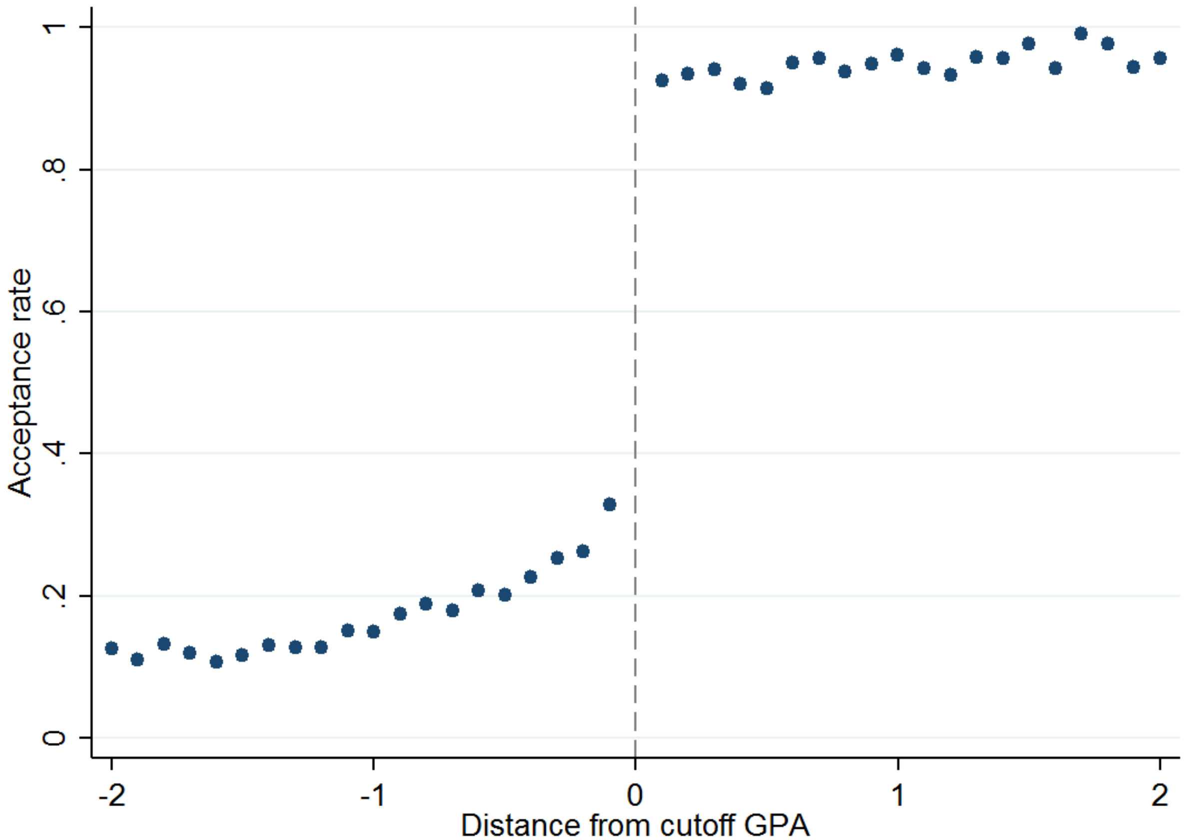


*Notes:* The figure shows the mean of an indicator for acceptance binned by tenths of a grade point on the x-axis. If students also apply for a standby seat (granted enrollment latest one year later), the standby cutoff is used. The accepted applicants below the cutoff are Quota Two. Applicants not accepted above the cutoff are due to lack of program-specific prerequisites. We are forced to exclude applicants exactly on the cutoff, as we only have GPAs and cutoffs with the first one decimal place in our data. CAS uses a lottery if there is a tie on the second decimal place.

Figure C.4: Comparison of AI relevance distributions

(a) AI Produce (b) AI Use

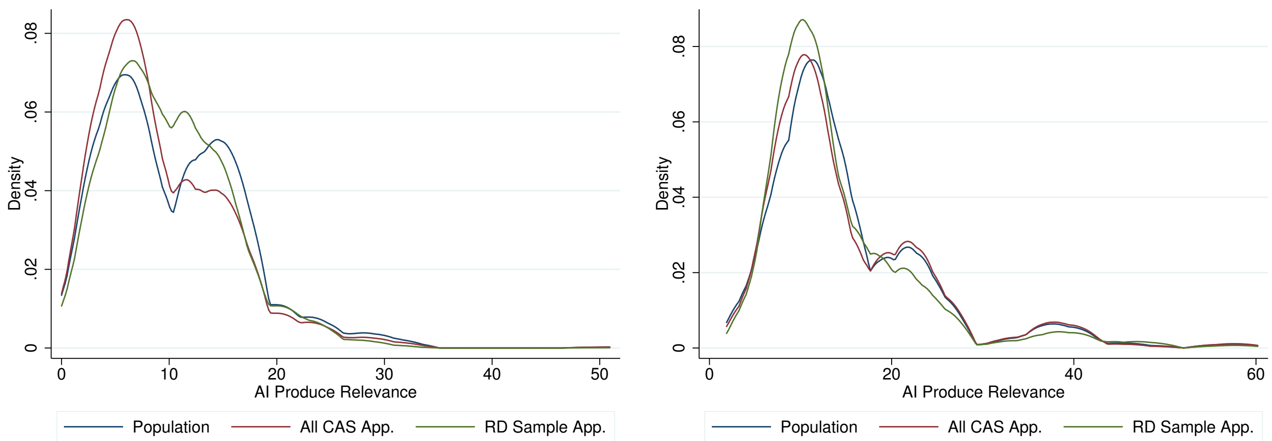


*Notes*: The figures plot epanechnikov kernel density estimates with a bandwidth of 2. The blue is the distribution of AI relevance among all bachelor's graduates in the country. The red line is all priorities listed by all applicants in the CAS system 1996-2007. The green line is the restricted CAS dataset used for RD estimation.

Table C.2: Correlation Between Earnings and AI Relevance in RD Samples

| | High AI Preferred | | Low AI Preferred | |
|---|---|---|---|---|
| | (1) Produce | (2) Use | (3) Produce | (4) Use |
| OLS Estimate (scaled) | 1065.8 | 442.3 | 1547.5 | 708.4 |
| | (75.43) | (34.63) | (106.9) | (115.3) |
| Observations | 5990 | 8100 | 8633 | 6530 |

*Notes*: Estimation of a linear regression of earnings and AI relevance within RD samples. We use gender, age at application, and high school GPA as control variables. To compare with the results in Table 2c, we scale the OLS coefficient with the change in AI relevance of completed degrees from Table 2b. Subsample of applicants who eight years later has completed a bachelor's degree categorized with the AI relevance measure. Excluding applicants exactly on the cutoff.

Figure C.5: Completion of Degree and Threshold Crossing

(a) Preferred Degree Is More AI Producer Relevant

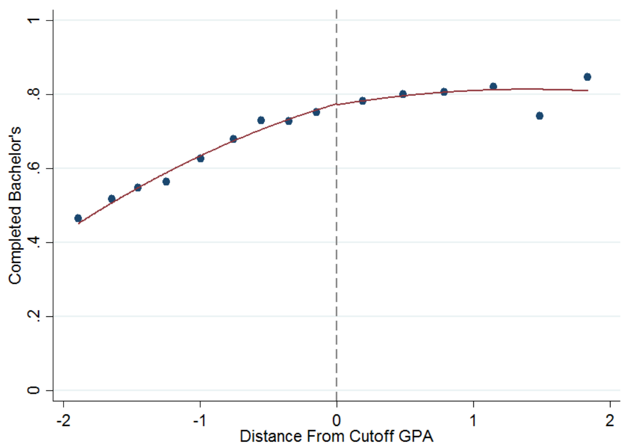


(b) Next-Best Is More AI Producer Relevant

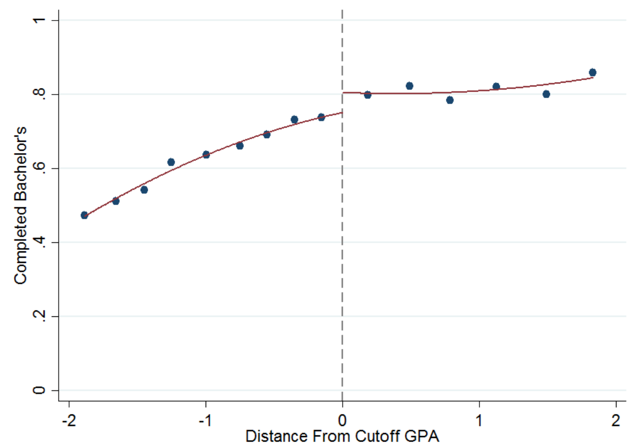


(c) Preferred Degree Is More AI User Relevant

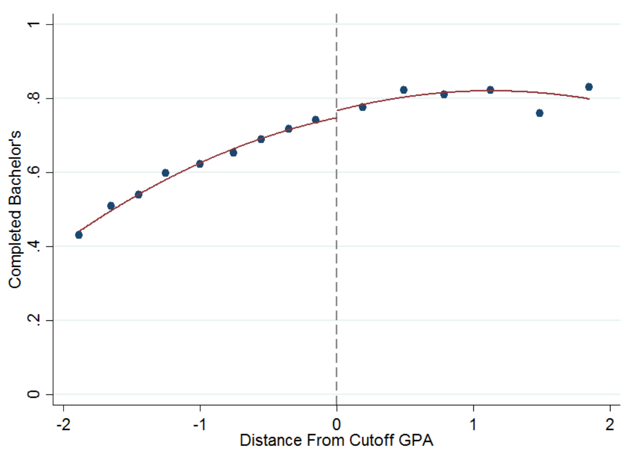


(d) Next-Best Is More AI User Relevant

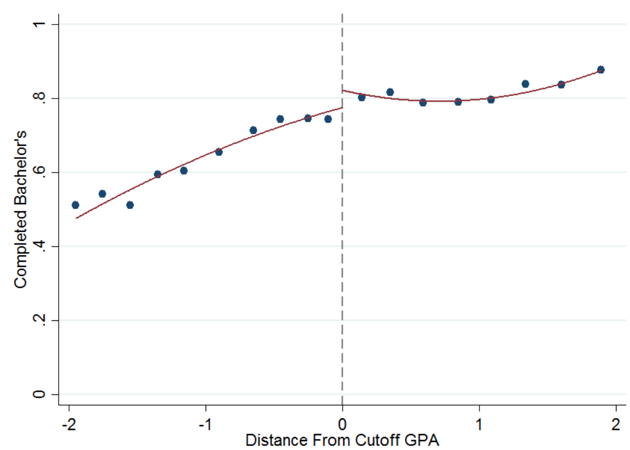


*Notes*: The figures plot completion rate of any bachelor's program, eight years after application, within bins of GPA minus cutoff, and fit estimated quadratics using all the underlying data. Excluding applicants exactly on the cutoff.

## C.3 Causal Effects of Admitting Students to Majors With Higher Average Earnings

A useful benchmark for the causal earnings effects of AI relevance is to be more agnostic about what types of degrees that lead to higher earnings. In this section, we rerun our RD analysis, now ranking majors according to their average earnings.

Instead of forming pairs with different AI relevances, we form pairs where average earnings differ between the preferred and next-best priorities. When the preferred program has the highest expected earnings, the difference in earnings between the preferred and second-best programs is $15,562 on average. In comparison, when ranking on AI producer relevance, this difference is only $3,788. Put differently, the agnostic ranking pushes students between majors of very different average earnings.

However, in Figure C.6 and Table C.3, we estimate that the causal effects of being admitted to a higher (lower) earnings degree only increase average earnings by $2,098 ($-452). The estimates, which are not statistically different from zero, are smaller than the AI producer relevance premiums estimated in Table 2c, despite the larger difference in average earnings between priorities. These results suggest that pushing students according to AI producer relevance, yields larger causal earnings returns than simply pushing them toward degrees with higher average salaries. One interpretation of the discrepancy between the causal earnings effects and the average earnings differences between majors is that earnings premiums on the AI producer margin are less driven by selection than the general earnings premiums between college majors.

Figure C.6: Earnings and Threshold Crossing - Expected Earnings Ranking

(a) Preferred is Higher Earning

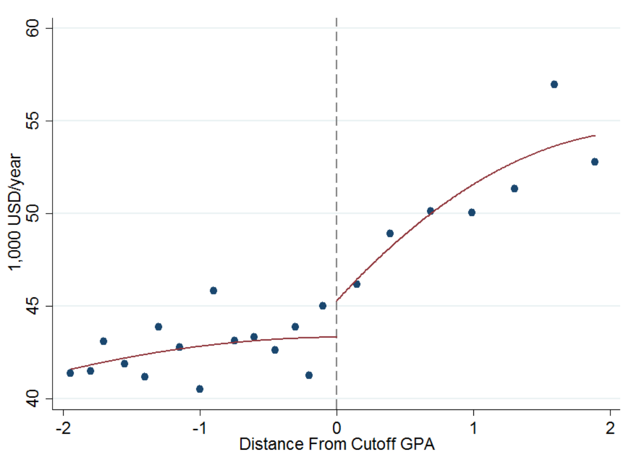


(b) Next-Best is Higher Earning

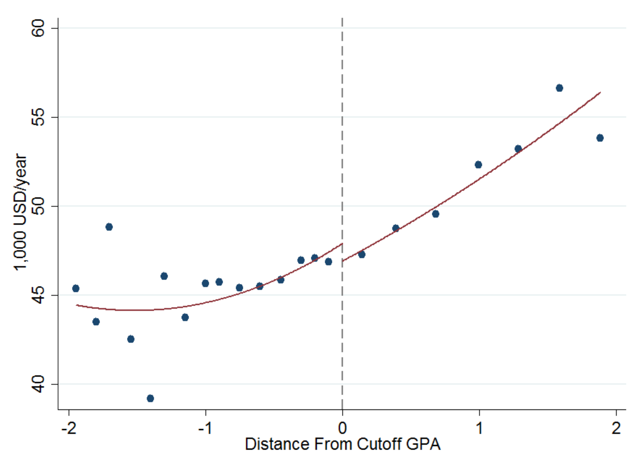


*Notes*: The figures plot the mean of earnings within bins of GPA minus cutoff, and fit estimated quadratics using all the underlying data. Excluding applicants exactly on the cutoff and restricting to earnings above zero. Earnings are the median of wage income plus self-employed business income in the years seven to nine years after application. Income variables are deflated using the CPI with 2015 as the base year.

Table C.3: Earnings and Threshold Crossing - Expected Earnings Ranking

| | High Earning Major Preferred | Low Earning Major Preferred |
|---|---|---|
| | (1) | (2) |
| RD Estimate | 2097.7 | -451.8 |
| | (1343.0) | (1189.7) |
| Mean below cutoff | 42898.6 | 45486.0 |
| Observations | 12208 | 15850 |

*Notes:* Estimation of equation (1) with gender and age as control variables. Excluding applicants exactly on the cutoff. Bandwidth is +/- 2.0 grade points.