Final Report: Reinforcing Learning? Student Responses to AI Labour Market Disruption

Kartik Akileswaran
kartik.akileswaran@post.harvard.edu

Andrei Andronic
andrei.andronic8@gmail.com

Aristotle Vossos
av2595@columbia.edu

Bouke Klein Teeselink
bouke.klein_teeselink@kcl.ac.uk

Supervised Program for Alignment Research — Spring 2026

Abstract

As large language models (LLMs) reshape the labour market, a less-explored but equally important area of concern is the educational sector: are prospective students adjusting their educational choices in response? If they are, these decisions will shape the future supply of skills in the workforce; if they are not, does this signal a possible lag in adjusting their expectations? Both cases could provide insights for policy intervention. We use the public release of ChatGPT in November 2022 as an exogenous shock. We construct an AI exposure index linking degree programmes to graduate occupations and combine it with UCAS application data covering 2019–2025. Using a Two-Way Fixed Effects (TWFE) framework, we estimate both baseline difference-in-differences and dynamic event study specifications and find that applications to more highly AI-exposed fields have increased relative to less exposed fields since 2022, with effects growing over time. We find no impact on whether students choose to attend university overall. These findings suggest that students may view AI exposure as an opportunity rather than a threat, though underlying pre-treatment demand trends in high-exposure fields such as computer science may also contribute to the observed pattern.

Keywords AI · Education · Skills · Labour markets · Economics of AI · Societal Impact of AI

1 Introduction

The release of ChatGPT in November 2022 brought the abilities of large language models (LLMs) into the public consciousness, while their integration into the workplace has generated widespread anxiety about the future of work and whether AI will displace human labour (The Economist, 2026). Within months, millions of workers gained access to tools capable of drafting documents, writing code, analysing data, and performing countless other cognitive tasks that previously required human expertise. Unlike previous automation waves that displaced physical and manual labour (Autor, Levy and Murnane, 2003; Acemoglu and Restrepo, 2022), generative AI shifts technological displacement to cognitive labour, threatening entry-level white-collar roles previously considered automation-resistant. Early evidence suggests these fears are not unfounded: firms highly exposed to LLM capabilities have reduced employment, with effects concentrated in junior positions (Klein Teeselink, 2025; Brynjolfsson, Chandar and Chen, 2025).

Thus far, most research has focused on labour demand; an equally important, but relatively understudied question is how AI affects labour supply. One of the most important factors in determining labour supply is education. What students choose to study will determine the skills they enter the labour force with, and by extension the types of roles and demands they are able to fulfill. Yet, little is known about how AI affects educational choices today.

In this paper, we study whether students adapt their educational choices in response to AI. Specifically, we examine the extent to which applications to subjects whose graduates work in more highly AI-exposed occupations have changed since the release of ChatGPT in November 2022. Answering this question will shed light on the future supply of skills in the UK and indicate whether students are anticipating labour market impacts or lagging behind them.

To measure how exposed each degree is to AI, we construct a subject-level exposure index. Because each degree is a potential pathway to many possible jobs, we first identify the occupations that graduates of each degree typically enter – for example, economics graduates may go into investment banking, consulting, or research, as well as other possible careers. We then take the AI exposure score for each of those occupations (building on the index from Gmyrek et al. 2025) and compute a weighted average to produce a single exposure score per subject. The resulting scores rank highest for fields feeding into digitised or language-intensive work such as statistics, information technology, and finance, and lowest for those feeding into hands-on or care-oriented occupations such as nursing and education.

To estimate whether students are changing their educational choices in response to AI, we combine our exposure measure with UCAS (The Universities and Colleges Admissions Service) data on UK university applications covering the period 2019–2025 and estimate a difference-in-differences model. The model compares application trends across fields with varying levels of AI exposure before and after ChatGPT's launch, with field and year fixed effects controlling for baseline differences in field popularity and common shocks to applications, respectively.

Our findings suggest that students have moved toward more highly AI-exposed fields, such as computer science, rather than away from them. At the same time, we find no effect on whether students decide to attend university overall. Our heterogeneity analyses indicate that this shift is broadly shared across gender, region, and ethnicity. Taken together, these patterns suggest that students either view AI exposure as an opportunity rather than a threat, or misinterpret early evidence regarding its labour market effects.

Our paper builds on various strands of existing literature. Human capital theory suggests that students take into account their future earnings when deciding to pursue higher education. The canonical framework laid out by Becker (1993) models education as an investment in which individuals compare expected rates of return across available options. One conclusion of this model is that insofar as future labour market prospects change, we should expect a change in what people choose to study. As the expected return on investment of their degrees changes, we would expect students to adjust their choice of study accordingly. For example, if students anticipate that AI will lead to fewer jobs in software engineering, they should adjust accordingly by moving away from fields that lead to software engineering careers.

A substantial body of evidence confirms that labour market conditions do influence field-of-study decisions. Students generally shift toward fields with better wages and job opportunities during periods of higher unemployment (Blom, Cadena and Keys, 2021). The empirical literature has consistently found, however, that students are only partially responsive to earnings signals, and that the degree of responsiveness depends on the quality of information available to them. Students hold highly inaccurate beliefs about the labour market prospects of different occupations and fields, systematically overestimating both wages and job availability, particularly for occupations they find attractive (Baker et al., 2018; Betts, 1996; Hastings et al., 2016). This overestimation is largest for students' most preferred occupations, a pattern that is consistent with motivated reasoning rather than rational updating (de Koning, Dur and Fouarge, 2025).

On the supply side of education, Conzelmann et al. (2023) and Grosz (2022) examine whether institutions adjust programme offerings and enrolment capacity in response to shifts in labour demand. Both find that institutional response is real but slow, with sub-baccalaureate programmes responding faster than four-year degree programmes and capacity adjustments lagging shifts in student demand by several years. Foote and Grosz (2020) further shows that local labour market downturns increase post-secondary enrolment and shift students toward more applied programmes, consistent with human capital theory's prediction that reduced opportunity costs during downturns increase the relative attractiveness of education. These findings imply that even if students respond quickly to AI-induced changes in career prospects, universities may be slow to expand or contract programmes accordingly, creating a period of mismatch between student demand and institutional supply that is itself of policy interest.

To link AI's influence to educational decisions, we also draw on existing research that examines AI's impact on productivity. Eloundou et al. (2024) evaluate 19,265 detailed work activities from the O*NET database and find that approximately 80% of the US workforce could have at least 10% of their tasks affected by LLMs – and that higher-income jobs face greater exposure on average. Gmyrek et al. (2025) present a framework of four progressively increasing exposure gradients, from low exposure and high task variability jobs, to high exposure and low task variability jobs, moving beyond the binary of automation versus augmentation.

Recent research also shows early signs of AI's impact on the labour market. Klein Teeselink (2025) provides UK-specific evidence, finding that highly exposed firms reduced junior employment by 0.4% per standard deviation increase in LLM exposure, while senior positions showed no significant change. Across 39 countries, Klein Teeselink and Carey (2026) find a 6.1% decline in job postings for high-exposure occupations globally. Our work thus bridges the gap between existing research on AI's impact on employment by examining whether these early results have translated into a measurable change in the types of subjects students choose to study.

While we would expect Generative AI-driven labour effects to play a role in students' choice of study, empirical literature on the subject is still nascent. Timmerman (2025) examine this question most closely, mapping college majors to occupations using AI exposure measures and finding that STEM fields are more highly exposed than non-STEM fields. However, they do not test whether students actually change their behaviour in response to this exposure. While prior research on physical automation has shown an impact on aggregate education demand — for instance, Di Giacomo and Lerch (2023) find that every additional industrial robot adopted per thousand workers increases local university enrolment by four students — evidence on how these technological shocks redirect students between fields remains limited. We bridge this gap by going beyond the decision of whether to pursue higher education, providing the first empirical evidence on what students choose to study in response to AI exposure.

2 Methodology

2.1 Data Sources and Sample

Our analysis combines three primary sources of data: raw data on UK applications from UCAS, LinkedIn profile data from Revelio Labs, and task-level exposure scores from Gmyrek et al. (2025).

Our key independent variable is a time-invariant, degree-level AI exposure index. This continuous treatment variable (ranging from 0 to 1) captures the average AI exposure of the specific labour markets that a field's graduates typically enter.

For our baseline specification, our GenAI exposure index is derived from exposure scores developed by Gmyrek et al. (2025). Their measure employs a human-AI hybrid methodology that integrates practical insights from a survey of 1,640 workers and iterative expert validation. Leveraging a highly granular database of nearly 30,000 tasks mapped to the global ISCO-08 framework, this approach mitigates the "techno-optimism" associated with purely algorithmic assessments and provides a robust approach compared to their earlier measure (Gmyrek, Berg and Bescond, 2023).

The construction of our novel exposure index relies on Revelio Labs LinkedIn data on UK workers, which provides each worker's education history alongside their full sequence of dated job-spells with occupation codes attached. We use the profile-level extract for the United Kingdom, restricted to individuals whose employment history contains at least one position starting on or after January 2021. For each spell of employment, we observe the start and end dates, the job title, the (imputed) salary, and standardised occupation codes. The educational record reports the degree level, field of study, university name and country, and graduation date for each credential listed on a profile. The two records share a common user identifier, which lets us reconstruct, for any combination of degree, field, university, and graduation cohort, the full distribution of occupational destinations its graduates hold over time. Crucially, by matching the standardised occupational codes (ISCO-08 and O*NET) of each recorded job spell to their corresponding Gmyrek et al. (2025) and Eloundou et al. (2024) task-level scores, we enrich the longitudinal employment data with empirical AI exposure metrics.

The final degree-level exposure indices are constructed by averaging these occupation-level scores across the realized career destinations of alumni. To ensure accurate representation and avoid over-weighting individuals holding multiple concurrent roles, we adopt a two-step chronological aggregation process. First, we expand each job spell across its active calendar years (up to 20 years post-graduation) and compute a within-person average of occupational exposure for each specific year. Second, we compute an across-person average for all individuals within the same graduation cohort and academic field.

Hence, our time-invariant composite indices represent the employment-weighted average of the generative AI exposure characterising the specific labour markets that graduates of each respective academic discipline systematically enter – with the variation in these degree-level scores driven entirely by cross-field differences in graduate career trajectories. Our baseline analysis generates an exposure index using the Gmyrek et al. (2025) exposure measures. However, we also perform the same process for the exposure data from Eloundou et al. (2024), where their α score captures tasks automatable by an LLM alone, while β captures those that can be performed by LLMs with software integration. As reported in Table 1 below, the resulting scores rank highest for fields feeding into digitised or language-intensive work such as statistics, information technology, and finance, and lowest for those feeding into hands-on or care-oriented occupations such as nursing and education.

Table 1: Average AI Exposure Scores by Subject Field

Subject	Baseline	Alpha (Direct)	Beta (Augmented)
Statistics	0.463	0.246	0.572
Information Technology	0.461	0.337	0.607
Finance	0.459	0.139	0.521
Economics	0.459	0.163	0.527
Accounting	0.452	0.138	0.522
Mathematics	0.451	0.229	0.555
Marketing	0.443	0.124	0.504
Business	0.434	0.146	0.503
Physics	0.420	0.233	0.533
Engineering	0.418	0.260	0.552
Law	0.405	0.143	0.487
Chemistry	0.402	0.181	0.492
Biology	0.388	0.170	0.478
Architecture	0.382	0.095	0.475
Education	0.366	0.132	0.417
Medicine	0.338	0.140	0.434
Nursing	0.288	0.098	0.392

Notes: The table reports the derived degree-level AI exposure indices for the 17 degree subjects we analyse. *Baseline* refers to the index derived from Gmyrek et al. (2025). *Alpha (Direct)* and *Beta (Augmented)* refer to the direct LLM automation and software-augmented subscores derived from Eloundou et al. (2024), respectively. All metrics are constructed using the employment-weighted averages of realized alumni occupational trajectories.

Our application data comes from UCAS, which operates the undergraduate admission process for UK universities. We use data for application volumes by undergraduate programme (academic field at a given institution) by year for 2019–2025. We restrict our sample to domestic applicants (domiciled in England, Scotland, Wales, or Northern Ireland) to avoid confounding from international student flows, whose education decisions may be influenced by other factors such as Brexit.

Our main dependent variable is the log of total domestic main-scheme UCAS applications to each field in each year. The log transformation allows us to interpret coefficients as approximate percentage changes and accounts for the large differences in baseline application volumes across fields.

On aggregate, applications for UK undergraduate education have remained robust despite the emergence of generative AI. Figure 2 illustrates the absolute volume of UCAS applications and acceptances, demonstrating a steady upward trajectory from 2018/19 that stabilises at historically high levels through 2023/24.

However, significant heterogeneity exists across individual academic disciplines. By indexing subject-level acceptances to a 2022 baseline, Figure 3 reveals divergent growth trajectories across broad CAH1 subject groups post-GPT-3.5. Hence, we confirm that on aggregate, university application growth kept pace, while subject-level variability hints at the potential impact of GenAI as well as a myriad of other factors.

Figure 1: Average Exposure Indices by Subject

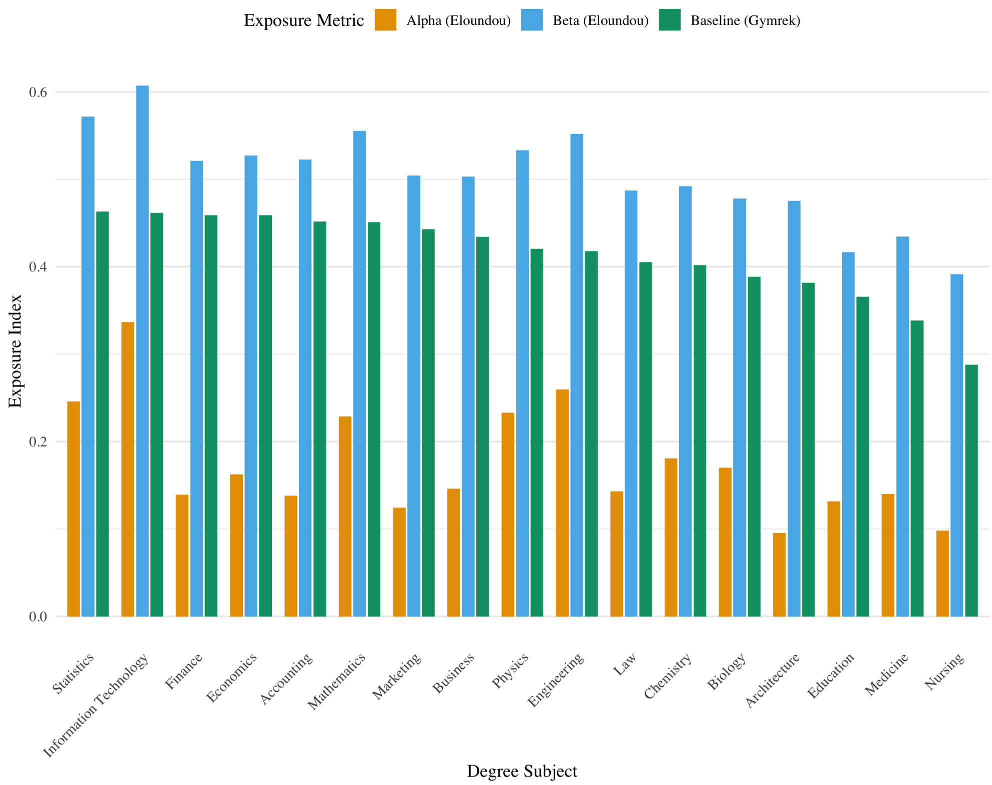

Notes: The figure visualises the Baseline, Alpha, and Beta degree-level AI exposure indices across the 17 subjects under analysis. The indices capture the average exposure of the jobs that graduates of each academic discipline systematically enter.

Figure 2: Indexed Growth of UK Undergraduate Applications (2018-2024)

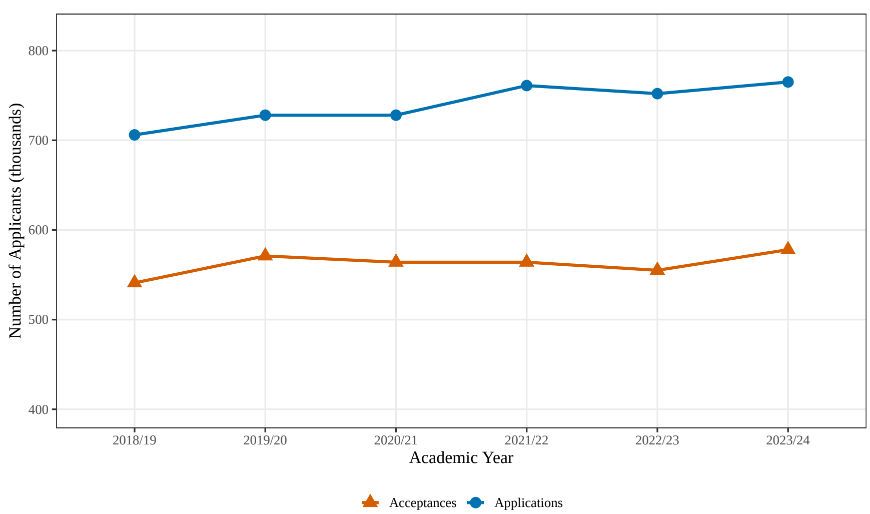

Notes: The figure plots the indexed growth of aggregate UK undergraduate applications and acceptances from 2018 to 2024. The data tracks the relative trajectory of total higher education demand and illustrates a sustained growth trend prior to and following the emergence of generative AI.

Figure 3: UK Undergraduate Subject Applications Growth Rates (Indexed)

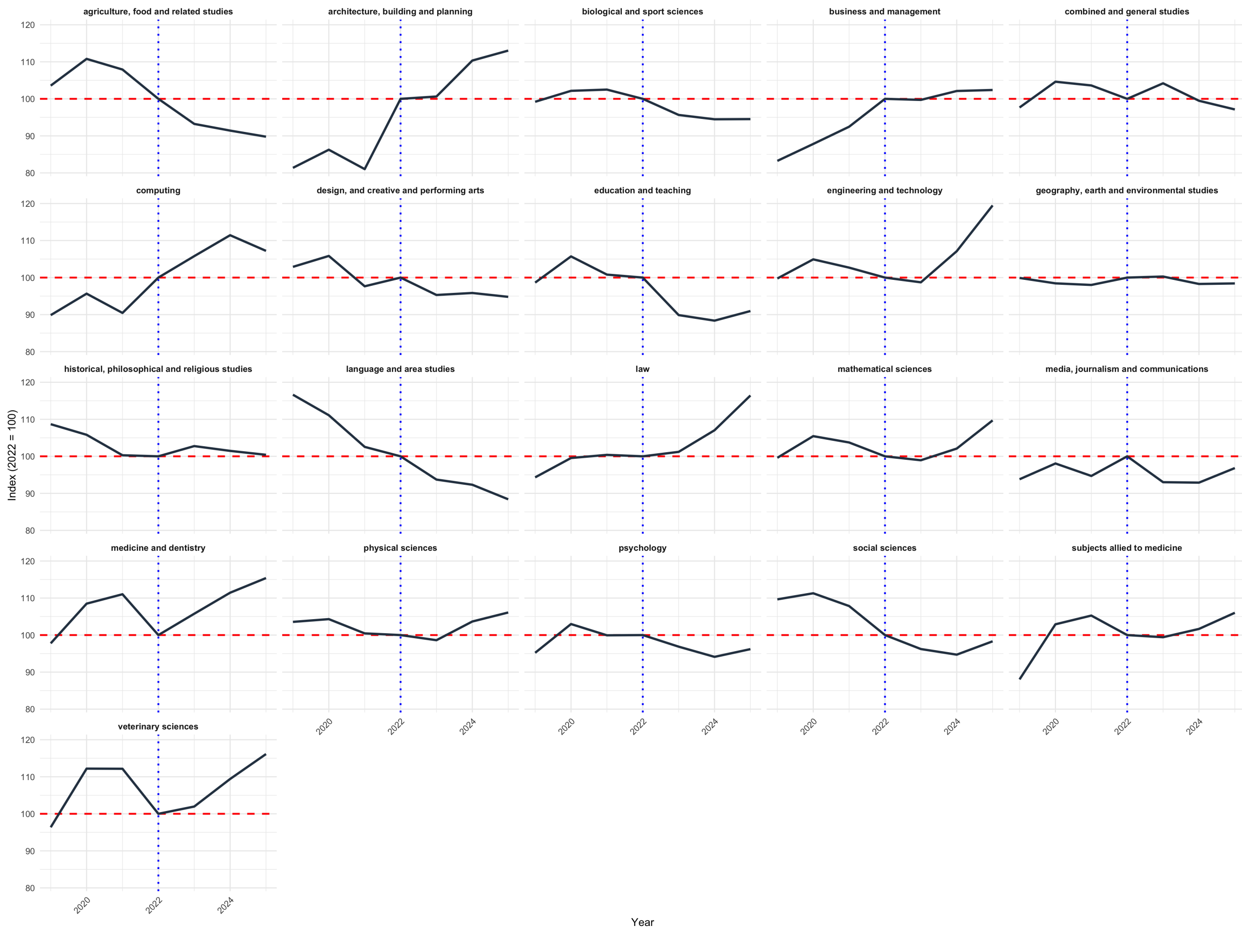

Notes: The figure visualises the relative trajectories of university acceptances across broad CAH1 subject groups. Acceptances for each discipline are indexed to their respective pre-treatment 2022 baseline (2022 = 100), indicated by the horizontal dashed red line. The vertical dotted blue line marks the late-2022 release of GPT-3.5.

UCAS classifies subjects using the Higher Education Classification of Subjects (HECoS) Common Aggregation Hierarchy at three levels of granularity, ranging from CAH1 (the broadest level) to CAH3 (the most granular). For example, CAH1 would group programmes into a category like "Business and Management", while CAH3 would use more specific groupings such as "Accounting", "Marketing", or "Finance". We manually map from CAH3 categories to the 17 subjects for which subject-level exposure measures are defined. The exposure-score subjects are a mix of granular classifications (like "Accounting", "Marketing", and "Finance") and broader classifications (like "Engineering"), so CAH3 enables us to maximise the amount of UCAS data used in our final panel dataset. While there are 106 distinct CAH3 subject classifications, our exposure scores are not defined for several subjects (such as Arts, Humanities, Geography, and others) and hence we drop UCAS records for programmes that cannot be clearly mapped. While we only mapped 36.1% of CAH3 subjects, this still enabled us to use 47% of the full UCAS dataset (see Table 2).

Table 2: UCAS Dataset Retention After Subject Mapping

	Unique CAH3 Subjects		Application Volume	
Status	Count	% of Total	Count	% of Total
Dropped (Unmapped)	106	63.9%	8,607,040	53.0%
Retained (Mapped)	60	36.1%	7,637,940	47.0%
Total Baseline Cohort	**166**	**100.0%**	**16,244,980**	**100.0%**

Hence, via the mapping of CAH3 subjects to Revelio Academic fields, we construct our final dataset. The resulting dataset is a balanced panel containing total applications for each of the 17 subject fields across the 7-year observational window (2019–2025). Each subject field is assigned its corresponding, time-invariant baseline exposure score, which serves as the continuous treatment variable when interacted with our post-2022 temporal shock dummy in the primary econometric specifications.

2.2 Estimation Strategy

We use a difference-in-differences design with continuous treatment intensity to estimate the effect of AI exposure on UK university application behaviour. Our identification strategy exploits variation in AI exposure across degree fields, combined with the public release of ChatGPT in November 2022 as a treatment point.

We select 2022 as the base year because it represents the last application cycle in which the majority of application decisions were made prior to ChatGPT's public availability; the main UCAS deadline falls in January, meaning the 2022 cohort's choices preceded the November 2022 launch. We use three years of pre-treatment data (2019–2021) and three years of post-treatment data (2023–2025).

For our main analysis, we estimate the following event study specification:

$$\log(\text{Applications}_{it}) = \sum_{\tau \neq 2022} \beta_\tau \cdot \text{Exposure}_i \times \mathbb{1}_{[t=\tau]} + \alpha_i + \delta_t + \varepsilon_{it} \tag{1}$$

where Applications_{it} denotes total domestic main-scheme UCAS applications to field i in year t. Exposure_i is the field's time-invariant AI exposure score. The term $\mathbb{1}_{[t=\tau]}$ denotes indicators for event time τ relative to the 2022 application cycle, which serves as the reference period.

We include field fixed effects to control for time-invariant differences across fields, and year fixed effects to account for common shocks to university applications. For example, field fixed effects absorb the fact that some subjects, such as computer science, have historically attracted far more applicants than others–ensuring that we do not confound baseline popularity with the effect of AI exposure. Year fixed effects account for aggregate shifts in application volumes–such as demographic changes in the size of the university-age cohort or broader economic conditions–that affect all fields equally. Standard errors are clustered at the field level.

Our identification strategy relies on the parallel trends assumption: absent the emergence of large language models, application trends across fields with different AI exposure levels would have evolved similarly, conditional on our fixed effects structure. Several features of our design support this assumption and the interpretation of our estimates as causal. First, our exposure measures are constructed from pre-treatment graduate career data, eliminating concerns about reverse causality whereby application shifts might influence measured exposure. Second, the sudden and largely unanticipated public release of ChatGPT provides a clear temporal cutoff point, after which we may assume that AI entered the public consciousness more broadly and began to influence labour markets. Third, our event study tests and confirms parallel pre-treatment trends across specifications.

To summarise the overall post-treatment effect, we also report average treatment effects on the treated from a static specification:

$$\log(\text{Applications}_{it}) = \beta \cdot \text{Exposure}_i \times \text{Post}_t + \alpha_i + \delta_t + \varepsilon_{it} \tag{2}$$

where Post_t is an indicator for application cycles from 2023 onward. Our baseline model estimates Equation 2 using the Gmyrek et al. (2025) exposure index. Further, as part of robustness checks in Section 3.2, we re-run the model using the alpha and beta scores from Eloundou et al. (2024) to examine how different definitions of AI exposure impact our results.

We perform a number of robustness checks to test the validity of our findings. For example, we re-estimate our analysis using alternative subject mapping approaches (averaging exposure scores across multi-field CAH1 groups versus our preferred CAH3 splitting approach). We also examine effects on total instead of log applications and re-run the analysis while weighting fields by number of applications. These robustness checks support our main findings.

We nonetheless acknowledge the possibility that other factors differentially affected high- and low-exposure fields around the time of ChatGPT's introduction. For instance, AI-exposed fields such as business, finance, and economics are also high-salary and quantitatively intensive, meaning our estimates could partly reflect differential responses to salary premia or shifting student preferences toward quantitative skills. Additionally, with 17 fields, our cluster count is relatively small, which may affect the reliability of cluster-robust inference.

3 Results

This section outlines our initial empirical findings on the impact of Generative AI on UK undergraduate education decisions. We first present our baseline DiD model outlined in Section 2 followed by an event study specification to verify our parallel trends assumption.

Table 3 reports the results of our regression of logarithm of applications on the outcome of interest Generative AI exposure. The model's DiD coefficient (β) of 2.344 is statistically significant at the 1% level ($p < 0.001$). Given the log-transformation used on our dependent variable, we can interpret this coefficient as a log-point change. We observe a positive relationship between exposure and application volumes, indicating that over the time period studied, fields with relatively higher GenAI exposure saw increases in applications. This relationship has a large magnitude: a 0.1-unit increase in our AI Exposure Index is associated with 26.4% more applications, on average per year, for a given subject.

To check for the validity of our causal estimates and validate our parallel trends assumption, we run an Event Study specification. Regression results with yearly coefficients are reported in Table 4 and these coefficients are visualised in Figure 4. The dashed vertical line at 2022 represents the timing of the exogenous shock (the November 2022 release of GPT-3.5) and marks the boundary between pre- and post-treatment outcomes. Visually and statistically, these results provide validation of our primary identifying assumption of parallel trends.

On the pre-treatment side of Figure 4, the point estimates for 2019 ($\beta_{2019} = 0.206$), 2020 ($\beta_{2020} = 0.019$), and 2021 ($\beta_{2021} = -0.346$) hover closely around the zero-effect line, with confidence intervals intersecting zero and rendering them statistically insignificant. This graphical evidence indicates that prior to the public release of Generative AI, a subject's baseline AI exposure score had no systematic correlation with its application growth trajectory, suggesting that the high- and low-exposure subjects were moving in parallel.

Table 3: Baseline Difference-in-Differences Estimates

Dependent Variable:	Log_Apps
Model:	(1)
Variables	
AI Exposure × Post-2022	2.344***
	(0.4399)
Fixed-effects	
Subject	Yes
Year	Yes
Fit statistics	
Observations	119
R^2	0.99485
Adjusted R^2	0.99360

Notes: The table shows the estimated effect of AI exposure on the number of applications. *Log_Apps* is the logarithm of the total number of applications. All effects are estimated using an exposure difference-in-differences design. *AI Exposure* is a degree's AI exposure, as described in Section 2.1. *Post-2022* is a dummy variable that takes the value of 1 for all years after 2022. All analyses include subject and year fixed effects.

The post-treatment regime reveals a stark shift following the exogenous shock (the November 2022 release of GPT-3.5). We observe strong positive coefficients with confidence intervals clearly detached from the zero line. Interestingly, this effect appears to be increasing over time, suggesting that the influence of AI exposure on academic decision-making has intensified as the technology has become more entrenched in the economy and education.

Our results seem counter-intuitive at first. Given both the public debate expressing concern around labour displacement due to Generative AI in jobs such as software engineering, as well as emerging empirical evidence outlined in Section 1, we might expect that the relationship between exposure and education demand is negative. One possible interpretation is that Generative AI opens up new avenues to use highly exposed degrees. For example, a computer science degree might now explicitly be sought by those interested in AI Engineering positions. Other potential factors are that Generative AI may lower the entry barrier to develop skills or apply to high-paying and high-exposure fields, or that students are initially viewing Generative AI as a 'gold rush', which inherently correlates with strong short-run demand.

Finally, past empirical literature suggests that students hold inaccurate beliefs about the labour market prospects of different occupations and fields, and systematically overestimate both wages and job availability, particularly for occupations they find attractive (Baker et al., 2018; Betts, 1996; Hastings et al., 2016). This overestimation is largest for students' most preferred occupations, a pattern that is consistent with motivated reasoning rather than rational updating (de Koning, Dur and Fouarge, 2025). If students cannot accurately perceive the AI-induced deterioration in their target fields, the behavioural response will be attenuated even if the underlying labour market signal is large.

Figure 4: Event Study - Dynamic Difference-in-Differences

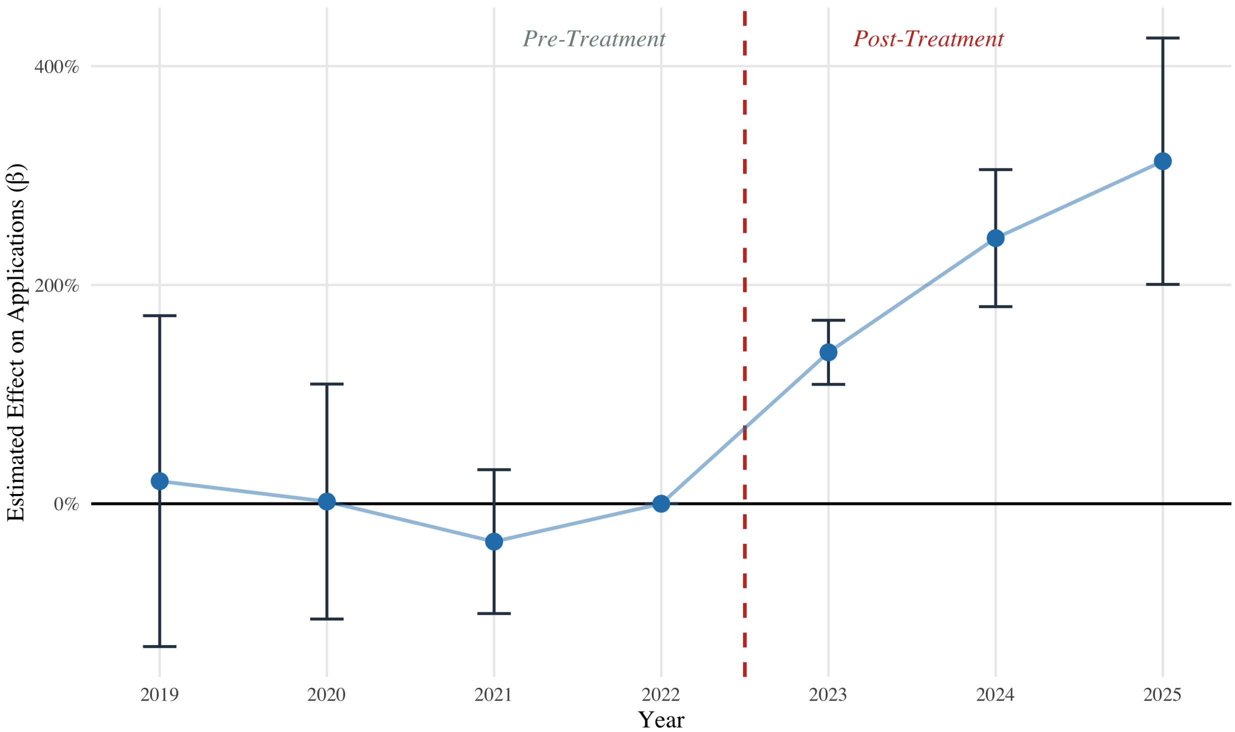

Notes: This figure plots the dynamic difference-in-differences event study coefficients for the effect of AI exposure on the log of total applications. The year 2022 serves as the omitted reference period. The vertical dashed line indicates the late-2022 release of ChatGPT. Error bars represent 95% confidence intervals using standard errors clustered at the subject level.

3.1 Heterogeneity Tests

To understand whether any subgroups of applicants are driving our main result, we estimate our difference-in-differences specification across four key factors: gender, nation (England, Scotland, Wales, or Northern Ireland), university prestige (Russell or non-Russell group), and ethnicity. For each dimension, we also estimate a pooled specification that interacts our treatment variable with group indicators, allowing us to formally test whether the differences across groups are statistically significant.

We find no meaningful differences in the response to AI exposure between male and female applicants. The estimated effect is 2.147 for men (SE 0.835, $p < 0.05$) and 2.252 for women (SE 0.642, $p < 0.01$). The formal interaction test confirms that this small difference is not statistically distinguishable from zero ($p = 0.91$), and is closely aligned with our main difference-in-differences model.

Splitting the sample by nation of domicile reveals an apparent gradient in the size of the effect: 2.342 for English applicants (SE 0.618, $p < 0.01$), 2.039 for Welsh applicants (SE 1.094, $p < 0.1$), 1.837 for Scottish applicants (SE 0.612, $p < 0.01$), and 1.388 for Northern Irish applicants (SE 0.928, not significant), although the interaction tests show that these differences are not statistically significant. Some of this attenuation likely reflects greater statistical noise in smaller-population nations, while Northern Ireland's distinctive labour market, with a smaller high-skilled service sector, a unique post-Brexit context, and the option for students to study and work in the Republic of Ireland, may also contribute to the more muted response.

We also examined whether the effect varied between more and less prestigious universities, and found a positive result only for Russell Group universities (1.473, SE 0.648, $p < 0.05$). However, the difference between Russell and Non-Russell Group institutions (1.476) was only marginally significant and reflects the relatively small number of clusters in our analysis. Moreover, data on programmes at the university level was only available with CAH1 specifications and for acceptances, rather than applications. This makes it difficult to compare this result with our main model, and may also reflect supply-side responses by Russell Group universities rather than demand-side shifts by applicants alone.

The effect of AI exposure is positive and statistically significant for every ethnic group in our data, but with a varying magnitude. Asian applicants show the smallest response (1.581, SE 0.700, $p < 0.05$), followed by Other (2.139, $p < 0.05$), Black (2.408, $p < 0.01$), White (2.595, $p < 0.01$), and Mixed (2.598, $p < 0.01$) applicants. The formal interaction test, with Asian applicants as the reference group, confirms the difference between Asian and Black applicants (0.827, $p < 0.01$) and between Asian and Mixed applicants (1.016, $p < 0.05$). The smaller response among Asian applicants may be due to Asian applicants already moving toward these fields prior to the introduction of ChatGPT, leaving less scope for further shifts in response to the AI shock. A complementary interpretation is that other ethnic groups, historically less concentrated in STEM, are converging toward these fields more rapidly.

Taken together, our heterogeneity results indicate that the overall shift toward AI-exposed fields is broadly shared across gender, region, and ethnicity.

3.2 Robustness Checks

To test the validity of our findings, we ran a total of six robustness checks.

First, while our baseline analysis uses the GenAI exposure score from Gmyrek et al. (2025), we replicated our model using the α and β exposure scores from Eloundou et al. (2024). The α score captures exposure of tasks that are automatable by standalone LLMs, while the β score captures tasks which can be automated with either standalone LLMs or LLMs with additional software integration. Our results (5) indicate the baseline result observed using the Gmyrek et al. (2025) exposure score holds up for β with a statistically significant coefficient of 1.978, but not with α. This suggests that students respond more to a more comprehensive conception of AI exposure rather than to exposure only driven by standalone LLMs. This can also be driven by the fact that a broader measure of exposure acts as a stronger treatment overall. Indeed, a key finding in Eloundou et al. (2024) is precisely that the proportion of jobs affected by GenAI rises significantly when considering complementary software integration.

Second, the baseline effect might be driven by a small set of subject outliers rather than a consistent trend across education choices. To explore this possibility, we re-estimate our baseline DiD iteratively, dropping one subject field from the regression each time in order to test coefficient stability. The point estimates (5) remain stable and statistically significant regardless of the subject excluded. This indicates that the aggregate effect we are seeing is broad-based across subjects, rather than driven by a small set of subjects.

Third, our baseline model could simply be capturing a pre-existing trend in education decisions that is driven by dynamics that pre-date ChatGPT's release. To test for this, we restricted the sample to the pre-shock time period (2019-2022) and applied a placebo treatment in the year 2021 to test if the effect was already present in the pre-treatment data. We find that the interaction coefficient (6) loses statistical significance, which supports our parallel trends assumption that AI exposure had no systematic impact on applications prior to the treatment year.

Fourth, we re-estimated our model using UCAS CAH1 data, which aggregates programmes into broader categories (see Section 2). The result maintains a positive coefficient of 0.733 (SE 0.393, p = 0.081) – directionally consistent with our baseline result – but is roughly one-third the magnitude of our CAH3 specification and only marginally significant (see Table 8). This is consistent with the loss of precision that occurs when distinct degree programmes with very different exposure profiles are collapsed under one broad category. For example, accounting and human resource management both fall under the broader CAH1 category of "Business and Management", despite their graduates entering occupations with very different AI exposure levels. This reinforces our choice of CAH3 as the preferred mapping for our main specification.

Fifth, the log transform of applications compresses large fields and amplifies small ones, so if a few small fields have volatile post-2022 counts, the baseline DiD coefficient could be artificially inflated. To address this issue, we re-run the DiD with raw application counts as the outcome, so the coefficient now measures additional applications per unit of AI exposure rather than a log-point change. The raw-count model produces a positive and statistically significant coefficient (see Table 7), which reinforces our main finding.

Sixth, our baseline OLS specification assigns equal weight to all subjects, so if the effect is concentrated in small niche fields, the headline coefficient overstates what is happening across the bulk of the applicant population. To test this, we re-run the log-applications DiD weighting each observation by the field's 2022 application count (see Table 9), giving larger fields proportionally more influence. The coefficient decreases to 2.017 but remains significant, confirming the result is not driven by idiosyncratic behaviour in small fields.

4 Discussion

Our finding that students moved toward, rather than away from, highly AI-exposed fields is initially counterintuitive given the public discourse around AI-driven job displacement. Several interpretations are consistent with these results. One possibility is that students view AI exposure as an opportunity rather than a threat, perceiving that AI complements rather than substitutes for human labour in many roles. A second, less benign possibility is that students hold lagging beliefs about labour market prospects in AI-exposed fields, perhaps systematically overestimating opportunities in fields that were previously attractive (de Koning, Dur and Fouarge, 2025). A third is that multiple shocks have taken place around the same time as the introduction of AI. These include the shift to remote/hybrid work, which began with the COVID-19 pandemic in 2020, and persistent high inflation in the UK, which began around September 2021, peaked in October 2022, and remains above the Bank of England's target range to this day. To the extent that students view their financial futures as more precarious, they may move towards safer career options, perhaps without accounting for recent trends in hiring and job opportunities in related sectors. Students' lagging beliefs and the flight to safety are not mutually exclusive explanations, and each may in fact be compounding the other's effect.

4.1 Policy Implications

For policymakers, our findings carry several important implications. The fact that students are shifting toward AI-exposed fields suggests that, contrary to some fears, AI has not deterred young people from pursuing careers in cognitive and white-collar work, for now; in fact, it appears to encourage them to pursue these careers. However, the consequences of this reallocation depend critically on whether students are responding to accurate signals about future labour market conditions. If, as the existing literature has documented, students systematically overestimate prospects in their preferred fields, today's shift toward AI-exposed degrees risks producing an over-supply of graduates to fields with declining entry-level employment, raising the threat of skills mismatch and graduate underemployment. This scenario underscores the importance of investing in labour market information and career guidance that provides students with accurate, up-to-date evidence on how AI is reshaping the occupations they aspire to enter.

Furthermore, undergraduate study in the UK may need to move in the direction of more flexibility, as AI is making labour markets and demand for skills more uncertain. Students tend to decide on which degree programmes to apply to in Year 12. Five years will pass between that decision and the day they graduate with that degree. Five years ago, ChatGPT had not yet been publicly released, but a student graduating next year will be graduating with a degree that they chose under a different set of labour market conditions than those they will encounter today. Policymakers – and universities themselves – should consider a more flexible system, where students can change degree programmes and where all students take a core set of courses that equip them with transferable skills that will be valuable in a future where AI is central in the workplace.

Educational preferences may also take time to adjust to changes in the labour market. Students in the UK select secondary school coursework to match the university programmes they intend to apply to in the future, so those applying to university in the years immediately following the introduction of GenAI may have been locked into their university subject choices – even if they would prefer to switch subjects. Therefore, the true effect of AI exposure on educational choices may lag by several years, and we might not see this in UCAS data until future application years.

4.2 Limitations

We acknowledge several limitations to our analysis. With 17 fields, our cluster count is small, which may affect the reliability of our inference. While we employ year fixed effects in an effort to absorb aggregate shocks that affected all fields equally during the study period, our estimates also cannot fully isolate the AI shock from various concurrent shocks or pre-existing dynamics that may be influencing our results. Furthermore, our exposure measure, while constructed using a rigorous task-based methodology, is rather limited. First, it captures current job content rather than how occupations may evolve over the working life of today's applicants. For example, AI may increase the optionality of various degrees – political science graduates may reasonably become software engineers, as the barrier to coding drops thanks to emerging AI tools. Second, it only captures the proportion of tasks for a given job that can potentially be completed by an LLM, so it may not truly capture a useful measure of exposure. For example, while many simple tasks may be automatable, the main value for a given role may come from a small subset of tasks that cannot yet be automated. Indeed, to truly understand potential labour market impacts, understanding the exact types of tasks that are impacted is of importance. Finally, we cannot fully isolate the demand-side response from potential supply-side adjustments by universities, which may have expanded capacity in highly exposed fields in the post-treatment period.

4.3 Future Research Directions

Several avenues for future research follow naturally from this work. First, more granular application data, containing university-by-degree-by-year data instead of degree-by-year-alone, would enable us to increase our sample size and use a more specific mapping of subjects to AI exposure, thus achieving more precision in our analysis. Second, cross-country comparisons would help establish whether the

UK pattern reflects a broader trend or particular features of the British higher education system, and our exposure index could be readily reconstructed using equivalent data for other national contexts. Third, postgraduate applicants may respond more sharply than undergraduates given their closer proximity to the labour market, and an analysis of master's-level application patterns could test this. Fourth, the distinction between augmentation and automation deserves deeper investigation, using measures designed to capture expected job impacts rather than task-level exposure alone. Additionally, as discussed in Section 4.2, better measures of exposure need to take into account the type of tasks that are impacted and their relative value in a role. Overall, measurement remains a key area requiring further research to better understand the impact of AI on education choices and labour outcomes. Last, whether students' choices reflect rational forward-looking optimisation or systematic misperception is an open question that experimental approaches that elicit beliefs about AI exposure and test whether they update in response to accurate information would be well-suited to address.

5 Conclusion

This paper examines whether UK university applicants alter their educational choices in response to the introduction of GenAI. Using a difference-in-differences design that exploits variation in AI exposure across degree fields, we find that applications to more highly AI-exposed fields have risen relative to less exposed fields following the release of ChatGPT in November 2022, with effects growing larger over time. Our headline estimate indicates that applications to fields such as computer science, mathematics, and finance – all relatively highly exposed to AI – increased more than applications to less highly-exposed fields. We find no effect on the extensive margin: students did not become more or less likely to attend university overall. Heterogeneity analyses reveal that this shift is broadly shared across gender, region, and ethnic background, but may be more concentrated at Russell Group universities. Our results are robust to a range of robustness checks (see Section 3.2). This research provides one of the first attempts to understand how AI is affecting labour supply. Our findings provide an early empirical foothold for further research that will only grow in importance as AI's capabilities – and its impact on the labour market – continue to expand.

References

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica*, 90(5): 1973–2016.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*, 118(4): 1279–1333.

Baker, Rachel, Eric Bettinger, Brian Jacob, and Ioana Marinescu. 2018. "The Effect of Labor Market Information on Community College Students' Major Choice." *Economics of Education Review*, 65: 18–30.

Becker, Gary S. 1993. *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education.* . 3rd ed., University of Chicago Press. First published 1964.

Betts, Julian R. 1996. "What Do Students Know about Wages? Evidence from a Survey of Undergraduates." *Journal of Human Resources*, 31(1): 27–56.

Blom, Erica, Brian C. Cadena, and Benjamin J. Keys. 2021. "Investment over the Business Cycle: Insights from College Major Choice." *Journal of Labor Economics*, 39(4): 1043–1082.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Stanford Digital Economy Lab Working Paper.

Conzelmann, Johnathan G., Steven Hemelt, Brad Hershbein, Shawn M. Martin, Andrew Simon, and Kevin Stange. 2023. "Skills, Majors, and Jobs: Does Higher Education Respond?" NBER Working Paper 31572.

de Koning, Bart K., Robert Dur, and Didier Fouarge. 2025. "Correcting Beliefs about Job Opportunities and Wages: A Field Experiment on Education Choices." IZA Discussion Paper 17951.

Di Giacomo, Giuseppe, and Benjamin Lerch. 2023. "Automation and Human Capital Adjustment: The Effect of Robots on College Enrollment." *Journal of Human Resources*. First published online December 2023.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384(6702): 1306–1308.

Foote, Andrew, and Michel Grosz. 2020. "The Effect of Local Labor Market Downturns on Postsecondary Enrollment and Program Choice." *Education Finance and Policy*, 15(4): 593–622.

Gmyrek, Paweł, Janine Berg, and David Bescond. 2023. "Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality." International Labour Organization ILO Working Paper 96.

Gmyrek, Paweł, Janine Berg, Karol Kamiński, Filip Konopczyński, Agnieszka Ładna, Balint Nafradi, Konrad Rosłaniec, and Marek Troszyński. 2025. "Generative AI and Jobs: A Refined Global Index of Occupational Exposure." International Labour Organization (ILO) ILO Working Paper 140, Geneva.

Grosz, Michel. 2022. “Do Postsecondary Training Programs Respond to Changes in the Labor Market?” *Journal of Human Capital*, 16(4): 461–487.

Hastings, Justine, Christopher A. Neilson, Anely Ramirez, and Seth D. Zimmerman. 2016. “(Un)informed College and Major Choice: Evidence from Linked Survey and Administrative Data.” *Economics of Education Review*, 51: 136–151.

Klein Teeselink, Bouke. 2025. “Generative AI and Labor Market Outcomes: Evidence from the United Kingdom.” King’s College London Working Paper, SSRN 5516798.

Klein Teeselink, Bouke, and Daniel Carey. 2026. “AI, Automation, and Expertise.” SSRN Working Paper 6134506.

The Economist. 2026. “Prepare for an AI jobs apocalypse.” *The Economist.* Online release May 14; May 16 print edition.

Timmerman, Jean Xiao. 2025. “Educational Exposure to Generative Artificial Intelligence.” *Federal Reserve FEDS Notes.*

A Regression Tables

Table 4: Dynamic Difference-in-Differences Estimates

Dependent Variable:	Log_Apps
Model:	(1)
Variables	
Exposure_Index × Year = 2019	0.2065
	(0.7134)
Exposure_Index × Year = 2020	0.0199
	(0.5066)
Exposure_Index × Year = 2021	-0.3464
	(0.3101)
Exposure_Index × Year = 2023	1.384***
	(0.1382)
Exposure_Index × Year = 2024	2.428***
	(0.2954)
Exposure_Index × Year = 2025	3.131***
	(0.5312)
Fixed-effects	
Subject	Yes
Year	Yes
Fit statistics	
Observations	119
R^2	0.99536
Adjusted R^2	0.99391

Clustered (Subject) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table reports the dynamic event study difference-in-differences estimates. The dependent variable, *Log_Apps*, is the logarithm of total applications. The omitted reference year is 2022. *Exposure_Index* represents our time-invariant field-level AI exposure index (based on Gmyrek et al. (2025) task-level exposure scores). All models include subject and year fixed effects.

B Robustness Checks

Table 5: Robustness Check: Alternative Exposure Measures

Dependent Variable:		Log_Apps	
Model:	(1)	(2)	(3)
Variables			
DiD_Gmyrek	2.344***		
	(0.4399)		
DiD_Alpha		0.8734	
		(0.5259)	
DiD_Beta			1.978***
			(0.4612)
Fixed-effects			
Exposure_Field	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	119	119	119
R^2	0.99485	0.99270	0.99459
Within R^2	0.35454	0.08546	0.32190

Clustered (Exposure_Field) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table compares the baseline model's exposure index using Gmyrek et al. (2025) exposure scores versus the index based on the α (representing pure LLM automation) and β (representing LLM plus software integration automation) exposure scores from Eloundou et al. (2024). Model (1) reports the baseline result from Table 3. Models (2) and (3) report the results using an index created with the *Alpha* and *Beta* exposure scores from Eloundou et al. (2024). *Post-2022* is a dummy indicating years following the ChatGPT release. All specifications include subject and year fixed effects with standard errors clustered at the subject level.

Table 6: Placebo Check in Time (Fake Treatment = 2021)

Dependent Variable:	Log_Apps
Model:	(1)
Variables	
AI Exposure × Post-2021 (Placebo)	-0.2864
	(0.5337)
Fixed-effects	
Subject	Yes
Year	Yes
Fit statistics	
Observations	68
R^2	0.99718
Adjusted R^2	0.99598

Clustered (Subject) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents a temporal placebo test restricting the sample to the pre-treatment period (2019–2022). A pseudo-treatment dummy, *Post-2021 (Placebo)*, takes the value of 1 for the years 2021 and 2022 to test for pre-existing trends prior to the release of ChatGPT. The dependent variable is the logarithm of total applications. The model includes subject and year fixed effects, with standard errors clustered at the subject level.

Table 7: Robustness Check: Log vs. Absolute Applications

Dependent Variables:	Log_Apps	Total_Apps
Model:	(1)	(2)
Variables		
AI Exposure × Post-2022	2.344***	161,922.1***
	(0.4399)	(46,635.6)
Fixed-effects		
Subject	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	119	119
R^2	0.99485	0.98194
Within R^2	0.35454	0.26217

Clustered (Subject) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table compares the baseline log-linear specification with a level-level specification to assess the absolute magnitude of the demand shift. Model (1) uses the logarithm of applications as the dependent variable. Model (2) uses the raw, unlogged count of total applications. Both models include subject and year fixed effects, with standard errors clustered at the subject level.

Table 8: CAH1 Robustness

	Dependent variable:
	log(Applications)
Post x AI Exposure	0.733*
	(0.393)
Field FE	Yes
Year FE	Yes
Observations	119
Within R2	0.0933
Note:	*p<0.1; **p<0.05; ***p<0.01

Notes: The table reports the estimated effect of AI exposure on log applications using the broader CAH1 subject classification instead of the more granular CAH3 mapping used in our baseline specification. *Post x AI Exposure* is the interaction between the field-level AI exposure index and a post-2022 indicator. The model includes subject and year fixed effects, with standard errors clustered at the subject level.

Table 9: Robustness Check: Sample Weighting

Dependent Variable:	Log_Apps	
Model:	(1)	(2)
Variables		
AI Exposure × Post-2022	2.344***	2.017***
	(0.4399)	(0.2151)
Fixed-effects		
Subject	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	119	119
R^2	0.99485	0.99103
Within R^2	0.35454	0.47560

Clustered (Subject) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table reports results from sample weighting to account for baseline field size differences. Model (1) is the unweighted baseline difference-in-differences estimate. Model (2) estimates the same specification using Weighted Least Squares (WLS), where observations are weighted by their pre-treatment application volume in 2022. Both models include subject and year fixed effects, with standard errors clustered at the subject level.

Figure 5: Robustness Check: Leave-One-Out Regression for Each Subject

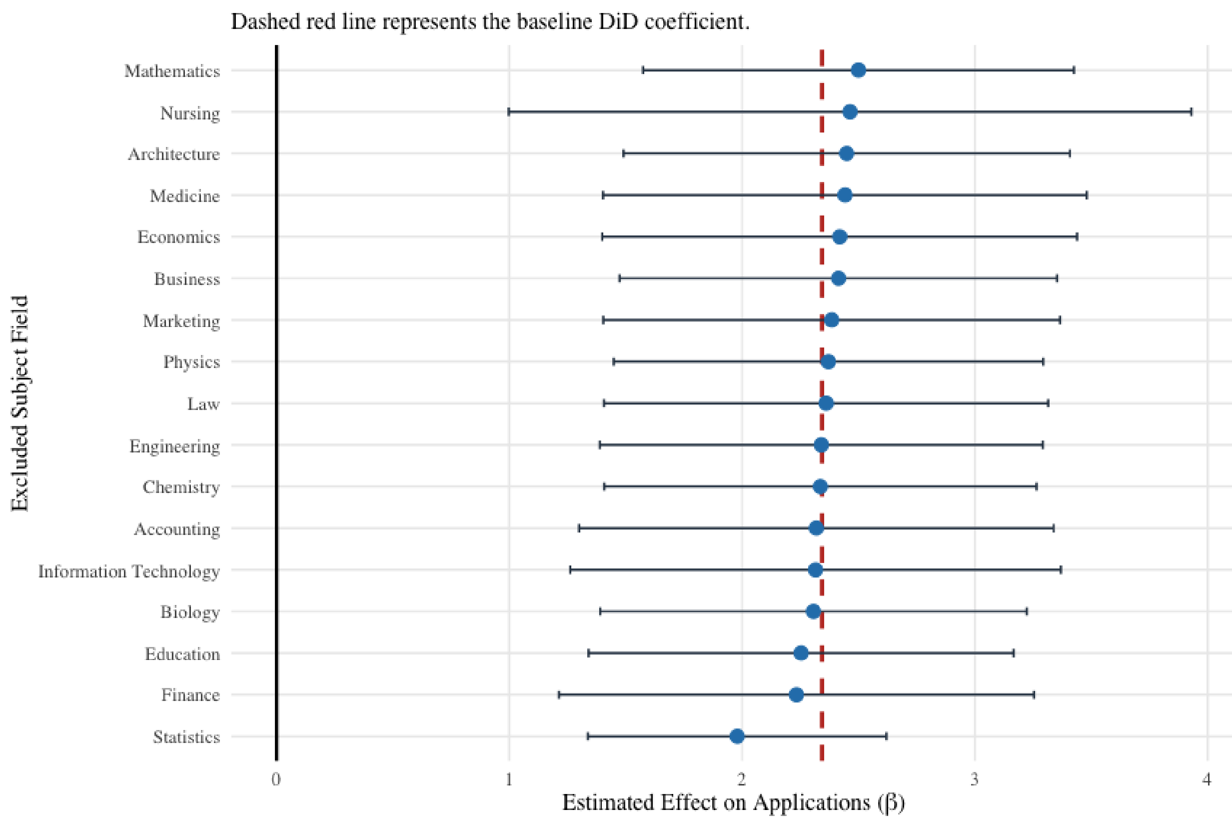

Notes: This figure visualises the stability of the baseline difference-in-differences estimate. Each point represents the interaction coefficient obtained from re-estimating the baseline model while iteratively excluding the specific subject field listed on the y-axis. The vertical dashed line denotes the baseline full-sample estimate. Error bars represent 95% confidence intervals with standard errors clustered at the subject level.

C Heterogeneity Tests

Table 10: Heterogeneity of Treatment Effect by Ethnicity

	Dependent variable:				
	Log Applications				
	Asian	Black	Mixed	Other	White
	(1)	(2)	(3)	(4)	(5)
AI Exposure x Post	1.581**	2.408***	2.598***	2.139**	2.595***
	(0.700)	(0.751)	(0.510)	(0.920)	(0.403)
Field FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	119	119	119	119	119
R2	0.9925	0.9940	0.9938	0.9906	0.9956
Adjusted R2	0.9907	0.9926	0.9923	0.9884	0.9946
Note:					*p<0.1; **p<0.05; ***p<0.01 Clustered (Field) standard errors in parentheses

Notes: The table reports the estimated effect of AI exposure on log applications, separately by ethnic group. Each column (1)–(5) reports results from an independent difference-in-differences regression restricted to applicants of the indicated ethnicity (Asian, Black, Mixed, Other, White). *AI Exposure x Post* is the interaction between the field-level AI exposure index and a post-2022 indicator. All specifications include subject and year fixed effects, with standard errors clustered at the subject level.

Table 11: Effect by University Prestige (Russell vs. Non-Russell)

	Dependent variable:	
	Log Applications	
	Russell Group	Non-Russell Group
	(1)	(2)
AI Exposure x Post	1.473**	−0.003
	(0.648)	(0.827)
Field FE	*Yes*	*Yes*
Year FE	*Yes*	*Yes*
Observations	84	84
R2	0.9893	0.9915
Adjusted R2	0.9863	0.9891
Note:		*p<0.1; **p<0.05; ***p<0.01
		Clustered (Field) standard errors in parentheses

Notes: The table reports the estimated effect of AI exposure on log applications, separately by university prestige. Column (1) restricts the sample to applicants accepted to Russell Group institutions (the UK's 24 leading research universities), while Column (2) restricts to applicants accepted to all other institutions. Because provider-level UCAS data only reports acceptances rather than applications, the dependent variable is the log of accepted applications. *AI Exposure x Post* is the interaction between the field-level AI exposure index and a post-2022 indicator. Both specifications include subject and year fixed effects, with standard errors clustered at the subject level.

Table 12: Effect by Nation (England, Wales, Scotland, N. Ireland)

Nation	Coefficient	Std. Error	CI Lower	CI Upper
England	2.342	0.618	1.132	3.553
Scotland	1.837	0.612	0.638	3.037
Wales	2.039	1.094	−0.105	4.184
Northern Ireland	1.388	0.928	−0.430	3.207

Notes: The table reports the estimated effect of AI exposure on log applications, separately by nation of domicile. Each row reports the coefficient, standard error, and 95% confidence interval from an independent difference-in-differences regression restricted to applicants domiciled in the indicated nation (England, Scotland, Wales, or Northern Ireland). All specifications include subject and year fixed effects, with standard errors clustered at the subject level.

Table 13: Effect by Gender (Male vs. Female)

	Dependent variable:	
	Log Applications	
	Men	Women
	(1)	(2)
AI Exposure x Post	2.147**	2.252***
	(0.835)	(0.642)
Year FE	Yes	Yes
Field FE	Yes	Yes
Observations	119	119
R2	0.9923	0.9946
Adjusted R2	0.9905	0.9933
Note:	*p<0.1; **p<0.05; ***p<0.01	

Notes: The table reports the estimated effect of AI exposure on log applications, separately by gender. Column (1) restricts the sample to male applicants, while Column (2) restricts to female applicants. *AI Exposure x Post* is the interaction between the field-level AI exposure index and a post-2022 indicator. Both specifications include subject and year fixed effects, with standard errors clustered at the subject level.