Swiss Leading House

Economics of Education • Firm Behaviour • Training Policies

Working Paper No. 255

Choice Architecture in Occupational Choices

Madison Dell, Enzo Brox, Patricia Palffy, Claudio Schilter and Uschi Backes-Gellner

Universität Zürich
IBW – Institut für Betriebswirtschaftslehre

UNIVERSITÄT BERN

Working Paper No. 255

Choice Architecture in Occupational Choices

Madison Dell, Enzo Brox, Patricia Palffy,
Claudio Schilter and Uschi Backes-Gellner

April 2026

Only cite with permission of the authors.

Please cite as:
"Choice Architecture in Occupational Choices." Swiss Leading House "Economics of Education" Working Paper No. 255, 2026. By Madison Dell, Enzo Brox, Patricia Palffy, Claudio Schilter and Uschi Backes-Gellner.

The Swiss Leading House on Economics of Education, Firm Behavior and Training Policies is a Research Program of the Swiss State Secretariat for Education, Research, and Innovation (SERI).

www.economics-of-education.ch

Choice Architecture in Occupational Choices

Madison Dell[1,3], Enzo Brox[2,3], Patricia Palffy[1,3], Claudio Schilter[2,3], Uschi Backes-Gellner[1,3]

[1]University of Zurich
[2]University of Bern
[3]Swiss Leading House VPET-ECON

April 17, 2026

Work in progress. Please only cite with permission of the authors.

Abstract: We study how choice architecture in online platforms shapes high-stakes occupational choices through two behavioral mechanisms: motivated reasoning and cognitive load. Using detailed process data from a large online job board and exploiting a quasi-experimental setting, we leverage two sources of exogenous variation in the presentation of occupation recommendations. First, we use random variation in the rank order of equally well-matched occupations to study the effects of motivated reasoning. Our results show that rank order strongly increases the level of users' engagement on the platform and, consequently, the number of occupations to which they apply. Second, we exploit a redesign that transformed the occupation recommendations from a static, text-heavy list into an interactive and visually enriched presentation. The redesign was neither announced nor anticipated, which allows for causal interpretation. We find that this small redesign significantly increases the number of occupations to which users apply, supporting our hypothesis that it reduces cognitive load, leading to increased use of a watch list that keeps more occupations in jobseekers' memory. Our findings provide large-scale field evidence showing that even small changes in platform design significantly and strongly shape consequential career choices.

JEL: D91, J24, D83
Keywords: Occupational choice, Choice architecture, Recommender systems, Motivated reasoning, Cognitive load

Acknowledgements: We thank the Swiss State Secretariat for Education, Research and Innovation for financial support through its Leading House on the Economics of Education. We thank Yousty for their partnership and for providing access to the data used in this study. We are grateful to Daniel Goller for valuable methodological guidance, Natalie Reid for helpful writing consultation, and Daniela Weigl for excellent research assistance. We also thank the members of the Leading House Advisory Board, Christopher Stanton, Eric Bettinger, Michael Thaler, Simon Janssen, and Jesper Eriksen for insightful comments and suggestions. In addition, we benefited from feedback received from participants at the University of Zurich Doctoral Seminar in Personnel Economics and Economics of Education, 2025 Netzwerktreffen, 2025 Leading House on Tour Workshop at the London School of Economics, 2025 VET4GROWTH Workshop at the University of Bergamo, 2025 Leading House administrative data course taught by Aline Bütikofer, 2026 University of Zurich PhD Discussion Workshop in Personnel Economics and Economics of Education, 2026 COPE Conference, and 2026 LEER Conference.

1. Introduction

Both how information is presented and how choices are structured influence the decisions that individuals make. The details of the decision-making context, known as choice architecture, include design features such as the sequence in which options appear and the format in which information is presented (Thaler & Sunstein, 2009). Even seemingly minor design features can greatly shape behavior by emphasizing options or simplifying complex information (Mertens et al., 2022). Choice architecture can expand or restrict choices, affecting how individuals integrate new information into their decisions. In online environments, these features are often implemented by platforms that actively structure how users search, evaluate, and select among options.

Two important behavioral mechanisms, biases and cognitive constraints, explain how choice architecture affects decisions. Biases, which are systematic deviations from rational information processing, often reflect identity, preferences, or beliefs. One important bias that influences how individuals interpret information is motivated reasoning, which leads individuals to place greater weight on information aligning with their beliefs or social identity and to discount conflicting information (Thaler, 2024). Motivated reasoning thus reinforces preexisting views rather than prompting individuals to update their beliefs with new information.

In contrast, cognitive constraints entail the limits on an individual's ability to process large amounts of information at one time. One important cognitive constraint is cognitive load (Sweller, 1988). According to cognitive load theory, the capacity of individuals' working memory is limited, so that excessive or poorly structured information can overwhelm attention and reduce learning (Sweller, 1988, 2011). In complex decision environments, reductions in cognitive load may operate primarily by expanding the variety of options that individuals consider rather than restricting the options to one top choice.

On both mechanisms, empirical evidence so far mainly comes from laboratory experiments (e.g., Thaler et al., 2024). While these studies provide strong internal validity, they offer limited insight into how these mechanisms operate in the complex, socially embedded environments in which many consequential decisions occur. High-stakes, real-world decisions such as occupational choice often involve stronger identity concerns and more complex information environments. Thus, little is known about whether laboratory experiments generalize to such decisions.

Our study investigates how motivated reasoning and cognitive load shape the impact of choice architecture on occupational choices. To provide causal evidence, we require a setting in which individuals make consequential decisions in an information-rich environment, free from prior labor market experience that could constrain their choices. Switzerland's vocational education and training (VET) system offers such a setting for several reasons. Approximately two-thirds of Swiss adolescents choose an apprenticeship after compulsory schooling ends at age 15 or 16. Since this is the first occupational choice an adolescent can make, this choice is unconfounded by prior work experience. Given that these three- to four-year apprenticeships set adolescents on a career path, this decision has long-term consequences. Moreover, with about 240 occupations to choose from, adolescents face a complex search problem and must confront a large amount of occupation information. These features make the Swiss VET system an ideal context for studying behavioral mechanisms in real-world occupational choices.

We leverage data from Yousty, Switzerland's largest private online job board for apprenticeship positions, which posts roughly 90% of all apprenticeship positions (Palffy et al., 2024). Given this near-complete market coverage, Yousty is widely used by Swiss adolescents entering VET. The platform recommends occupations based on users' stated interests and displays a match score indicating how closely each occupation aligns with those interests. Users can then view, save, and apply to occupations directly on the Yousty platform, and we use the resulting process data for our analyses.

To study the causal effects of motivated reasoning and cognitive load in occupational choices, we exploit two sources of exogenous variation. First, we exploit a randomly assigned rank order among recommended occupations with identical match scores. Specifically, if the assessment of alignment between an individual's occupational interests and the requirement profiles of particular occupations yields two occupations with the exact same match score, one occupation is randomly shown first and the other second in the online presentation of the individual's recommended occupations. Second, we exploit a design change in the presentation of recommended occupations that was unanticipated in the field and transformed the recommendation results from a static list with text-heavy occupation descriptions to an interactive and visually enriched (Tinder-like) presentation of occupation information. The exogenous variation of these design features enables us to causally identify how order and presentation influence Yousty users' occupational choices.

Regarding the first design feature, we find that rank order plays a substantial role in Yousty users' occupational choices. Even when occupations have an identical match score, users are more likely to apply to apprenticeships in higher-ranked occupations than in lower-ranked occupations and are significantly more likely to click on and save those higher-ranked occupations. The influence of rank varies systematically with occupation characteristics, consistent with users interpreting ambiguous rank information differently across contexts. Specifically, we find that rank effects are stronger for high-paying occupations and for occupations that align with gendered occupational choices linked to social identity. These patterns are consistent with motivated reasoning (Thaler, 2024): users appear to place greater weight on rankings that reinforce preferences toward which they are already inclined.

The second design feature, an interactive and visually enriched presentation of occupation recommendations, significantly increases the number of occupations to which users apply. The increase in the number of occupations to which users apply coincides with a large and precisely estimated increase in their watch list usage. After the redesign, users are substantially more likely to save at least one occupation and save more occupations on average. In line with cognitive load theory, we argue that the interactive and visually enriched redesign reduces cognitive load at any given point in time during the search process and thereby increases the use of the watch list, which keeps more occupations in memory for later applications.

Our paper makes three contributions to three different literatures. First, it contributes to the behavioral economics literature by providing large-scale field evidence on how choice architecture shapes occupational choice in a high-stakes, real-world setting. Whereas much of the existing evidence on motivated reasoning and cognitive load comes from controlled laboratory environments (e.g., Bordalo et al., 2016; Taber & Lodge, 2006; Thaler, 2024), our setting allows us to observe how these mechanisms operate when individuals face complex information, strong identity considerations, and substantial long-term consequences. Occupational choice is an especially suitable domain for studying behavioral mechanisms because it is both consequential and cognitively demanding. Adolescents must navigate many potential occupations, process heterogeneous information, and weigh personal interests, identity-relevant cues, and social norms (Palffy et al., 2023). Moreover, the long-term implications of initial career choices amplify the importance of understanding how biases and cognitive constraints influence decision-making. By moving beyond abstract decision tasks and into a domain where mistakes carry real costs, our study strengthens the external validity of behavioral economic theories about decision-making.

Second, our paper contributes to the labor economics literature by investigating how occupation recommender tools influence the decisions of jobseekers, specifically adolescents making their first labor market choice. While prior research has shown that recommender tools can expand job search (e.g., Belot et al., 2019; Oswald-Egg, 2025) and improve labor market outcomes such as employment rates and earnings (Altmann et al., 2022; Bächli et al., 2025; Belot et al., 2025a; Belot et al., 2025b), existing work has primarily evaluated their effectiveness rather than the mechanisms through which they operate. Yet recommender tools inherently rely on choice architecture: they rank options, filter information, and structure complex decision environments. Our study shows how two design features interact with motivated reasoning and cognitive load to shape occupational choices. By exploiting exogenous variation in ranking and in the visual presentation of recommendations, we identify how individuals respond not only to the content of recommendations but to the way that content is structured.

Third, our paper contributes to the platform design literature by showing how choice architecture in online environments influences users' choice sets and decisions. Platforms increasingly act as intermediaries that structure choices through ranking algorithms, information filtering, and interface design, shaping how users search for and evaluate options (e.g., Dinerstein et al., 2018; Ursu, 2018). A growing strand of literature highlights how specific design features, such as interactive decision aids, affect consumer behavior and market outcomes in online environments (e.g., Häubl & Trifts, 2000; Morath and Münster, 2018). Still, little is known about how these features interact with users' behavioral tendencies in high-stakes contexts. Our findings demonstrate that design elements such as the ordering of recommended options or the shift from text-heavy to visually structured information can alter job search and application behavior. By identifying the behavioral mechanisms through which these features operate, our study emphasizes the importance of considering biases and cognitive constraints when designing recommendation interfaces. More generally, our results offer guidance for platforms seeking to promote welfare-enhancing decisions: design features that enhance salience (such as rank ordering) can unintentionally reinforce identity-driven biases, whereas features that lower cognitive demands can broaden search and improve match quality. These insights extend beyond occupational choice to other online matching markets where platform design fundamentally shapes how individuals navigate complex decisions.

The paper proceeds as follows. Section 2 describes our data and provides the necessary institutional details of our causal identification strategy. Section 3 presents and interprets our main results. Section 4 discusses our contributions, gives policy recommendations and implications for future research, and concludes.

2. Data and Methods

2.1 Data

To empirically investigate how motivated reasoning and cognitive load shape the impact of choice architecture on occupational choices, we require a setting in which both sequencing and information presentation vary exogenously and can be directly observed. Yousty's online occupation recommender tool provides precisely this setting. It records detailed process data on how users interact with a search, matching, and recommender tool during their apprenticeship search, including the options they encounter and the order and format in which these options appear. Using these data, we quantify the impact of two design features of Yousty's occupation recommender tool. After completing a 33-question test, each user receives 20 occupation recommendations that best match their stated interests.[1] The dataset covers 246,869 users who completed a test between 2019 and 2024 and includes the match score and rank for each recommended occupation.

Our empirical strategy exploits two sources of plausibly exogenous variation embedded in the recommender tool. First, the assignment of a rank order of recommended occupations with an identical match score provides random variation in rank. When multiple occupations share identical match scores, the platform assigns rank based on a random occupation identifier.

[1] Appendix Figure A1 shows how the test questions look. The test did not change during the redesign.

Because the rank determines the sequence in which occupations appear but contains no additional information about match quality, variation in rank is effectively random among equally well-matched occupations. This feature allows us to identify the causal effect of rank positions within the recommendation list.

Second, we exploit a major change in how recommended occupations are presented. This change was not anticipated by users and was implemented on June 1, 2023. Prior to this date, recommended occupations were displayed in a static, text-heavy list showing all 20 occupations simultaneously.[2] After the redesign, the recommendations shifted to an interactive and visually enriched presentation of occupation information: Yousty users view one occupation at a time, accompanied by concise text, images, and videos.[3] To view the next occupation, users must either press the heart symbol to save the occupation to their watch list (i.e., "favorites" to revisit later)[4] or swipe the occupation away by pressing the arrow if it does not interest them. This change reduced the amount of information presented at once, lowering cognitive load while enhancing engagement by requiring users to make a decision about each occupation. Because the change was not known to the users and the implementation date was unexpected, we can interpret the design change as exogenous to users' behavior. This enables causal inference regarding the effects of information presentation, as the changes can be considered quasi-random.

In our empirical part, we focus on one of the most consequential decisions an adolescent can make before entering the labor market, their applications to apprenticeship positions. To explore how the redesign affected occupational choices, we construct an outcome measuring

---

[2] Appendix Figures A2 and A3 show how the occupation recommendations looked before the redesign.
[3] Appendix Figure A4 shows how the occupation recommendations look since the redesign. To see the occupation recommendations in the redesigned format, anyone can take the occupation interest test (called the Berufs-Finder), which is available here (in German, French, or Italian): https://www.yousty.ch/de-CH/berufswahl
[4] Appendix Figure A5 shows how the watch list looks since the redesign.

the number of occupations to which users submit apprenticeship applications. Importantly, this is measured as the number of distinct occupations (not the number of applications) because we are interested in whether the redesign induced users to apply to apprenticeships in a broader set of occupations or restricted their choices to just the first adequate option that comes to mind.

In a second step, we examine an important mechanism that cognitive load theory helps explain: the role of saving occupations to a watch list. We hypothesize that the interactive and visually enriched redesign reduces cognitive load and increases use of the watch list, which keeps more occupations in memory for later applications. Although the watch list feature existed before the redesign, it was infrequently used. We hypothesize that information overload led Yousty users to restrict their choices to a bare minimum. The redesigned recommendations help combat overload by presenting only one occupation at a time and structuring the choice to be as simple a decision as possible: users either save the occupation if interested or swipe to see the next occupation if not. To study this mechanism, we estimate whether the redesign increases occupations saved to the watch list. We distinguish between effects on the extensive margin (i.e., whether a user saves any occupations to the watch list, measured by a binary indicator) and the intensive margin (i.e., how many occupations are saved, measured by the number of occupations saved, conditional on saving at least one occupation).

Table 1 reports summary statistics for our outcome variable (i.e., the number of occupations to which users submit apprenticeship applications) and proposed mechanism motivated by cognitive load theory (i.e., the number of occupations saved to the watch list). The application outcome is highly zero-inflated, since only 7.1% of Yousty users apply to any apprenticeship in a given time period. The mean number of occupations to which Yousty users apply is 0.095 occupations per user, but it is highly dispersed (SD = 0.401) with a long right tail up to 12 occupations. Saving occupations to the watch list is a much more common phenomenon than applying to a real apprenticeship position in an occupation, as 18.9% of

Yousty users save one or more occupations to the watch list. The mean number of occupations saved to the watch list is 4.381 occupations per user, but it is highly dispersed (SD = 3.416) and ranges from 1 to 20 different occupations (conditional on a user saving any occupations to the watch list).

Table 1: Summary Statistics

Variable	N	Mean	SD	Min	Max
Outcome Variable: Occupations to which Yousty Users Apply					
Number of Occupations in Applications	246,869	0.095	0.401	0	12
Mechanism Variable: Occupations Saved to Watch List					
Saved Any Occupations to Watch List	246,869	0.189	0.392	0	1
Number of Occupations on Watch List	46,680	4.381	3.416	1	20

Notes: N = observation count. SD = standard deviation. Min = minimum value. Max = maximum value. Full sample is 246,869. "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list.

2.2 Methods

Estimating the Rank Effect

To identify the impact of rank, the central challenge is that the displayed rank of an occupation is mechanically related to its match score, which reflects the platform's assessment of preference alignment. Naive comparisons across ranks therefore conflate the effect of being displayed at a higher position with differences in underlying fit. Our empirical strategy isolates the role of rank by conditioning flexibly on the match score and by exploiting within-person variation across the top 20 recommended occupations.

We estimate variants of the following baseline specification:

$$Y_{ij} = \beta \cdot Rank_{ij} + f(Score_{ij}) + \alpha_i + \gamma_j + \varepsilon_{ij}, \tag{1}$$

where Y_{ij} denotes one of the three engagement outcomes. The function $f(\cdot)$ controls flexibly for the match score using score-bin fixed effects. Individual fixed effects α_i absorb all person-specific characteristics that affect overall platform engagement, such as motivation, ability, or background. Occupation fixed effects γ_j capture baseline differences in popularity or application intensity across occupations. Identification in equation (1) comes from within-person comparisons across occupations that differ in rank but have similar match scores.

To strengthen credibility, we implement three complementary robustness exercises. First, we exploit a salient feature of the recommendation algorithm, which is that multiple occupations often receive exactly the same match score for a given user. In our data, score ties frequently involve more than two occupations and can include multiple occupations within a user's top 20 list. We exploit this structure by estimating an alternative specification that restricts identification to comparisons *within score ties*:

$$Y_{ij} = \beta \cdot Rank_{ij} + \alpha_{i,s} + \gamma_j + \varepsilon_{ij}, \tag{2}$$

where $\alpha_{i,s}$ denotes individual-score fixed effects, defined at the level of individual i and exact score s. This specification absorbs all variation across different scores and identifies β solely from differences in rank among occupations that have the same match score for the same user.

Second, we estimate nonparametric specifications that replace the linear rank term with rank indicators, yielding a transparent rank gradient that illustrates how engagement decays with lower placement in the recommendation list.

Third, we further refine the identification strategy by exploiting variation in the ordering of occupations *within score ties*. Specifically, we replace the overall rank measure with the order in which occupations appear among those that share the same match score for a given user. This within-tie ordering reflects the sequence in which equally scored occupations are displayed. To isolate this variation, we include fixed effects for the overall rank position, which absorb differences in engagement that arise from being placed higher or lower in the recommendation list. Under this specification, identification comes from differences in the relative ordering of occupations within score ties, holding both the match score and the absolute position in the list constant. This approach isolates variation in display order that is orthogonal to both underlying match quality and general position effects.

Estimating the Redesign Effect

To estimate the causal impact of the 2023 platform redesign on Yousty users' occupational choices, we employ a difference-in-discontinuity (DD-RD) design (e.g., Butts, 2023; Goller et al., 2025; Grembi et al., 2016). The intuition of this approach is to isolate the effect of the redesign from seasonal patterns in apprenticeship search and application behavior by comparing the discontinuity in outcomes at the redesign threshold (a sharp cutoff on June 1, 2023) with the corresponding discontinuities observed at the same date in non-treated years. The running variable is defined as days relative to June 1st in each calendar year,[5] which allows for flexible time trends on either side of the cutoff while differencing out seasonal changes that are common across years.

Identification relies on the assumption that, absent the redesign, outcomes would evolve smoothly across the June 1st threshold. This assumption is plausible in our setting because the redesign date was unknown to users and never announced in advance, preventing users from manipulating test timing around the cutoff. Using the following equation, we estimate local linear regressions on either side of the cutoff, restricting the sample to observations within a symmetric bandwidth around the redesign date:

$$Y_i = \beta_0 + \beta_1(D_i) + \beta_2(P_i) + \beta_3(D_iP_i) + \beta_4(T_i) + \beta_5(D_iT_i) + \beta_6(P_iT_i) + \beta_7(D_iP_iT_i) + \varepsilon_i$$

In this model, Y_i represents the outcome of interest for individual i, D_i is the running variable measuring the number of days between an individual's test date and the redesign date (with D_i = 0 on June 1st), P_i is an indicator equal to one if the test date falls on or after June 1st in a given year, and T_i is an indicator for the treated year (2023). The coefficient β_6 captures

[5] The running variable represents the first action users take in the application process, which is taking a test on occupational interests that then provides them with occupation recommendations. The redesigned recommendations were shown to users beginning on June 1, 2023. To ensure clean identification, we exclude users who completed their test during May, the month that the recommendations transitioned from the old version to the redesigned version. Excluding May avoids potential contamination from partial exposure during the transition month.

the main parameter of interest: the causal effect of the redesign, interpreted as the difference in the discontinuity at the June 1st cutoff in the treated year relative to the corresponding discontinuities in other years.

3. Results

3.1 Random Assignment of Ranks: The Role of Motivated Reasoning

Does Rank Affect Outcome Measures?

We begin by documenting the baseline relationship between rank and user engagement with occupations. Table 2 presents our main results, estimating equation (1) with individual fixed effects, occupation fixed effects, and flexible controls for the match score. Across all three outcome measures, we find a strong and statistically significant effect of rank. Occupations that appear higher in the recommendation list are substantially more likely to receive clicks, saves, and applications than lower-ranked alternatives, even when they have similar match scores.

The magnitude of the estimates is economically meaningful. Moving an occupation down by one rank position reduces overall activity by approximately 0.4-0.5 percentage points. This corresponds to approximately 2% relative to the sample mean. This 2% effect is consistent across our measures, although it is measured relative to a lower mean in the case of applications. These patterns are remarkably stable across specifications. In columns 2, 4, and 6, instead of controlling for the match score using score-bin fixed effects, we exploit within-match score variation by including fixed effects for the interaction of the individual and the score FE. Reassuringly, the estimated rank effects remain highly similar in magnitude and statistical significance across all outcomes. Since rank contains no additional information once the match score is held fixed, these findings suggest that users respond directly to the ordering of recommendations rather than to differences in assessed fit.

Figure A6 in the Appendix provides a graphical illustration of the baseline rank effects by plotting predictive margins from the baseline specification for each of the three outcomes.

The figure shows a clear and monotonic decline in engagement as rank worsens across the top 20 recommendations. Higher-ranked occupations receive substantially more overall activity, are more likely to be saved to the watch list, and have a higher probability of receiving an application. The gradients are steepest at the top of the recommendation list, with engagement dropping sharply between the first few ranks and flattening out at lower positions.

Table 2: Plain Rank Effect

	Activity		Watch List		Apply	
	(1)	(2)	(3)	(4)	(5)	(6)
Rank	-0.0047***	-0.0049***	-0.0042***	-0.0042***	-0.0002***	-0.0002**
	(0.0002)	(0.0004)	(0.0002)	(0.0003)	(0.0000)	(0.0001)
Individual FE	✓	✓	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓	✓	✓
Score-Bin FE	✓		✓		✓	
Indiv. x Score FE		✓		✓		✓
N	652,220	600,841	652,220	600,841	652,220	600,841

Notes: Columns (1)-(2) use a count measure of activity, defined as the number of distinct actions taken with respect to an occupation: viewing any job advertisement of the occupation, viewing the occupation profile, saving the occupation to the watch list, applying to a trial apprenticeship, or applying to an apprenticeship. Columns (3)-(4) use a dummy for saving the occupation to the watch list. Columns (5)-(6) use a dummy for applying to an apprenticeship. Standard errors, clustered at the user level, are reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

In Table 3, we further probe the robustness of the findings by adopting an alternative definition of rank that focuses on the ordering of occupations *within individual-match score ties*. Specifically, we replace the overall rank measure with the order in which occupations appear among those that share the same match score for a given user, while controlling for fixed effects for the overall rank position. This approach isolates variation in display order that is orthogonal to both the score and the absolute position in the list. The results closely mirror those of Table 2.

Table 3: Effect of Order within Ties

	Activity		Watch List		Apply	
	(1)	(2)	(3)	(4)	(5)	(6)
Order	-0.0049***	-0.0069***	-0.0042***	-0.0061***	-0.0002**	-0.0003**
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Individual x Score FE	✓	✓	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓	✓	✓
Rank FE		✓		✓		✓
N	600,841	600,840	600,841	600,840	600,841	600,840

Notes: Columns (1)-(2) use a count measure of activity, defined as the number of distinct actions taken with respect to an occupation: viewing any job advertisement of the occupation, viewing the occupation profile, saving the occupation to the watch list, applying to a trial apprenticeship, or applying to an apprenticeship. Columns (3)-(4) use a dummy for saving the occupation to the watch list. Columns (5)-(6) use a dummy for applying to an apprenticeship. Standard errors, clustered at the user level, are reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Taken together, this section establishes a clear and robust empirical pattern: rank exerts a strong influence on users' engagement with occupations. The consistency of the estimates across multiple specifications and alternative rank definitions indicates that this relationship is not an artifact of score measurement or functional form assumptions but instead reflects a genuine behavioral response to how options are ordered on the platform.

Is Rank Used as Information?

The baseline analysis documents a robust relationship between rank and engagement with occupation recommendations. A natural question is whether this relationship reflects a purely mechanical salience effect, where higher-ranked occupations receive more attention or perceived fit simply because they are more prominent, or whether individuals use the ambiguous rank signal to form beliefs about occupational quality in ways that vary systematically across contexts.

To shed light on this distinction, we examine whether rank effects vary with observable indicators of occupational quality that are plausibly relevant for users' decisions, such as expected wages. If rank operates purely through salience or attention, its effect should be largely uniform across occupations and independent of such characteristics. In contrast, if users interpret rank as a signal of quality or suitability whose relevance varies systematically across contexts, rank effects should interact with these indicators in structured ways.

We implement this idea by estimating specifications of the form:

$$Y_{ij} = \beta_1 Rank_{ij} + \beta_2 Q_j + \beta_3 (Rank_{ij} \times Q_j) + \alpha_{i,s} + \gamma_j + \varepsilon_{ij}, \tag{3}$$

where Q_j denotes an occupation-level quality measure, such as the median wage associated with occupation j. As before, $\alpha_{i,s}$ are individual-score fixed effects, and γ_j are occupation fixed effects.

Under a pure salience mechanism, the interaction coefficient β_3 should be close to zero, indicating that rank effects are similar across occupations perceived as high- and low-quality. By contrast, a positive (negative) β_3 would indicate that rank matters more (less) for occupations with higher perceived quality, consistent with individuals using rank as an informational cue whose relevance varies across contexts.

Our interpretation hinges on systematic heterogeneity in how rank information is used. Motivated reasoning predicts that individuals selectively rely on ambiguous signals to justify ex ante more desirable choices. In this context, favoring occupations with higher predicted earnings may be psychologically more comfortable than favoring low-paying alternatives.[6] We therefore hypothesize that users place greater weight on rank information when evaluating occupations with higher predicted earnings.

Table 4 provides strong support for this hypothesis. Across all outcomes, the interaction between rank and an indicator for occupations in the top quartile of predicted earnings is negative and highly statistically significant, indicating that the rank gradient is substantially steeper for high-earning occupations. Quantitatively, the estimated rank effect at least doubles

[6] We use the same earnings data as Brenoe & Wasserman (2024), which provide average predicted earnings for a representative individual (i.e., a 30-year-old, childless, unmarried person living in non-rural Zurich surveyed in January 2019) for nearly all apprenticeship occupations.

for occupations with high predicted earnings relative to the baseline. This pattern suggests that users are more responsive to rank when the occupation is ex ante more attractive along a salient dimension, consistent with rank being interpreted as a quality signal rather than operating purely through mechanical salience. Table A1 in the Appendix shows that these findings are robust to allowing for non-linearities in the effect of rank. When we include both rank and rank squared, as well as their interactions with the high-earnings indicator, we continue to find pronounced heterogeneity.

Table 4: Heterogeneity in Rank Effect by Predicted Earnings

	Activity		Watch List		Apply	
	(1)	(2)	(3)	(4)	(5)	(6)
Rank	-0.0046***	-0.0033***	-0.0037***	-0.0034***	-0.0003***	0.0001
	(0.0004)	(0.0004)	(0.0003)	(0.0003)	(0.0001)	(0.0001)
Top Earnings Quartile	0.2230***	–	0.1284***	–	0.0347***	–
	(0.0050)	–	(0.0033)	–	(0.0012)	–
Rank x Top Earn. Q.	-0.0079***	-0.0073***	-0.0042***	-0.0041***	-0.0015***	-0.0013***
	(0.0004)	(0.0004)	(0.0002)	(0.0002)	(0.0001)	(0.0001)
Individual x Score FE	✓	✓	✓	✓	✓	✓
Occupation FE		✓		✓		✓
N	568,698	568,698	568,698	568,698	568,698	568,698

Notes: Columns (1)-(2) use a count measure of activity, defined as the number of distinct actions taken with respect to an occupation: viewing any job advertisement of the occupation, viewing the occupation profile, saving the occupation to the watch list, applying to a trial apprenticeship, or applying to an apprenticeship. Columns (3)-(4) use a dummy for saving the occupation to the watch list. Columns (5)-(6) use a dummy for applying to an apprenticeship. Standard errors, clustered at the user level, are reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We next turn to heterogeneity along social identity dimensions. If social norms or self-image concerns matter, users may be reluctant to express interest in occupations that conflict with prevailing gender norms. In such cases, motivated reasoning predicts that rank information should be less effective in justifying engagement with gender-atypical occupations. The results in Table 5 are consistent with this prediction. The positive triple interaction (between rank, gender, and gender typicality of the occupation) indicates that rank effects are muted when male users evaluate stereotypically female occupations and analogously for female users evaluating stereotypically male occupations.

Table 5: Heterogeneity in Rank Effect by Gender Typicality

	Activity		Watch List		Apply	
	(1)	(2)	(3)	(4)	(5)	(6)
Rank x Male x Female Occ	0.0044***	0.0046***	0.0036***	0.0037***	0.0003**	0.0003**
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Individual x Score FE	✓	✓	✓	✓	✓	✓
Occupation FE	✓		✓		✓	
Occupation x Gender FE		✓		✓		✓
N	600,841	600,840	600,841	600,840	600,841	600,840

Notes: Columns (1)-(2) use a count measure of activity, defined as the number of distinct actions taken with respect to an occupation: viewing any job advertisement of the occupation, viewing the occupation profile, saving the occupation to the watch list, applying to a trial apprenticeship, or applying to an apprenticeship. Columns (3)-(4) use a dummy for saving the occupation to the watch list. Columns (5)-(6) use a dummy for applying to an apprenticeship. All specifications control for the main effects of rank, gender, and gender typicality of the occupation, as well as all pairwise interactions among these variables. Standard errors, clustered at the user level, are reported in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

The interpretation of the result in Table 5 assumes that users have, on average, gender-typical preferences. In Table 6, we challenge this assumption. Using responses to the 33 questions in the occupation interest test, we classify individuals according to whether their preferences align with gender-typical patterns.[7] Table 6 shows that the attenuation of rank effects for gender-atypical occupations is driven almost entirely by users with gender-congruent preferences. For individuals whose preferences are aligned with gender norms, rank information is substantially less influential when evaluating occupations that conflict with those norms. By contrast, for users with gender-non-congruent preferences, rank effects are markedly weaker. This pattern suggests that rank is selectively used to justify choices that are already psychologically or socially acceptable.

[7] We estimate a logistic regression predicting gender from all 33 answers and use the resulting predicted probability of being female as a continuous measure of gender typicality. Individuals are classified as having gender-congruent preferences if this predicted probability aligns with their observed gender and as gender-non-congruent otherwise.

Table 6: Heterogeneity in Rank Effect by Gender Typicality and Preference Congruence

	Activity		Watch List		Apply	
	(1)	(2)	(3)	(4)	(5)	(6)
Rank x Male x Female Occ	0.0074***	0.0027***	0.0056***	0.0026***	0.0006***	0.0000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)
Individual x Score FE	✓	✓	✓	✓	✓	✓
Occupation x Gender FE	✓	✓	✓	✓	✓	✓
Sample	congruent	non-congruent	congruent	non-congruent	congruent	non-congruent
N	318,310	282,524	318,310	282,524	318,310	282,524

Notes: The congruent (non-congruent) sample consists of individuals whose preferences are (not) congruent with gender-typical preferences. Columns (1)-(2) use a count measure of activity, defined as the number of distinct actions taken with respect to an occupation: viewing any job advertisement of the occupation, viewing the occupation profile, saving the occupation to the watch list, applying to a trial apprenticeship, or applying to an apprenticeship. Columns (3)-(4) use a dummy for saving the occupation to the watch list. Columns (5)-(6) use a dummy for applying to an apprenticeship. All specifications control for the main effects of rank, gender, and gender typicality of the occupation, as well as all pairwise interactions among these variables. Standard errors, clustered at the user level, are reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

3.2 Interactive and Visually Enriched Presentation: The Role of Cognitive Load

Table 7 presents difference-in-discontinuity estimates of the causal effect of the 2023 redesign on apprenticeship applications. It shows a significant jump in the number of occupations after the June 1st redesign. Across different bandwidths, the discontinuity at the June 1st cutoff is significantly larger in the treated year, with estimates ranging from 0.032 to 0.052 additional occupations. While modest in absolute terms, these effects are sizable relative to the average discontinuity in control years (0.085 occupations). Figure 1 provides descriptive evidence on the timing of this change, showing a modest increase around the June 1st cutoff in 2023 that is not present in other years. Overall, the results indicate that the redesign led users to apply to a broader set of occupations.[8]

[8] In the Appendix, we present several robustness checks and supplemental analyses that yield qualitatively similar conclusions to our main results: using alternative years in the control group (Tables A2-A3); including the transition month of May (Tables A4-A6); running a traditional sharp RD (using the Stata command *rdrobust*; Calonico et al., 2014) that does not account for seasonality (Table A7); and producing RD plots (using the Stata command *rdplot*; Calonico et al., 2014) for visual inspection (Figure A7).

Table 7: DD-RD Estimates for Occupations to which Yousty Users Apply

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Outcome Variable: Occupations to which Yousty Users Apply				
Number of Occupations in Applications	0.032**	0.033**	0.036**	0.052**
	(0.012)	(0.014)	(0.018)	(0.024)

Notes: Control group years include 2020-2022 and 2024. Month of May is excluded in all years. Bandwidth around June 1st indicated in each column heading. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Figure 1: Mean Outcomes by Week (Since June 1st), Number of Occupations in Applications

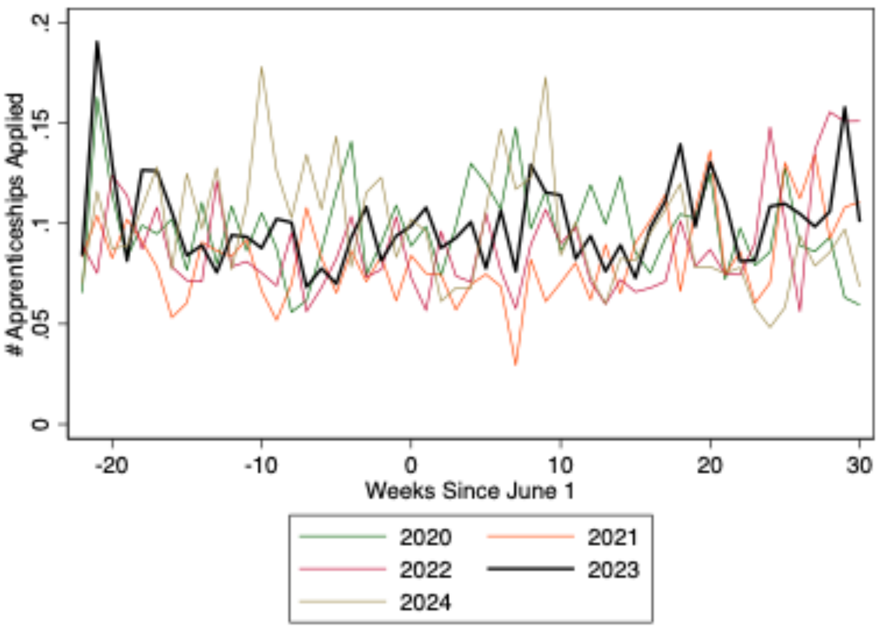

Table 8 presents difference-in-discontinuity estimates for the mechanism variable, occupations saved to the watch list. It shows a significant increase in the number of occupations after the June 1st redesign. Across different bandwidths, the discontinuity at the June 1st cutoff is substantially larger in the treated year in comparison other years, with estimates ranging from 0.345 to 0.444 for the probability of saving at least one occupation. We also observe large effects at the intensive margin: among users who save any occupations, the number of occupations saved increases by approximately 1.7 to 2.0 occupations. These effects are stable across bandwidth choices and are visually apparent in Figures 2 and 3, where we plot mean outcomes in each week (relative to June 1st) by year. These figures show a sharp increase in saving behavior at the June 1st cutoff in 2023 that is absent in other years, followed by a persistently higher level of saving in the post-redesign period, including throughout 2024. These results are consistent with our hypothesis that the redesigned recommendations reduce cognitive load and increase use of the watch list, which keeps more occupations in memory for later applications.[9]

Table 8: DD-RD Estimates for Occupations Saved to Watch List

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Mechanism Variable: Occupations Saved to Watch List				
Saved Any Occupations to Watch List	0.444***	0.433***	0.387***	0.345***
	(0.013)	(0.014)	(0.018)	(0.025)
Number of Occupations on Watch List	1.737***	1.818***	2.015***	1.704***
	(0.249)	(0.281)	(0.351)	(0.468)

Notes: Control group years include 2020-2022 and 2024. Month of May is excluded in all years. Bandwidth around June 1st indicated in each column heading. "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list. Standard errors in parentheses. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Figure 2: Mean Outcomes by Week (Since June 1st), Saved Any Occupations to Watch List

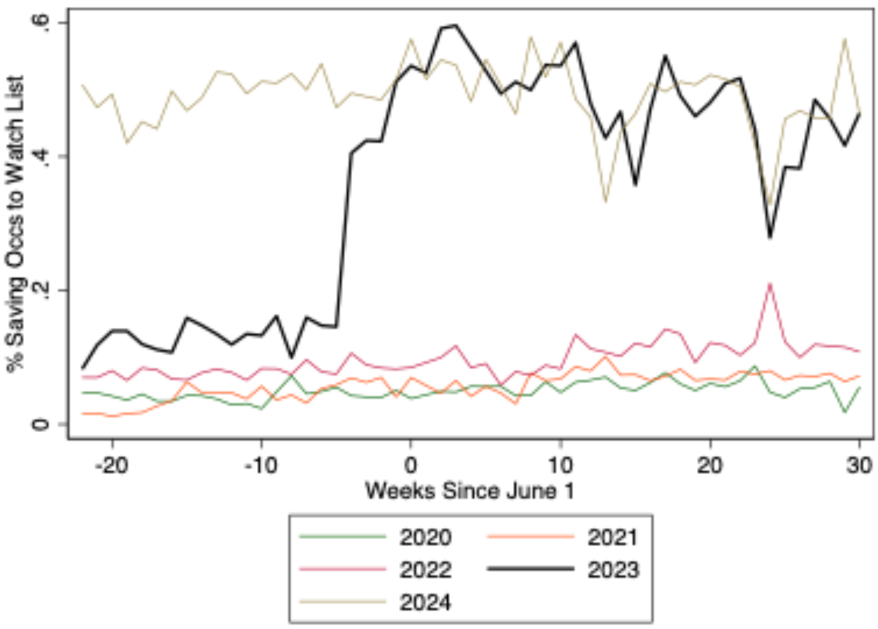

[9] In the Appendix, we present several robustness checks and supplemental analyses that yield qualitatively similar conclusions to our main results: using alternative years in the control group (Tables A2-A3); including the transition month of May (Tables A4-A6); running a traditional sharp RD (using the Stata command *rdrobust*; Calonico et al., 2014) that does not account for seasonality (Table A7); and producing RD plots (using the Stata command *rdplot*; Calonico et al., 2014) for visual inspection (Figures A8-A9).

Figure 3: Mean Outcomes by Week (Since June 1st), Number of Occupations on Watch List

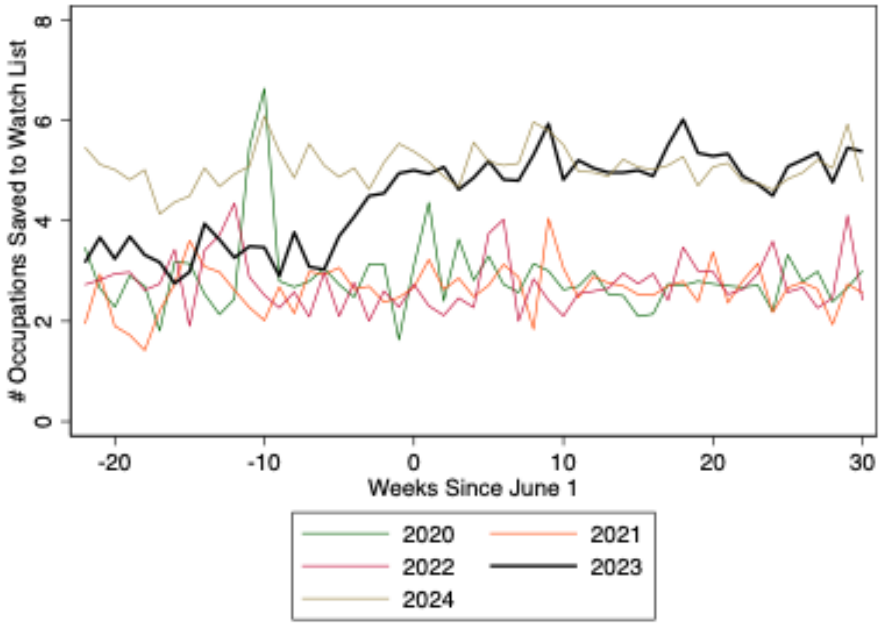

Note: "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list.

Taken together, the results indicate that the 2023 redesign led Yousty users to apply to a broader set of recommended occupations rather than restricting their options from the very beginning. The increase in occupations to which users applied can be explained by reduced cognitive load, which led to a large and persistent increase in saving occupations to a watch list. This suggests that the redesign primarily influenced occupational choices by encouraging users to save more occupations, keeping options top of mind to revisit later when applying. Overall, the findings are consistent with the redesign reducing cognitive load and translating increased saving behavior into a broader set of applications.

4. Discussion and Conclusion

We study how choice architecture shapes high-stakes occupational choices by isolating two behavioral mechanisms: motivated reasoning and cognitive load. Leveraging exogenous variation in the ranking and presentation of occupation recommendations on Switzerland's largest apprenticeship platform, we provide causal evidence on how small changes in design features influence Yousty users' occupational decision-making process. Our findings show that such small choices matter not only in the lab but also for consequential real-life decisions such as occupational choices in apprenticeship applications, which strongly affect long-term labor market outcomes.

We document a substantial effect of rank order on Yousty users' engagement and application behavior, even when occupations are equally well matched to users' stated interests. We then show that this rank effect varies systematically with occupation characteristics, suggesting that users interpret the inherently ambiguous ranking information differently across contexts. Rank effects are stronger for high-paying occupations and for occupations that align with gendered occupational norms. This pattern is consistent with motivated reasoning, whereby users place greater weight on ambiguous information when it reinforces pre-existing preferences.

Consistent with cognitive load theory, we find that the interactive and visually enriched presentation of occupation recommendations increased the number of occupations to which users submit apprenticeship applications. These application effects can be explained by a substantial increase in watch list usage, as Yousty users were significantly more likely to save occupations at all and saved a larger number of occupations on average. We argue that increased usage of the watch list helps Yousty users keep occupations top of mind to revisit later when applying to apprenticeships. These results indicate that reducing cognitive load through interface design can meaningfully affect downstream, high-stakes outcomes such as occupational choices.

Our findings also highlight broader implications for the design of online platforms that mediate high-stakes decisions. Even seemingly small and low-cost changes in how information is ordered or presented can have meaningful effects on behavior, shaping not only which options users consider but also how extensively they search. This underscores that platform design is not neutral: it actively structures decision environments and can amplify or mitigate behavioral biases. In our setting, ranking interacts with motivated reasoning in ways that may reinforce existing preferences, while reducing cognitive load can broaden search and promote more extensive exploration.

These insights are particularly relevant as online platforms increasingly guide decisions in domains such as education, labor markets, and finance. For platform designers, our results suggest that improving decisions quality requires careful attention to both the informational content of recommendations and the way this information is delivered. For policymakers, they highlight that subtle design features can have aggregate effects on allocation and inequality, especially when they interact with identity-driven preferences.

More broadly, we contribute real-world rather than laboratory evidence on behavioral mechanisms in consequential decisions, demonstrating that cognitive constraints and biases identified in laboratory settings also play a meaningful role in shaping real-world occupational choices. Taken together, these findings highlight the importance of carefully designing online choice environments to support more informed and effective decision-making in consequential domains.

References

Altmann, S., Glenny, A. M., Mahlstedt, R., & Sebald, A. (2022, December). *The direct and indirect effects of online job search advice.* (IZA Discussion Paper No. 15830). https://ssrn.com/abstract=4310239

Bächli, M., Lalive, R., & Pellizzari, M. (2025, February). *Helping jobseekers with recommendations based on skill profiles or past experience: Evidence from a randomized intervention.* (IZA Discussion Paper No. 17713). https://ssrn.com/abstract=5147072

Belot, M., Kircher, P., & Muller, P. (2019). Providing advice to jobseekers at low cost: An experimental study on online advice. *The Review of Economic Studies, 86*(4), 1411–1447. https://doi.org/10.1093/restud/rdy059

Belot, M., de Koning, B. K., Fourarge, D., Kircher, P., Muller, P., Phlippen, S. (2025a). *Advising job seekers in occupations with poor prospects: A field experiment.* (NBER Working Paper No. 33819). https://www.nber.org/papers/w33819

Belot, M., Kircher, P., & Muller, P. (2025b). Do the long-term unemployed benefit from automated occupational advice during online job search? *The Economic Journal*, ueaf041. https://doi.org/10.1093/ej/ueaf041

Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2016). Stereotypes. *The Quarterly Journal of Economics*, *131*(4), 1753-1794. https://doi.org/10.1093/qje/qjw029

Brenoe, A., & Wasserman, M. (2024). "Math skills, perceptions of fit, and occupational choice." AEA RCT Registry. May 31. https://doi.org/10.1257/rct.10215-2.2

Butts, K. (2023). Geographic difference-in-discontinuities. *Applied Economics Letters*, *30*(5), 615–619. https://doi.org/10.1080/13504851.2021.2005236

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust data-driven inference in the regression-discontinuity design. *The Stata Journal*, *14*(4), 909–946. https://doi.org/10.1177/1536867X1401400413

Dinerstein, M., Einav, L., Levin, J., & Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review, 108*(7), 1820-1859. https://doi.org/10.1257/aer.20171218

Goller, D., Gschwendt, C., & Wolter, S. C. (2025). This time it's different – Generative artificial intelligence and occupational choice. *Labour Economics*, 102746. https://doi.org/10.1016/j.labeco.2025.102746

Grembi, V., Nannicini, T., & Troiano, U. (2016). Do fiscal rules matter? *American Economic Journal: Applied Economics*, 1–30. http://dx.doi.org/10.1257/app.20150076

Häubl, G., & Trifts, V. (2000). Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science, 19*(1), 4-21. https://doi.org/10.1287/mksc.19.1.4.15178

Mertens, S., Herberz, M., Hahnel, U. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences, 119*(1), e2107346118. https://doi.org/10.1073/pnas.2107346118

Morath, F., & Münster, J. (2018). Online shopping and platform design with ex ante registration requirements. *Management Science, 64*(1), 360-380. https://doi.org/10.1287/mnsc.2016.2595

Oswald-Egg, M. E. (2025). *New information, new interests? Impact of an occupation finder on the occupational choices of dual vocational education and training students.* (CES Working Paper Series No. 23). https://doi.org/10.3929/ethz-b-000717755

Palffy, P., Lehnert, P., & Backes-Gellner, U. (2023). Social norms and gendered occupational choices of men and women: Time to turn the tide? *Industrial Relations, 62*(4), 380–410. https://doi.org/10.1111/irel.12332

Palffy, P., Lehnert, P., & Backes-Gellner, U. (2024). *What works for women does not work for men: A large field experiment on countering gendered occupational choices.* (Swiss Leading House "Economics of Education" Working Paper No. 207). https://ideas.repec.org/p/iso/educat/0207.html

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4

Sweller, J. (2011). Cognitive load theory. In J. P. Mestre & B. H. Ross (Eds.), *The psychology of learning and motivation: Cognition in education* (pp. 37–76). Elsevier. https://doi.org/10.1016/B978-0-12-387691-1.00002-8

Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, *50*(3), 755-769. https://doi.org/10.1111/j.1540-5907.2006.00214.x

Thaler, M. (2024). The fake news effect: Experimentally identifying motivated reasoning using trust in news. *American Economic Journal: Microeconomics, 16*(2), 1–38. https://doi.org/10.1257/mic.20220146

Thaler, R., & Sunstein, C. S. (2009). *Nudge: Improving decisions about health, wealth, and happiness.* Penguin Books.

Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science, 37*(4), 530-552. https://doi.org/10.1287/mksc.2017.1072

Appendix: Figures and Tables

Figure A1: Sample Question from Occupation Interest Test

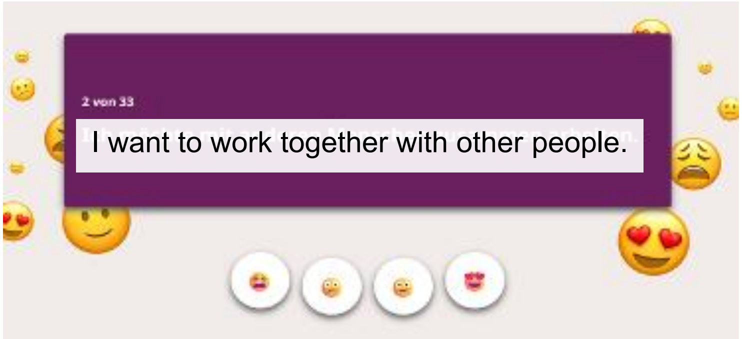

Figure A2: Sample Occupation Recommendations from Before the Redesign (Top 20)

Figure A3: Sample Occupation Recommendations from Before the Redesign (Top 3)

Figure A4: Sample Occupation Recommendations from After the Redesign

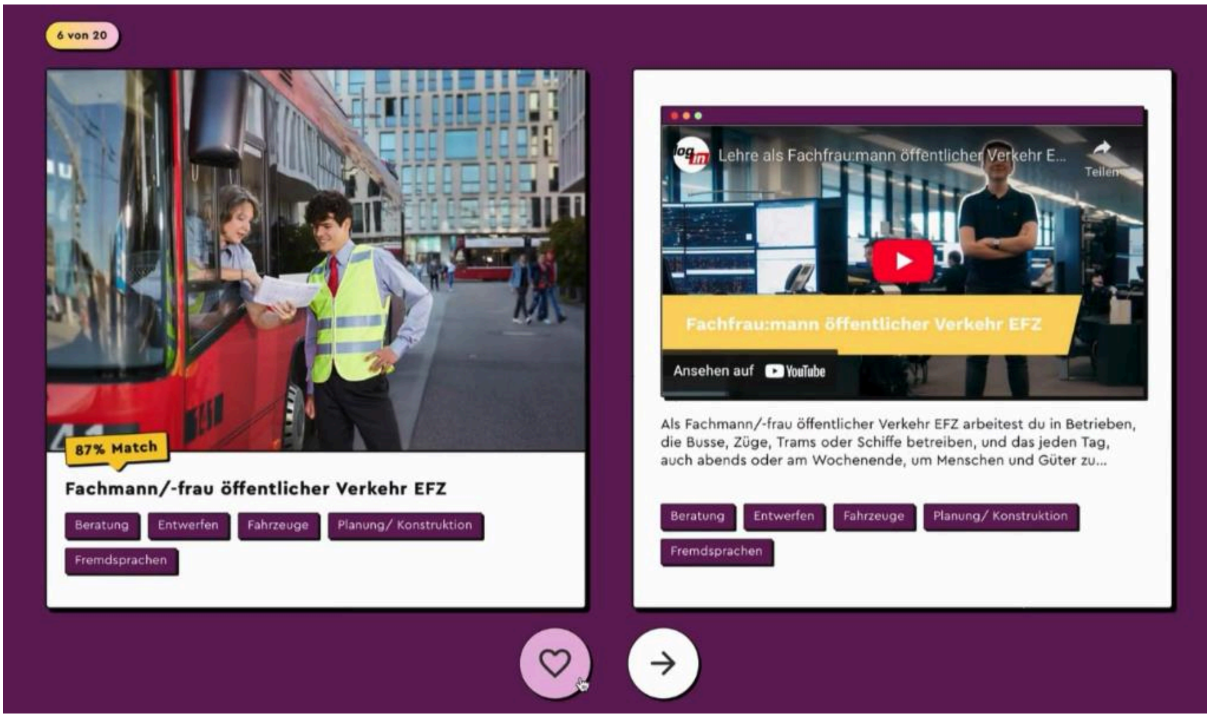

Figure A5: Sample Watch List from After the Redesign

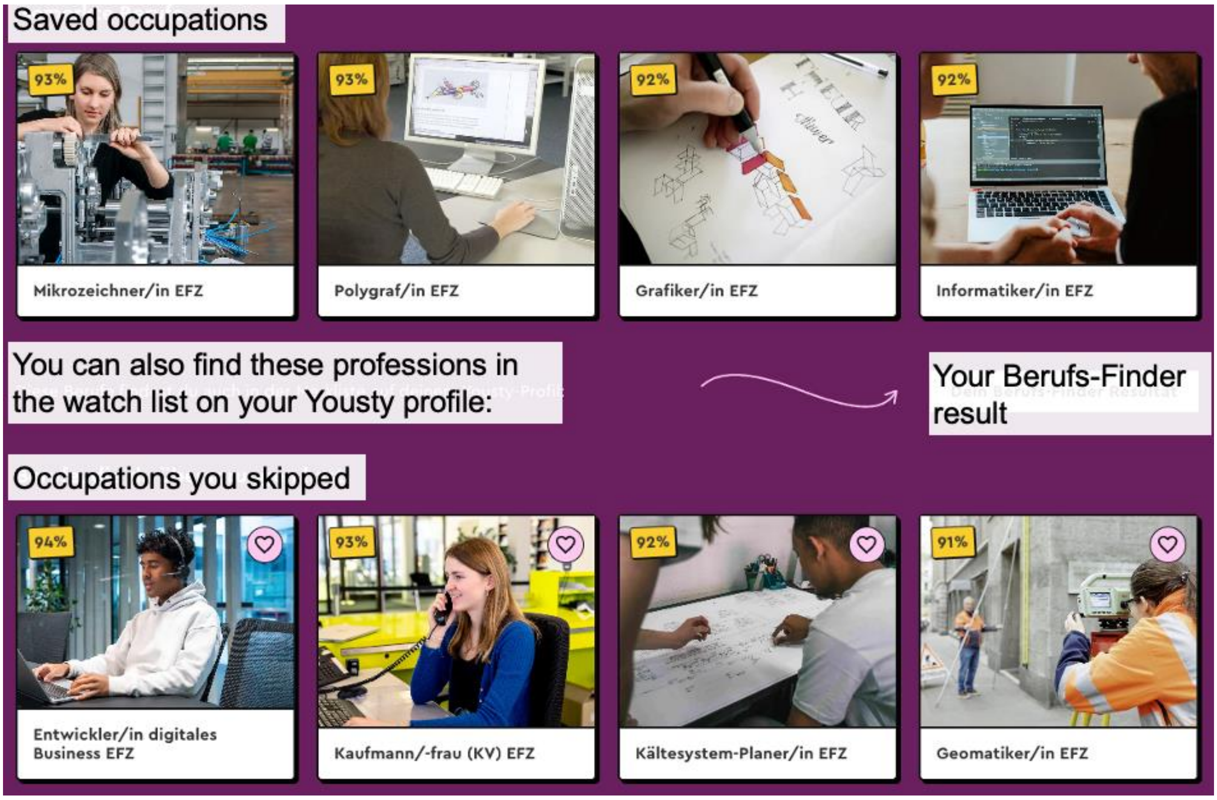

Figure A6: Baseline Rank Effects, Predictive Margins with 95% Confidence Intervals

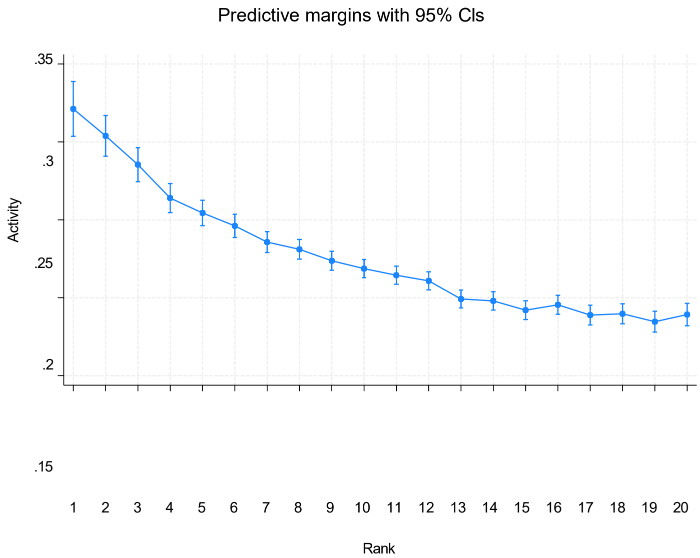

(a) Activity

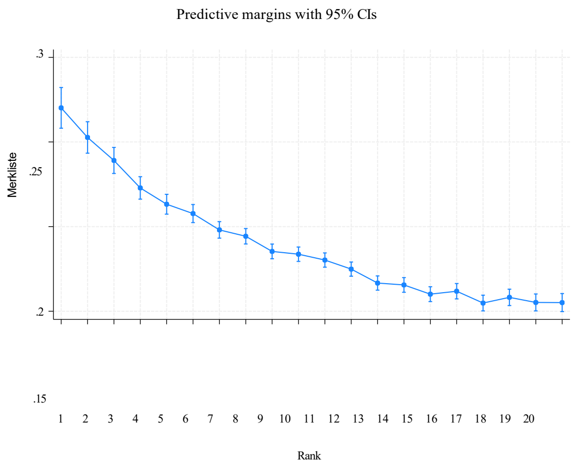

(b) Watch List

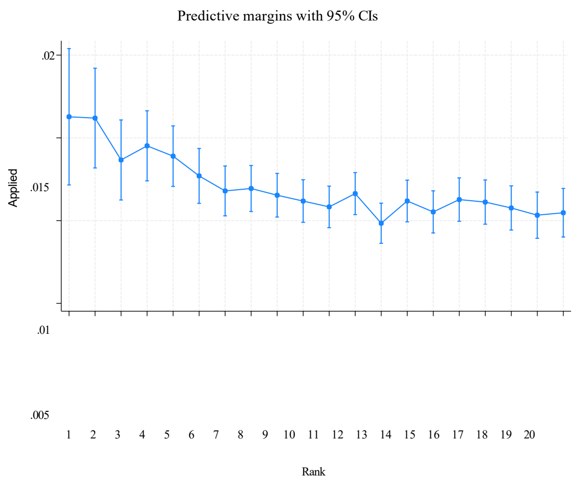

(c) Apply

Figure A7: RD Figures (outcome = Number of Occupations Applied, control = 2020-2022, 2024)

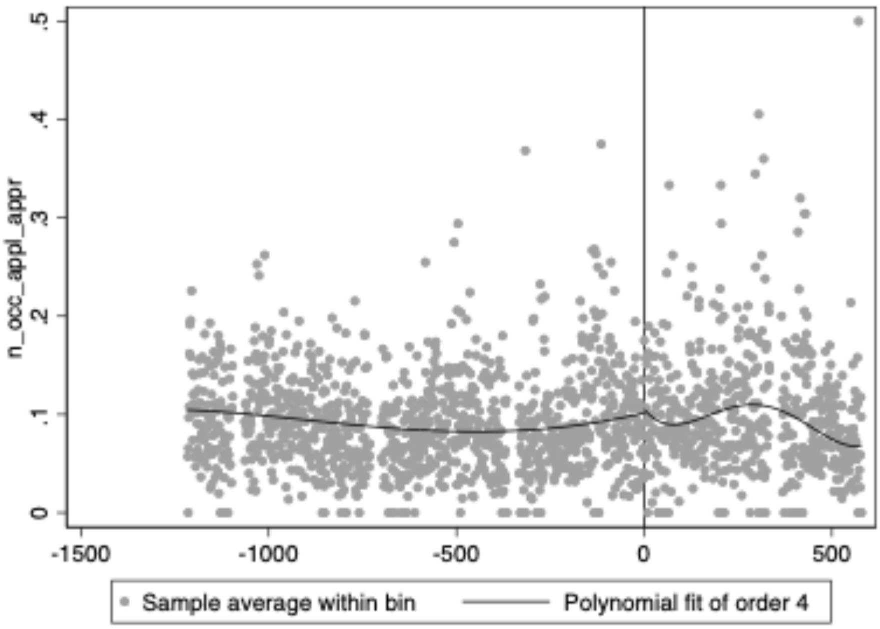

Figure A8: RD Figures (outcome = Saved Any Occupations, control = 2020-2022, 2024)

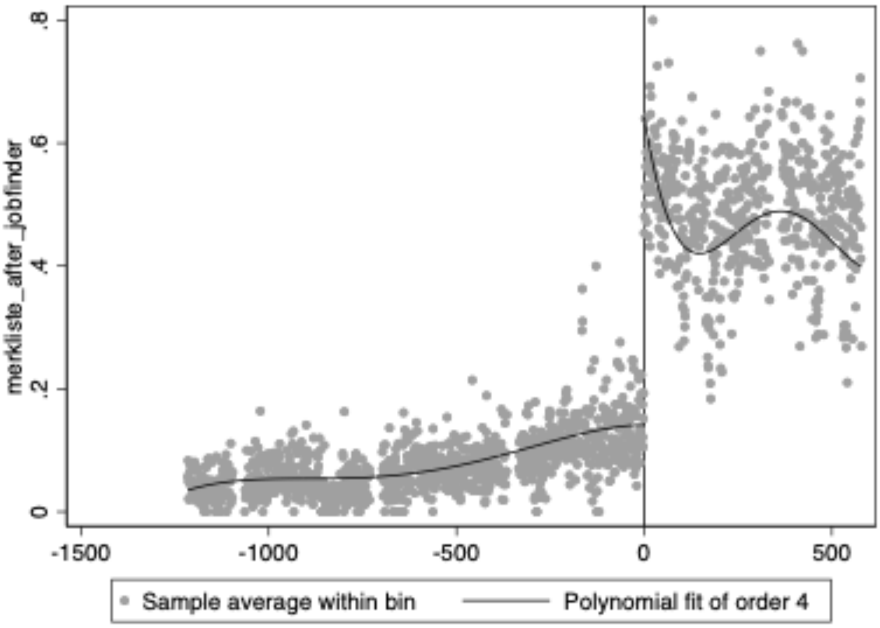

Figure A9: RD Figures (outcome = Number of Occupations Saved, control = 2020-2022, 2024)

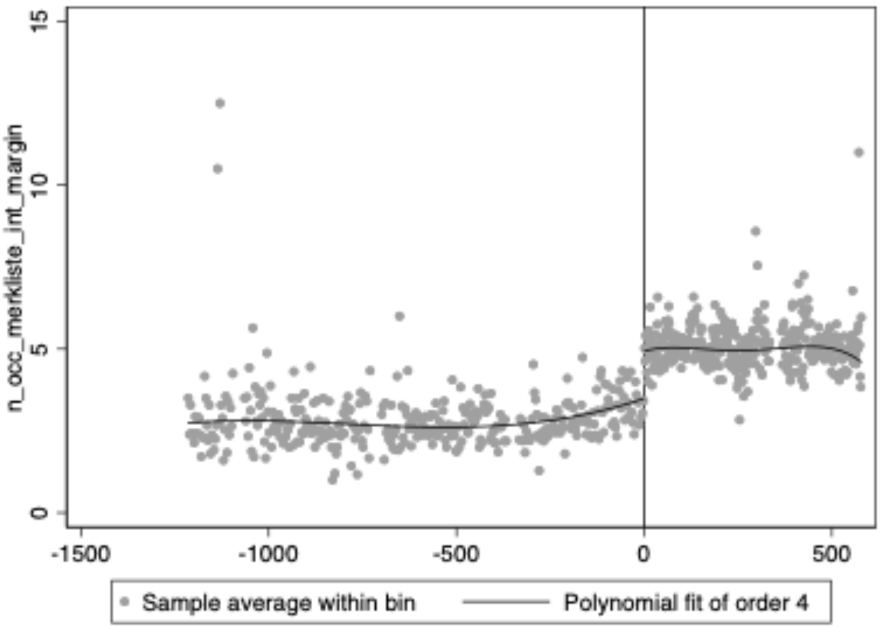

Note: "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list.

Table A1: Heterogeneity in Quadratic Rank Effect by Predicted Earnings

	Activity		Watch List		Apply	
	(1)	(2)	(3)	(4)	(5)	(6)
Rank	-0.0118***	-0.0111***	-0.0093***	-0.0103***	-0.0007**	0.0000
	(0.0013)	(0.0013)	(0.0010)	(0.0010)	(0.0003)	(0.0003)
Top Earnings Quartile	0.2510***	–	0.1433***	–	0.0412***	–
	(0.0094)	–	(0.0061)	–	(0.0024)	–
Rank x Top Earn. Q.	-0.0144***	-0.0143***	-0.0076***	-0.0074***	-0.0030***	-0.0030***
	(0.0018)	(0.0017)	(0.0012)	(0.0011)	(0.0004)	(0.0004)
Rank Squared	0.0003***	0.0003***	0.0002***	0.0003***	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Rank Sq. x Top Earn. Q.	0.0003***	0.0003***	0.0002***	0.0001***	0.0001***	0.0001***
	(0.0001)	(0.0001)	(0.0001)	(0.0000)	(0.0000)	(0.0000)
Individual x Score FE	✓	✓	✓	✓	✓	✓
Occupation FE		✓		✓		✓
N	568,698	568,698	568,698	568,698	568,698	568,698

Notes: Columns (1)-(2) use a count measure of activity, defined as the number of distinct actions taken with respect to an occupation: viewing any job advertisement of the occupation, viewing the occupation profile, saving the occupation to the watch list, applying to a trial apprenticeship, or applying to an apprenticeship. Columns (3)-(4) use a dummy for saving the occupation to the watch list. Columns (5)-(6) use a dummy for applying to an apprenticeship. Standard errors, clustered at the user level, are reported in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Table A2: DD-RD Estimates, Alternative Control Years

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Panel A: Outcome Variable = Occupations to which Yousty Users Apply				
Number of Occupations in Applications	0.034***	0.030**	0.037*	0.044*
	(0.013)	(0.015)	(0.019)	(0.026)
Panel B: Mechanism Variable = Occupations Saved to Watch List				
Saved Any Occupations to Watch List	0.439***	0.426***	0.388***	0.346***
	(0.012)	(0.014)	(0.018)	(0.024)
Number of Occupations on Watch List	1.732***	1.774***	1.987***	1.702***
	(0.247)	(0.280)	(0.351)	(0.467)

Notes: Control group years include 2019-2022 and 2024. Month of May is excluded in all years. Bandwidth around June 1st indicated in each column heading. "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Table A3: DD-RD Estimates, Alternative Control Years

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Panel A: Outcome Variable = Occupations to which Yousty Users Apply				
Number of Occupations in Applications	0.046***	0.044***	0.046**	0.057**
	(0.014)	(0.016)	(0.021)	(0.028)
Panel B: Mechanism Variable = Occupations Saved to Watch List				
Saved Any Occupations to Watch List	0.421***	0.406***	0.354***	0.325***
	(0.016)	(0.018)	(0.023)	(0.032)
Number of Occupations on Watch List	1.734***	1.783***	2.004***	1.659***
	(0.259)	(0.291)	(0.363)	(0.483)

Notes: Control group years include 2022 and 2024. Month of May is excluded in all years. Bandwidth around June 1st indicated in each column heading. "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Table A4: DD-RD Estimates, Including May

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Panel A: Outcome Variable = Occupations to which Yousty Users Apply				
Number of Occupations in Applications	0.026**	0.013	0.022	-0.008
	(0.012)	(0.014)	(0.019)	(0.026)
Panel B: Mechanism Variable = Occupations Saved to Watch List				
Saved Any Occupations to Watch List	0.224***	0.146***	0.026	-0.045*
	(0.013)	(0.015)	(0.019)	(0.027)
Number of Occupations on Watch List	0.314	-0.030	-0.194	-0.050
	(0.202)	(0.225)	(0.287)	(0.385)

Notes: Control group years include 2020-2022 and 2024. May is included in all years. Bandwidth around June 1st indicated in column heading. "Number of Occupations Saved to the Watch List" is conditional on a saving any occupations to the watch list. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Table A5: DD-RD Estimates, Including May

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Panel A: Outcome Variable = Occupations to which Yousty Users Apply				
Number of Occupations in Applications	0.025*	0.009	0.020	-0.007
	(0.013)	(0.015)	(0.020)	(0.027)
Panel B: Mechanism Variable = Occupations Saved to Watch List				
Saved Any Occupations to Watch List	0.218***	0.137***	0.025	-0.046*
	(0.012)	(0.014)	(0.019)	(0.027)
Number of Occupations on Watch List	0.314	-0.074	-0.225	-0.044
	(0.200)	(0.224)	(0.287)	(0.384)

Notes: Control group years include 2019-2022 and 2024. May is included in all years. Bandwidth around June 1st indicated in column heading. "Number of Occupations Saved to the Watch List" is conditional on a user saving any occupations to the watch list. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Table A6: DD-RD Estimates, Including May

	(1) 120 days	(2) 90 days	(3) 60 days	(4) 30 days
Panel A: Outcome Variable = Occupations to which Yousty Users Apply				
Number of Occupations in Applications	0.034**	0.019	0.029	-0.000
	(0.014)	(0.016)	(0.021)	(0.029)
Panel B: Mechanism Variable = Occupations Saved to Watch List				
Saved Any Occupations to Watch List	0.198***	0.103***	0.003	-0.066*
	(0.016)	(0.019)	(0.024)	(0.035)
Number of Occupations on Watch List	0.361*	-0.025	-0.108	0.037
	(0.213)	(0.237)	(0.301)	(0.404)

Notes: Control group years include 2022 and 2024. May is included in all years. Bandwidth around June 1st indicated in column heading. "Number of Occupations Saved to the Watch List" is conditional on a user saving any occupations to the watch list. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Table A7: Sharp RD Estimates, Applications to Apprenticeships

	(1)	(2)	(3)
Panel A: Outcome Variable = Occupations to which Yousty Users Apply			
RD Estimate	0.0285**	0.0279**	0.0241*
	(0.0120)	(0.0124)	(0.0140)
Observations	234,435	202,574	127,524
Control Years	2019-2022 & 2024	2020-2022 & 2024	2022 & 2024
Bandwidth	105.5	98.37	84.99
Eff. Obs. Left	9,183	8,383	6,994
Eff. Obs. Right	14,379	12,131	7,220
Panel B: Mechanism Variable = Saved Any Occupations to Watch List			
RD Estimate	0.435***	0.435***	0.421***
	(0.0192)	(0.0194)	(0.0204)
Observations	234,435	202,574	127,524
Control Years	2019-2022 & 2024	2020-2022 & 2024	2022 & 2024
Bandwidth	141.1	138	105.8
Eff. Obs. Left	10,834	10,463	9,183
Eff. Obs. Right	19,972	19,415	14,379
Panel C: Mechanism Variable = Number of Occupations Saved to Watch List			
RD Estimate	1.825***	1.826***	1.838***
	(0.192)	(0.192)	(0.193)
Observations	44,175	42,479	38,135
Control Years	2019-2022 & 2024	2020-2022 & 2024	2022 & 2024
Bandwidth	137.3	136.5	131.5
Eff. Obs. Left	1,390	1,378	1,346
Eff. Obs. Right	9,197	9,140	8,992

Notes: "Number of Occupations Saved to the Watch List" is conditional on a user having saved any occupations to the watch list. Standard errors in parentheses. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.