College as a Marriage Market*

Lars Kirkebøen, Edwin Leuven, Magne Mogstad, and Jack Mountjoy

March 20, 2025

Abstract

College graduates tend to marry each other. We use detailed Norwegian data to show that strong assortativity further arises by institution and field of study, especially among high earners from elite programs. Admission discontinuities reveal a dominant role for causal college match-making instead of inevitable selection: applicants who prefer a given program but are marginally rejected from it are only slightly more likely than random to match with a partner from that program, while the marginally admitted do so far more often. Marginal admission into elite professional programs propels women into elite household formation: they earn substantially more themselves and match with higher-earning elite partners. Marginal elite admission for men yields no change in partner earnings. College match-making effects are concentrated among students who attend the same institution at the same time, and are larger when opposite-sex peers are more abundant. These findings point to substantial search costs, especially in the elite marriage market.

Keywords: assortative mating, homogamy, marriage market, matching, household formation, search costs, returns to college, field of study, college major, college selectivity, elite education, professional education. **JEL codes:** D13, I23, I24, I26, J12, J24, J31

*Kirkebøen: Research Department, Statistics Norway. E-mail: kir@ssb.no. Leuven: Department of Economics, University of Oslo. Also affiliated with CESifo, IZA, and Statistics Norway. Email: edwin.leuven@econ.uio.no. Mogstad: Department of Economics, University of Chicago; Statistics Norway; NBER. Email: magne.mogstad@gmail.com. Mountjoy: Booth School of Business, University of Chicago; Statistics Norway; NBER. Email: jack.mountjoy@chicagobooth.edu. We thank Thomas Le Barbanchon, Mikkel Gandil, and seminar participants for valuable feedback and suggestions. We are grateful for support from the Norwegian Research Council (projects no. 237840 and 275906) and the Robert H. Topel Faculty Research Fund at the University of Chicago Booth School of Business.

1 Introduction

People have a strong tendency to sort into internally homogeneous marriages. A long-standing literature spanning many disciplines has studied such assortative mating because of its potentially important consequences for inequality within and across generations, as well as the reproduction of populations more generally (Schwartz, 2013).

Although education is one of the traits most intensively studied in the assortative mating literature, college graduates are commonly treated as a uniform group.[1] An emerging body of descriptive work, however, points to the possibility that the *type* of college education (institution or field of study) is an important but neglected pathway through which individuals further sort into homogeneous marriages.[2]

Assortative mating by college type may be economically consequential for at least two reasons. First, typical labor market outcomes vary substantially across different fields and institutions (Lovenheim and Smith, 2023). Matching by college type thus opens up far more opportunities for generating inequality across households compared to matching on schooling level alone. Second, to the extent that colleges enable such homogamy as match-makers, the nearly exclusive focus of the returns-to-college literature on individual earnings may substantially understate the role of higher education in the production of household inequality, especially if more lucrative programs also tend to have stronger match-making effects.

In this paper, we study college as a marriage market. We first document strong assortativity by institution and field of study among the college-educated, especially among high earners from elite programs. We then exploit admission discontinuities to reveal a dominant role for causal college match-making in the production of this assortativity, rather than the inevitable pairing of underlying types that merely correlate with field and institution. Finally, we explore the consequences of college match-making for earnings and the formation of elite households, and investigate mechanisms that shed light on how the marriage market works.

The context of our study is Norway's post-secondary education system. As detailed in Section 2, our work draws on two strengths of this environment. First, Norwegian population register data allow us to observe not only individuals' specific

[1] See literature reviews in Blossfeld (2009), Han and Qian (2024), and Eika et al. (2019).

[2] See Nielsen and Svarer (2009); Bičáková and Jurajda (2017); Eika et al. (2019); Ford (2020); Han and Qian (2024).

education types (institution and field) and labor market outcomes, but also if and whom they marry (or cohabit with). Second, a centralized admission process creates instruments for college type from discontinuities that effectively randomize applicants near unpredictable admission cutoffs into different institutions and fields of study.[3]

In Section 3, we use this data to document striking patterns of homogamy by college type. Among all college-educated couples, one out of every three completed degrees from the same institution. One out of six majored in the same field. One out of eight completed the same program, sharing degrees in the same field from the same institution. These rates of homogamy remain roughly constant across most of the household income distribution, but then increase dramatically in the top few percentiles. Among college-educated couples in the top one percent of household income, fully half are homogamous by institution, 40 percent by field, and one third by both field and institution. We show that the entirety of this spike in college homogamy at the top of the income distribution is driven by couples who completed degrees from elite professional programs in law, medicine, business, and engineering. We then compare the observed homogamy rates to random and maximal benchmarks, showing that assortativity by program, field, institution, and elite education achieve substantial fractions of their maximum feasible extent.

These descriptive results raise the question of why college graduates are so likely to match with someone from the same institution and field of study. One possibility is pure selection: individuals may match on underlying traits like ability, personality, taste, and family background that merely correlate with college type. Another possibility is causation, where enrolling in a particular program actively influences one's match. Such causal college match-making effects could operate through a number of mechanisms, including search costs and type changes.[4]

To disentangle these explanations and quantify their relative contributions, we exploit features of the college admission system that effectively randomize applicants near unpredictable admission cutoffs into different programs. In Section 4, we describe these institutional features and their utility in separately identifying causal

[3]Kirkeboen et al. (2016) use these discontinuities to show that earnings payoffs vary substantially by field of study. The results highlight the limitations of treating the college-educated as a uniform group.

[4]Both theory and evidence suggest marriage decisions are increasingly driven by returns to matching on similarities (e.g. due to leisure complementarities), rather than potential gains from trade (see the review in Juhn and McCue, 2017).

effects of a given type of education from systematic selection into it. We find that crossing the admission threshold into a preferred program has no substantial impacts on the "extensive margins" of education and matching, with little change in the probability of enrolling in any college, completing any degree, matching with any partner, or matching with any college-educated partner. What does change is the type of education, and partner, acquired: crossing the threshold into a preferred program substantially increases the probability of getting admitted to that program, enrolling in it, completing a degree from it, and, importantly, matching with a partner who completed a degree from it.

To translate these reduced-form discontinuities into treatment effects of enrollment, we estimate an instrumental variables model that uses threshold crossing into an applicant's preferred program as an instrument for enrolling in it. We use this model to decompose the high rates of observed college homogamy into causal college-match making versus selection. We explicitly measure selection by comparing the match outcomes of untreated compliers – applicants who prefer a given program but are exogenously denied access due to falling just short of the admission threshold – versus the match outcomes we would expect if those applicants simply matched with partners randomly. This comparison thus reveals any systematic matching proclivities among those who prefer a particular program, independent of actually getting treated by it.

The results, presented in Section 5, reveal a dominant role for causal college match-making over selection in the production of college homogamy. Applicants who prefer a given program become much more likely to match with a partner who completed that program as a direct result of actually enrolling in it. Untreated applicants who identically prefer the same program but are marginally denied access to it are only slightly more likely than random to match with a partner from that program, revealing little role for selection. This pattern persists when considering applicants on the margin between different fields (regardless of institution), different institutions (regardless of field), and elite versus non-elite programs. Enrolling in an elite professional program generates especially large match-making effects for female applicants and those from more educated families.

College enrollment itself, rather than selection into it, thus primarily drives homogamous matching among the college-educated. But are these college match-making effects economically consequential? We find that the types of applicants who experi-

ence larger match-making effects of enrolling in their preferred option also tend to experience larger gains in household income. Elite professional programs, in particular, propel marginally admitted women into elite household formation: they earn substantially more themselves and match with substantially higher-earning elite-educated partners, becoming much more likely to join the top few percentiles of the household earnings distribution. Marginally admitted elite men, on the other hand, see statistically insignificant changes in both their own earnings and those of their partners. Results also diverge by parental education. For applicants with at least one college-educated parent, elite enrollment substantially increases their earnings and their likelihood of entering the top ranks of the household income distribution. For applicants with no college-educated parents, elite enrollment yields smaller earnings gains and no greater chance of entering the top earnings percentiles.

In Section 6, we investigate the mechanisms behind these college match-making effects to shed light on how this marriage market works. If enrolling in a particular option changed an individual's type in a broad matching market with low cost of finding partners, then we would expect some of the enrollment effects to be driven by matches across different institutions and cohorts, since the particular location and timing of the education should not matter much for matching on the new type. Furthermore, we would not expect the density of opposite-sex peers within a particular institution to matter much for matching in a broader market. Instead, we show that college match-making effects are concentrated among students who attend the same institution at exactly the same time, and are larger at institutions where opposite-sex peers are more abundant, suggesting that search costs are an empirically important feature of the marriage market. These results do not imply that type changes are an unimportant consequence of enrollment; the earnings effects we estimate indeed show transformations in labor market outcomes from enrolling in the preferred option, especially elite programs. But our results do suggest that acting on such type changes in the marriage market is more difficult outside of the particular pool of peers who find themselves on campus together at the same time, and even inside that pool when opposite-sex peers are more scarce.

Altogether, our findings suggest that colleges are effectively local marriage markets, mattering greatly for whom one marries, not because of the pre-determined traits of the admitted students but as a direct result of attending a particular institution at a given time. Our ability to credibly distinguish causality from selection

and highlight key mechanisms makes these findings relevant for students, researchers, and policymakers. For students, if strong assortative mating by institution and field simply reflected selection, then it would be immaterial for their incentives and returns to educational choices; even in the absence of enrolling in a particular field or institution, they would be just as likely to match with similar partners. For researchers, our findings of the dominance of causality over selection, economic consequences for household earnings and inequality, and the presence of search costs contribute useful new facts to guide theoretical and empirical work on marriage markets, educational decisions, and economic inequality. Finally, policymakers concerned with the returns to higher education and the causes of social stratification, including the formation of elite households, may also be interested in these findings that credibly distinguish the causal contributions of colleges from selection into them, revealing trade-offs in boosting household incomes in ways that can exacerbate dynastic inequalities.

This paper connects and advances multiple literatures. First, we broaden the returns-to-college literature, which focuses almost exclusively on individual earnings, with rare causal evidence on the marriage market and household earnings effects of college choices.[5] Only a handful of related studies draw credibly causal inferences about how college choices affect marriage market outcomes. Artmann et al. (2021) use admission lotteries to four oversubscribed post-secondary programs in the Netherlands and find that field of study matters for partner choice. Kaufmann et al. (2021) study admission into elite programs in Chile using a regression discontinuity design, finding sizable effects on partner socioeconomic status. Barrios-Fernandez et al. (2024) also use admission discontinuities to study effects of elite programs in Chile, with a focus on intergenerational consequences for the applicant's children but a brief analysis of the characteristics of the co-parent.

[5]See Mountjoy (2024) and the literature reviewed therein. Related papers to this one include Kirkeboen et al. (2016), who study (individual) earnings payoffs to different fields of study in Norway; Zimmerman (2019), who studies the (individual) earnings returns to elite professional programs in Chile, finding a similar pattern to this paper of larger returns for applicants from more advantaged backgrounds; and the literature on causal (individual) earnings returns to attending more selective institutions, including Dale and Krueger (2002), Hoekstra (2009), Mountjoy and Hickman (2021), Chetty et al. (2023), and Bleemer (2024). None of these papers study household returns inclusive of effects on matching and spousal income; Smith et al. (2025) is a rare exception with data on predicted household income based on credit report variables. Oreopoulos and Salvanes (2011) review the literature on non-pecuniary returns to education, in which causal evidence on the effects of college choices on marriage market outcomes is scarce; see Goldin (1992) for descriptive historical evidence and Ge et al. (2022) for rare estimates of the marriage market effects of institutional selectivity.

Another body of work seeks to quantify the importance of search frictions and meeting opportunities for assortative mating.[6] Most related to this paper is Nielsen and Svarer (2009), who use Danish data to document the extent to which individuals match on education length and type.[7] They find that around half of educational homogamy can be attributed to the tendency of individuals to marry someone who went to the same educational institution or to an institution nearby. Nielsen and Svarer (2009) conclude that this could be due to search frictions, or due to selection of people with the same preferences into the same institution. To address this identification challenge, a few studies have taken advantage of detailed data from the dating market.[8] Hitsch et al. (2010) find that search frictions may play an important role in explaining observed matching patterns by education at an online dating site. Belot and Francesconi (2013) use data from a speed dating agency to identify the role of opportunities separately from that of preferences. Their findings suggest the role of individual preferences is outweighed by that of opportunities.

We connect and advance these literatures in several ways. First, we are able to comprehensively measure patterns of college homogamy across the entire population, and establish new facts about elite homogamy at the top of the household earnings distribution, through the use of population register data that links every individual not only to their own detailed education and earnings records but also to those of their partner, if any. Second, we are able to go beyond descriptive patterns and decompose total homogamy into selection versus treatment by exploiting discontinuities in the centralized admission system. Third, we leverage exogenous variation along multiple margins of college choice, allowing us to separately study impacts of enrollment in preferred programs, fields, institutions, and elite versus non-elite education. Fourth, this variation also allows us to explore which specific mechanisms of the enrollment treatments drive college match-making effects to understand how the marriage market

[6] A large theoretical literature studies equilibrium sorting patterns in marriage markets with search frictions. See, for example, Burdett and Coles (1997); Eeckhout (1999); Bloch and Ryder (2000); Shimer and Smith (2000); Atakan (2006); Jacquet and Tan (2007); and the review by Chiappori (2020). A complementary literature evaluates different methods of measuring of assortative mating; see the recent contribution by Chiappori et al. (2025).

[7] See also Mansour and McKinnish (2018), who use survey data from the US to document that same-occupation matching is strongly related to the sex composition of the occupation. To distinguish between a preferences explanation and a search cost explanation, they investigate whether women accept lower-wage husbands if they match within-occupation compared to if they do not, and how this wage gap varies with the sex composition of the occupation.

[8] See also Fisman et al. (2006, 2008); Bruze (2011); Banerjee et al. (2013); Lee (2016).

functions. Finally, we use high-quality tax data to explore consequences for household earnings, decomposed into own plus partner earnings, and shed new light on the role of higher education in the production of elite households.

2 Setting and data

In this section, we describe the Norwegian higher education sector, introduce our data sources, and define key variables.

2.1 Higher education in Norway

The Norwegian post-secondary sector consists of about ten public universities and a much larger number of public and private university colleges. The vast majority of students attend a public institution, and even the private institutions are funded and regulated by the Ministry of Education and Research. A post-secondary degree normally lasts 3-5 years. Applicants apply to programs, which specify both a field and an institution (e.g. Business at the University of Oslo). The universities all offer a wide selection of fields. The university colleges rarely offer fields like Law, Medicine, Science, or Technology, and instead focus on professional degrees in fields like Engineering, Health, Business, and Teaching. There are generally no tuition fees, and most students are eligible for financial support for living expenses (part loan, part grant) from the Norwegian State Educational Loan Fund.

The main universities are located in the major cities of each of the five regions: Bergen and Stavanger (West), Oslo (East), Kristiansand (South), Trondheim (Central) and Tromsø (North). In addition, there are a few other universities and many university colleges spread across the country. Figure 1 displays the distribution of the post-secondary student population across Norwegian municipalities in the years 1998–2004. About 60 percent live in Oslo, Bergen or Trondheim, the three biggest cities in Norway, with sizable student populations in several other municipalities, including Tromsø, Kristiansand, and Stavanger.

2.2 Data sources and key variable definitions

Our analysis employs several nationwide administrative data sources that we can link through unique individual identifiers. We start with the Central Population

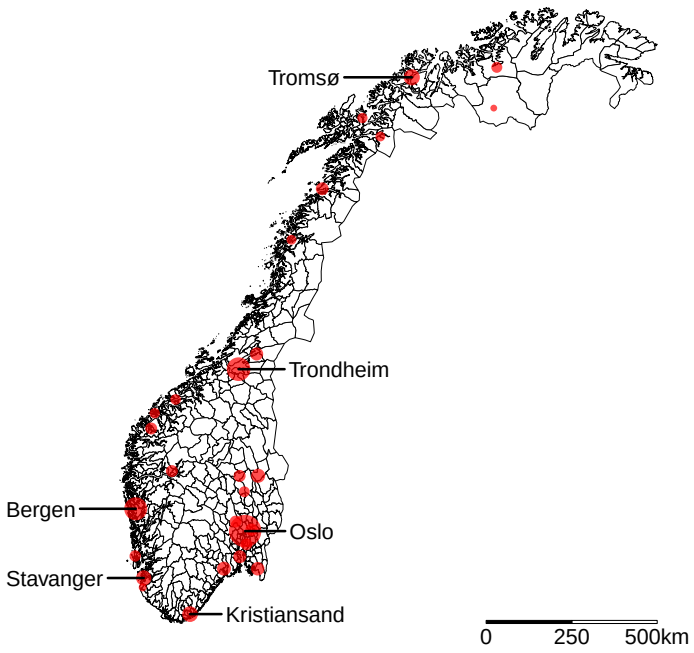


Figure 1. Higher education enrollment across Norway

Register, which provides demographic information and household membership for every Norwegian resident from 1967 to 2018. Following Norwegian official statistics, we define a match as an adult couple who are either married or registered cohabitants. For each adult, we define their partner from their first observed match, if any.

We link each individual and partner to records of college enrollment spells and completed degrees from the National Education Register. These records are reported directly from schools and, importantly for our purposes, include detailed indicators of field and institution, the combination of which defines a program. We observe roughly 50 distinct fields, 100 institutions, and 1,400 programs across the entire population. We define a college-educated couple as homogamous if they both completed post-secondary degrees from the same program, field, or institution, respectively.

Following Cattan et al. (2022) and Bütikofer et al. (2018), we define elite professional programs as those in Medicine or Law at any institution, Business at the Norwegian School of Economics (NHH), and Engineering at the Norwegian University of Science and Technology (NTNU). These highly selective programs enroll only a tiny fraction of a given birth cohort (roughly 3 percent) but produce a disproportionately large share of private and public sector leaders and top-income earners

(Kirkebøen, 2010). Despite the seemingly equalizing features of the Norwegian higher education system—zero tuition fees, subsidized living expenses, an absence of elite private feeder high schools, and transparent admission policies with no regard to legacy status or extracurriculars—Cattan et al. (2022) show that elite enrollment in Norway exhibits a steep socioeconomic gradient similar to elite programs in much more unequal countries like the US (Chetty et al., 2023), the UK (Britton et al., 2021), and Chile (Barrios-Fernandez et al., 2024).[9] Figure A1 plots the average high school GPA and average adult earnings associated with each post-secondary program, with elite programs highlighted near the top of both dimensions.

We measure annual pre-tax individual earnings and partner earnings (the sum of which defines household earnings) from tax records through 2018, and we observe employers from the matched employer-employee register. For our regression discontinuity sample of applicants, described further in Section 4, we also merge in application data from the Norwegian Universities and Colleges Admission Service, which cover nearly all applications to post-secondary education in Norway for the years 1998-2004.

3 Patterns of college homogamy

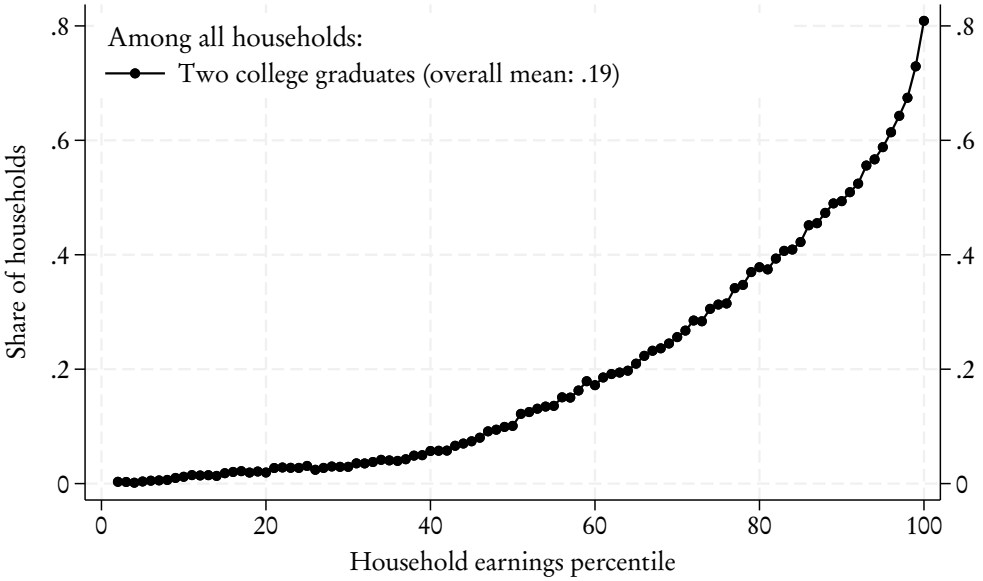
In this section, we describe patterns of college homogamy on average and across the household income distribution, revealing high rates of homogamy that increase dramatically among top earners from elite programs. We then compare the observed homogamy rates to random and maximal benchmarks, showing that assortativity measures by program, field, institution, and elite education achieve substantial fractions of their maximum feasible extent.

3.1 College homogamy across the household income distribution

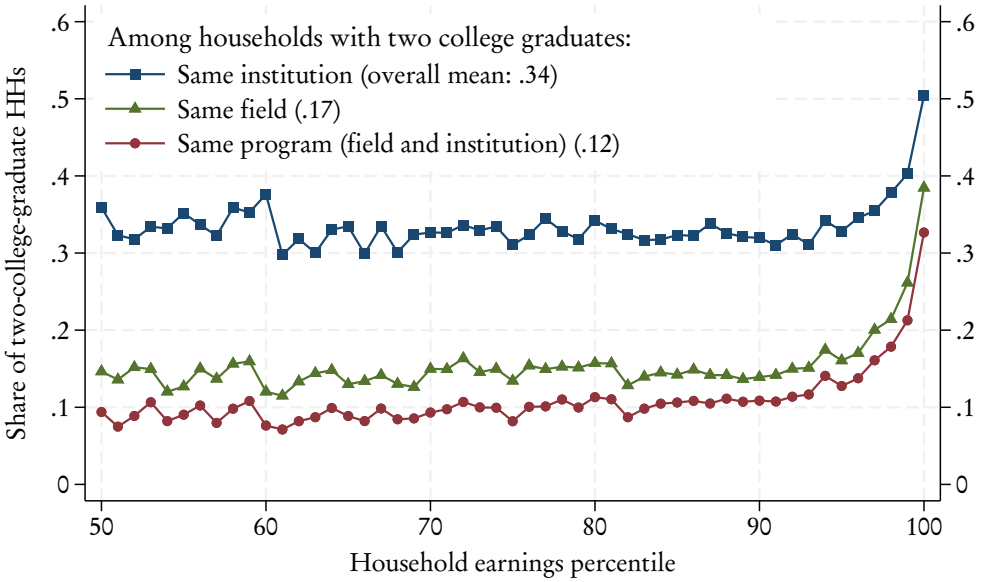
Figure 2 visualizes college homogamy across the household earnings distribution. We consider all households in Norway in 2018 (including singles) in which the registered household head is 30-45 years old. We average annual pre-tax household earnings over the years in which the household head is 30-45 years old, and stratify by the percentile of this earnings measure within the birth cohort of the household head. The top panel of Figure 2 shows that, on average, 19 percent of all households are

[9]Broader patterns of Scandanavian educational mobility across generations are also quite similar to the United States, despite large differences in income mobility (Landersø and Heckman, 2017).

(a) Homogamy with respect to college completion



(b) Homogamy with respect to institution, field, and program



Note: The sample includes all households in Norway in 2018 in which the registered household head is 30-45 years old. The horizontal axis is the within-birth-cohort percentile of annual pre-tax household earnings averaged over the years in which the household head is 30-45. The top panel includes single households. The bottom panel keeps the same income measure but zooms in and conditions the vertical axis shares on households in which both partners are domestic college graduates.

Figure 2. College homogamy by household income

composed of two college graduates, but this rate spans from zero at the bottom of the income distribution to 80 percent at the top.

The bottom panel of Figure 2 conditions on households with two domestic college graduates to ask how their homogamous composition changes with income. On average, one out of every three college-educated couples (34 percent) completed degrees from the same institution. One out of six (17 percent) majored in the same field. One out of eight (12 percent) completed the same program, which means the couple share degrees in the same field from the same institution. These rates of homogamy remain roughly constant across most of the household income distribution, but then increase dramatically at the top few percentiles. Among college-educated couples in the top one percent of household income, half are homogamous with respect to institution, and nearly 40 percent are homogamous with respect to field. One third are homogamous with respect to both, meaning the rate of program homogamy among the richest households is nearly three times the average.

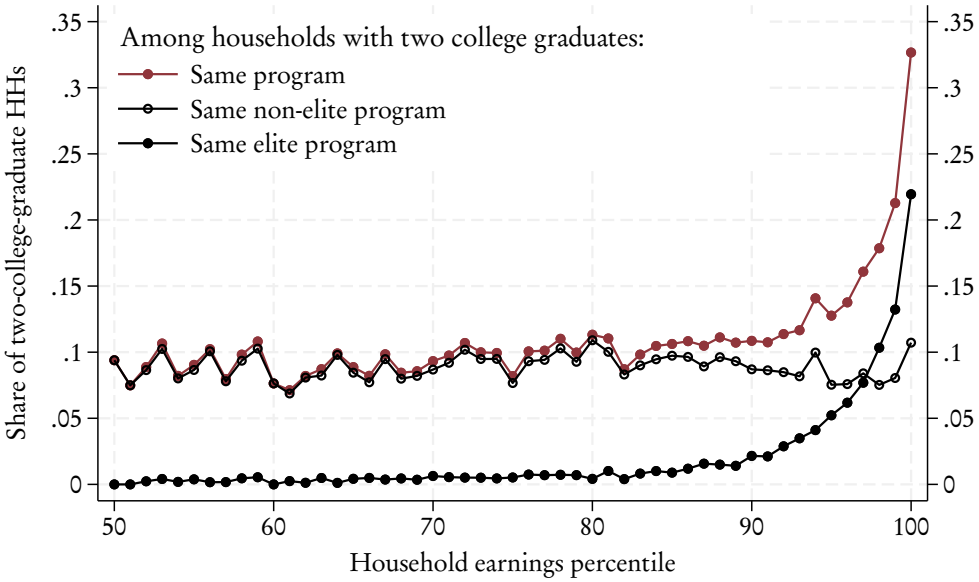
The top panel of Figure 3 decomposes program homogamy from Figure 2(b) into whether the shared program is non-elite versus elite. Up until roughly the 85th percentile of household income, virtually all college homogamy is driven by couples from non-elite programs. Beyond the 85th percentile, the rate of non-elite program homogamy continues at its roughly constant value, while the rate of elite program homogamy surges upward. Thus, the entirety of the spike in program homogamy at the top of the income distribution is driven by elite-educated couples.

The bottom panel of Figure 3 further explores the surge in elite homogamy at the top of the income distribution by considering couples who are from any elite program. Overall, only 5 percent of college-educated couples are two graduates from elite programs. Such double-elite couples are nearly non-existent outside of the top decile of household income, and then proliferate in the final few percentiles. Among the highest-earning college-educated couples, fully one third are double-elite, with the majority of those sharing the same institution *and* field, and thus program.

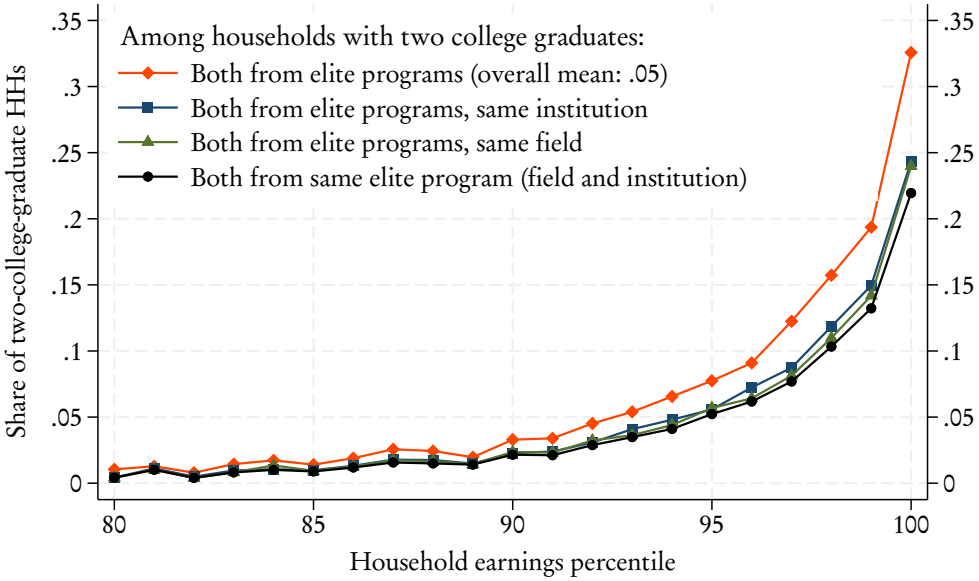
3.2 Interpreting the magnitudes of college homogamy

In Table 1, we interpret the documented magnitudes of college homogamy through comparisons to random and maximal benchmarks, leading to a natural measure of the strength of assortativity. The first column reports the overall observed homogamy rates among couples with two college graduates: 12 percent share the same program,

(a) The role of elite programs in high-income homogamy



(b) Elite homogamy



Note: These figures keep the same income percentile measure as Figure 2. The vertical axis shares condition on households in which both partners are domestic college graduates.

Figure 3. Elite college homogamy

17 percent share the same field, 34 percent share the same institution, and 5 percent are both elite-educated.

These homogamy rates do not, in themselves, necessarily imply that assortativity is stronger with respect to, say, institution as compared to field. This is because homogamy rates depend not only on the degree of sorting but also on the prevalence of men and women with each type of education.[10] First, homogamy rates could be larger with respect to institution even if college-educated men and women matched randomly. The second column of Table 1 reports the random homogamy rate

$$h^r = \sum_e s_e^m \times s_e^w, \tag{1}$$

where e indexes the values of a given education type (e.g. different institutions), s_e^m denotes the share of college-educated matched men with education type e, and s_e^w denotes the share of college-educated matched women with education type e. Under random matching, program homogamy would be rare (0.8 percent of college-educated couples), given the disperse distribution of programs. Homogamous random matches would be a bit more common with respect to fields, institutions, and elites, but the observed homogamy rates are still much larger than these random benchmarks, leading to the absolute assortativity measure in the fourth column of Table 1.

Homogamy rates are also constrained by sex imbalances within each education type. This is most easily illustrated with two types of education. If men and women had equal distributions of education, then perfect assortativity would be feasible (i.e., the homogamy rate could be one). On the other hand, if all men had education of one type while all women had education of the other type, then no assortativity would be feasible (i.e., the homogamy rate must be zero). More generally, the number of men who can match to a woman with the same education type is obviously bounded by the number of men with that type, but further bounded by the number of women with that type if there are fewer such women than men. Adding up across types, the maximal rate of feasible homogamy is therefore given by

$$h^m = \sum_e \min(s_e^m,\, s_e^w) \tag{2}$$

[10]Fewer than one percent of couples in our sample are same-sex, so we consider men and women as the two sides of the market for parsimony in constructing the random and maximal benchmarks in Table 1.

Table 1. Homogamy and assortativity among college-educated couples

	Homogamy			Assortativity	
	Observed h	Random h^r	Maximal h^m	Absolute $h - h^r$	Rescaled $\frac{h-h^r}{h^m-h^r}$
Same program	.123	.008	.643	.114	.180
Same field	.165	.038	.670	.128	.202
Same institution	.343	.053	.845	.290	.366
Both elite	.049	.018	.108	.031	.346

Note: The sample comprises all couples in Norway in 2018 in which the household head is 30-45 years old and both partners are domestic college graduates. Observed homogamy is the share of couples who both completed degrees from the same program, field, or institution, or are both elite-educated. Random homogamy is given by $h^r = \sum_e s_e^m \times s_e^w$, where e indexes the values of a given education type (e.g. different institutions), s_e^m denotes the share of college-educated matched men with education type e, and s_e^w denotes the share of college-educated matched women with education type e. Maximal homogamy is $h^m = \sum_e \min(s_e^m, s_e^w)$.

and reported in the third column of Table 1. Since men and women tend to sort systematically into different fields, sex imbalances within fields lead to lower ceilings on field and program homogamy. Men and women are relatively more balanced within institutions, leading to a higher ceiling on institutional homogamy. Elite homogamy is capped at 11 percent: 16 percent of college-educated matched men are elite-educated, but only 11 percent of women are. Thus, at most 11 percent of college-educated couples could have both an elite man and an elite woman.

To combine these ingredients into an informative measure of assortativity, we follow Liu and Lu (2006) and recenter each observed homogamy rate h relative to random matching and then scale by the positive range from random to maximal:

$$R = \frac{h - h^r}{h^m - h^r}. \tag{3}$$

This measure can be interpreted as the share of maximum feasible assortativity that is achieved by the observed homogamy rate. The absence of assortativity would correspond to $R = 0$ (purely random matching), while $R = 1$ would denote the achievement of maximum feasible assortativity given the marginal distributions of male and female education types. The final column of Table 1 shows that assortativity by program achieves 18 percent of its feasible maximum, and assortativity by field achieves 20 percent. Assortativity is especially strong with respect to institution and

elite status, with both achieving over a third of their maximum feasible extent.[11]

4 Decomposing homogamy into selection versus treatment

The previous section documented the strong tendency of similarly-educated college graduates to match with each other, especially among the elite. How much of this observed homogamy is actively caused by colleges being marriage markets, and how much is selection? To fix ideas, consider that 33 percent of elite college graduates in our data match with an elite partner, while only 9 percent of non-elite graduates do so. On one hand, an elite education itself may cause an individual to become more likely to match with an elite partner (for reasons explored further in Section 6). On the other hand, the types of individuals who select into elite programs may have underlying traits, like intelligence and family wealth, that destine them towards matching with elite partners regardless of their actual education.

To disentangle these two explanations, we exploit features of the college admission system that effectively randomize applicants near unpredictable admission cutoffs into different programs. In this section, we describe these features and their utility in separately identifying causal effects of a given type of education from systematic selection into it. We then describe our estimation sample of applicants, empirically assess the validity of the proposed research design, document what does and does not change across the admission cutoffs, and lay out our instrumental variables model.

4.1 Admission process and identification strategy

The college admission process in Norway is centralized. Individuals submit their application to a single organization, the Norwegian Universities and Colleges Admission Service, which handles admission into all universities and most university colleges. Applicants rank up to fifteen programs in their application. For many programs, demand exceeds supply, with the latter determined by annual funding decisions from the Ministry of Education and Research. Offers of admission are made according

[11] Table A1 conducts a variation on this exercise where the matching market is taken to be the regional county, such that potential matches happen within but not across the 20 regions of Norway. The results end up being very similar to those in Table 1. As expected, random matches within region are a bit more likely to be homogamous with respect to institution. At the same time, the ceiling on feasible institutional homogamy is a bit lower, overall leading to a similar rescaled measure of institutional assortativity (34 percent).

to a sequential dictatorship mechanism where the order is determined by application scores. These scores are a function of the applicant's high school grade point average and a few other secondary criteria, including fulfillment of military service, choosing specific subjects in high school, age at application, previous education, and under-representation in a given field. The applicant with the highest application score receives an offer from her most preferred program; the second highest applicant receives an offer from her most preferred program among the programs that still have slots; and so on. This process continues until slots or applicants run out.[12] If students want to transfer to a different field or institution, they typically need to participate in the subsequent year's admission process on equal terms with other applicants.

Under this system, applicants scoring above a certain threshold are much more likely to receive an offer from a preferred program compared to applicants with the same preferences but marginally lower application scores. We first consider how to use these discontinuities to identify causal impacts of programs. To this end, panel (a) of Table 2 is sufficient. This panel illustrates an applicant on the margin of getting different field offers from the same institution. Suppose the applicant has an application score of 49. In this case, she would marginally qualify for an offer from her 3rd ranked program: Field 2 at Institution A. This defines her *preferred program* in the local program ranking around her application score. Now consider the applicant at the bottom of panel (a) who has a slightly lower application score of 47. This applicant has the same local program ranking, but does not receive an offer from the preferred program (Field 2 at Institution A) due to the marginally lower score. By comparing the outcomes of applicants like these – with the same locally preferred program but application scores just above versus below the program's admission threshold – we can identify the effect of crossing the threshold into the preferred program, as long as applicants cannot perfectly manipulate their scores to fall just above the threshold

[12]This procedure generates a first set of offers which are sent out to the applicants in late July. Applicants then have a week to accept the offer, if they get one. Irrespective of whether they accept, applicants can choose to remain on a waiting list for preferred program options, or withdraw from the application process. The slots that remain after the first round are then allocated in a second round of offers in early August among the remaining applicants on the waiting list. These new offers are generated following the same sequential dictatorship mechanism as in the first round. Since applicants in this second round can only move up in the offer sequence, second round offers will either correspond to first round offers, or be an offer for a higher ranked program. By choosing to remain on a waiting list, an applicant accepts that their first round offer is automatically discarded if they get a higher-ranked offer in the second round. In mid-August, the applicants begin their study in the accepted program.

Table 2. Illustration of identifying variation in programs, fields, and institutions

(a) Fields

Program Ranking	Inst.	Field	Cutoff
1st best	A	1	57
2nd best	B	1	52
3rd best	A	2	48
4th best	A	3	45
	Application score = 49		
Local Ranking	Inst.	Field	Offer
Preferred	A	2	Yes
Next-best	A	3	No
	Application score = 47		
Local Ranking	Inst.	Field	Offer
Preferred	A	2	No
Next-best	A	3	Yes

(b) Institutions

Program Ranking	Inst.	Field	Cutoff
1st best	B	1	52
2nd best	A	2	48
3rd best	B	2	46
4th best	B	3	43
	Application score = 49		
Local Ranking	Inst.	Field	Offer
Preferred	A	2	Yes
Next-best	B	2	No
	Application score = 47		
Local Ranking	Inst.	Field	Offer
Preferred	A	2	No
Next-best	B	2	Yes

(which we investigate below).[13]

Furthermore, this admission process creates exogenous variation not only in programs but also in fields and institutions. This allows us to quantify the relative importance of field versus institution for matching outcomes. To see this, consider both panels of Table 2. In panel (a), the two applicants are on the margin of getting offers from different fields (2 vs. 3) but the same institution (A). The applicant with the score of 49 marginally qualifies for an offer from her *preferred field* (Field 2) in the local ranking of fields around her application score. The other applicant with the score of 47 has the same local preferences over fields, but is not offered Field 2 because of the slightly lower score. By comparing the outcomes of these applicants we can identify the effect of getting an offer from the preferred field.

Panel (b) illustrates a case where two applicants are on the margin of getting offers from the same field (2) but different institutions (A vs. B). The applicant with the

[13]Note that some applicants can be on two local margins simultaneously. For example, the applicant with a score of 49 in panel (a) can be on the margin between (B, 1) and (A, 2), as well as the margin between (A, 2) and (A, 3). In our analysis below we stack both margins. However, only about 15–20 percent of applicants are observed on two margins, and our estimates do not materially change if we exclude these applicants (see robustness analyses in Section 5).

score of 49 marginally qualifies for an offer from Institution A, her *preferred institution* in the local ranking of institutions around her application score. The other applicant with the score of 47 has the same local preferences over institutions, but is not offered Institution A because of the slightly lower score. By comparing the outcomes of these applicants we can identify the effect of getting an offer from the preferred institution.

4.2 Estimation sample

To implement this research design, we use the near-universe of post-secondary application records for the years 1998-2004 from the Norwegian Universities and Colleges Admission Service. We observe each applicant's ranking of programs, application scores, offers received, and enrollment decisions, as well as the final-round admission cutoff (if any) for each program each year. 1998 is the first cohort with application data available, and stopping at 2004 allows us to study outcomes in a balanced panel spanning 13 years after application.[14] We retain each applicant's first observed application. We also require that they – at the time of application – have no post-secondary degree already, are younger than 27 years, and are not already matched with a spouse or registered partner. We drop applicants who have an application score further than 2 standard deviations away from the admission threshold of their locally preferred option, and a small number who have missing information on completed education 13 years after applying.

Our main estimation sample comprises 110,382 unique applicants who apply to at least two programs where the preferred program has a binding admission cutoff and the next-best alternative has a lower cutoff (or none). This ensures that we have information on the preferred program and a source of exogenous variation into it. When estimating effects of field, institution, and elite education, which are groups of programs, we construct analogous estimation samples of applicants who face a binding cutoff into into their preferred option (e.g. a given institution) and a less selective next-best alternative that is outside of that preferred option. These estimation samples comprise 98,858 applicants on the margin between different fields, 96,049 between different institutions, and 22,545 between elite vs. non-elite education.

Table 3 summarizes our main estimation sample. Applicants are, on average,

[14]The panel is nearly perfectly balanced. Only a small number of individuals (about 2 percent) drop out at some point during the 13 year period, mainly due to emigration.

Table 3. Estimation sample summary statistics

	Sample mean (SD)		
Pre-determined characteristics:			
- Age	20.7 (1.8)		
- Female	0.62		
- Immigrant	0.04		
- College-educated parent	0.50		
Applications:			
- Number of programs ranked	6.6 (4.2)		
- Number of distinct fields ranked	3.5 (2.1)		
- Number of distinct institutions ranked	3.9 (2.8)		
- Applied to an elite program	0.20		
- Rank of locally preferred program	1.8 (1.6)		
Offers:			
- Rank of best offer	2.8 (3.1)		
- Offered 1st rank	0.41		
- Offered 2nd rank	0.20		
- Offered 3rd rank	0.10		
- No immediate offer	0.20		
Education:	Enroll immediately	Ever enroll	Complete degree
- Any college	0.77	0.97	0.83
- Preferred program	0.32	0.44	0.30
- Preferred field	0.36	0.55	0.39
- Preferred institution	0.41	0.56	0.41
- Any elite program	0.05	0.16	0.11
Matching:			
- Any partner	0.81		
- College-educated partner	0.47		

Note: This table presents summary statistics of the main estimation sample, consisting of 110,382 applicants. Application outcomes correspond to the first observed application. Ever enroll, complete degree, and matching are measured 13 years out from the first observed application.

between 20 and 21 years old when we observe them applying for the first time.[15] 62 percent are female, 4 percent are immigrants, and half have at least one college-educated parent. Applicants list about 7 programs on average, across 3–4 different fields and institutions, and 20 percent apply to an elite program. The locally preferred option around the applicant's application score typically corresponds to their 2nd ranked program. 41 percent of applicants are offered their first-ranked program, with the average offer corresponding to the 3rd ranked program. 20 percent of applicants receive no offer in their first application attempt, but nearly all eventually enroll somewhere, as shown in the next panel.

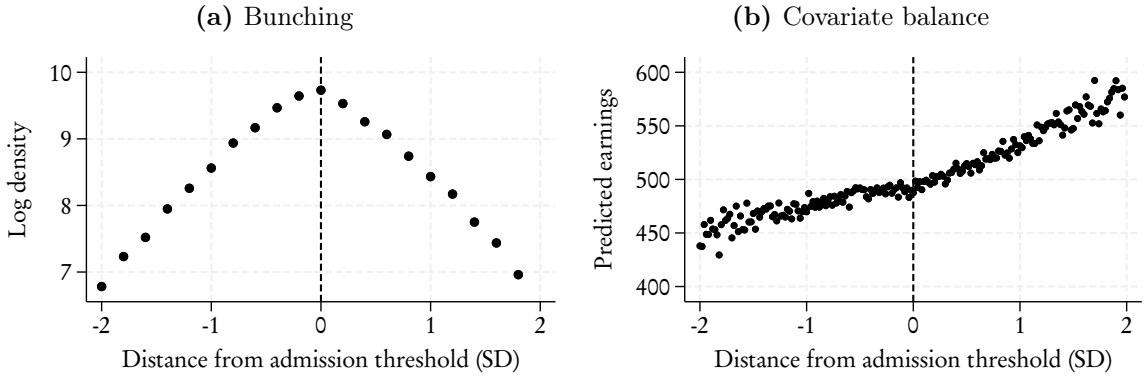
77 percent of applicants enroll in college immediately following their first application attempt. 32 percent immediately enroll in their locally preferred program around their application score. 36 percent enroll in their locally preferred field, and 41 percent in their locally preferred institution. 5 percent immediately enroll in any elite program. 13 years later, 97 percent of applicants have enrolled in some higher education; 44 percent in their preferred program from the initial application, about 55 percent in their initially preferred field or institution, and 16 percent in any elite program. 83 percent of all applicants eventually earn a college degree; conditional on ever enrolling in a preferred option, around 70 percent complete that specific option. 81 percent of applicants match with a partner within 13 years after applying, and 47 percent have a college-educated partner.

4.3 Assessing the validity of the regression discontinuity design

Figure 4 probes the validity of the research design. If applicants are able to sort themselves just above the relevant cutoff to gain admission into their preferred option, we would expect to see discontinuities in the density of applicants and in their observed characteristics at the cutoffs. To investigate this, we pool across programs in our estimation sample, normalizing the running variable so that zero on the horizontal axis represents the admission cutoff to the preferred program, and observations to the left (right) of zero have application scores that are below (above) the cutoff.

The left panel of Figure 4 plots the density of applications, showing no indication that applicants are able to strategically position themselves just above the cutoff. The right panel investigates covariate balance. We construct a composite index of pre-

[15] In Norway, students graduate from high school in the year they turn 19, after which many serve in the military, travel, or work for a year or two before enrolling in college.



Note: Panel (a) plots the log density of applications around the admission threshold. Panel (b) plots an index of pre-determined covariates, constructed as predicted earnings from a regression of earnings (measured 13 years after application) on age at application, gender, immigrant status, mother and father education levels, raw high school GPA, home county, and application cohort.

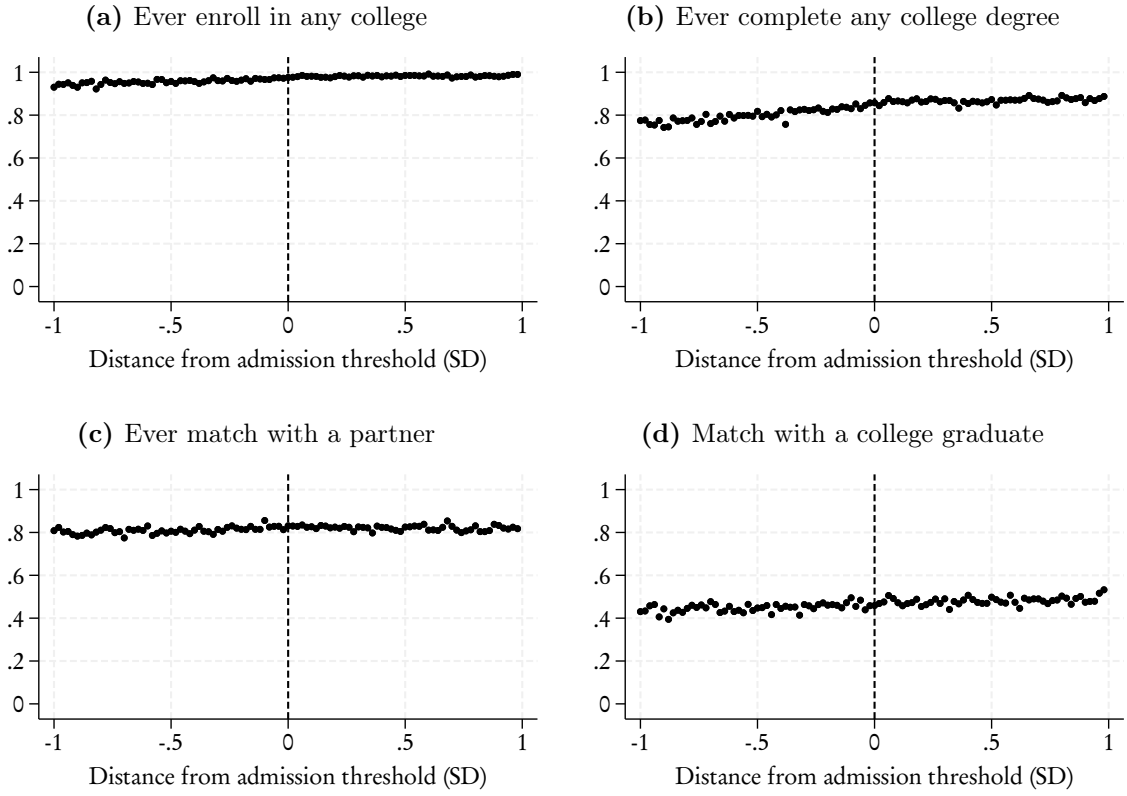
Figure 4. Assessing the validity of the regression discontinuity design

determined applicant characteristics – age at application, gender, immigrant status, mother and father education levels, raw high school GPA, home county, and application cohort – by regressing earnings 13 years after application on these variables and plotting the predicted values from the regression. The right panel of Figure 4 shows no indication that applicants just above versus just below the admission threshold are observably different from each other.

4.4 What does and does not change across the admission thresholds

Since applicants do not appear to sort around the admission thresholds, we can infer causal effects by examining how outcomes change across the thresholds. Before getting to our instrumental variables model and main results, we first establish what kinds of outcomes do and do not change across the thresholds. Figure 5 shows that crossing the threshold into a preferred program does not meaningfully affect outcomes along the "extensive margins" of education and matching. We see no substantial changes in the probability of ever enrolling in college (panel a), ever completing a degree (panel b), matching with any partner (panel c), or matching with any college-educated partner (panel d).

What does change is the *type* of education, and partner, acquired. Figure 6 considers education type and confirms that crossing the admission threshold of a preferred program substantially increases the probability of getting an offer from that program

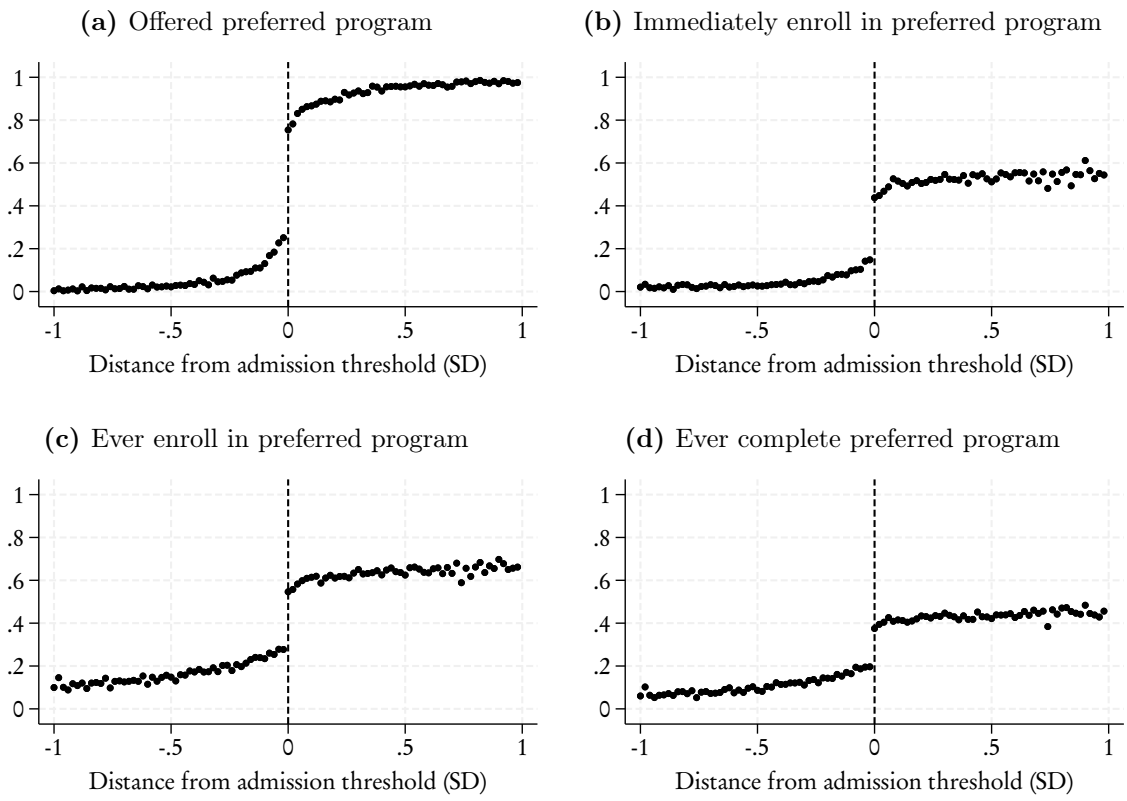


Note: Each outcome is measured cumulatively as of 13 years after the initial application.

Figure 5. What does not change across the admission thresholds: extensive margins of education and matching

(panel a). This increases the probability of enrolling in that program, with a similarly-sized jump regardless of whether enrollment is measured immediately following the initial application (panel b) or ever within 13 years of the initial application (panel c).[16] Not everyone who starts the preferred program completes it, but a majority do,

[16] Admission offers and enrollment compliance are fuzzy rather than deterministic through the cutoff for several reasons. First, the admission process proceeds in multiple rounds; our data correspond to the final round of centralized offers. Some applicants who score above the second-round cutoff but below the first-round cutoff do not choose to be on the waiting list (actively, or passively by not responding) and will therefore not get an offer. On the other side of the threshold, if slots remain at the end of the centralized admission process, some programs admit additional below-threshold applicants through decentralized local rounds. Second, institutions have some limited discretion when evaluating applicants with special circumstances, like illness or disability. Third, enrollments are measured 1-2 months into the academic year, by which time some applicants who accepted an offer have already dropped out. Finally, the centralized application system makes the cost of applying very small. Together with the absence of tuition fees, this likely attracts a non-negligible number of applicants who are only weakly attached to their locally preferred option.

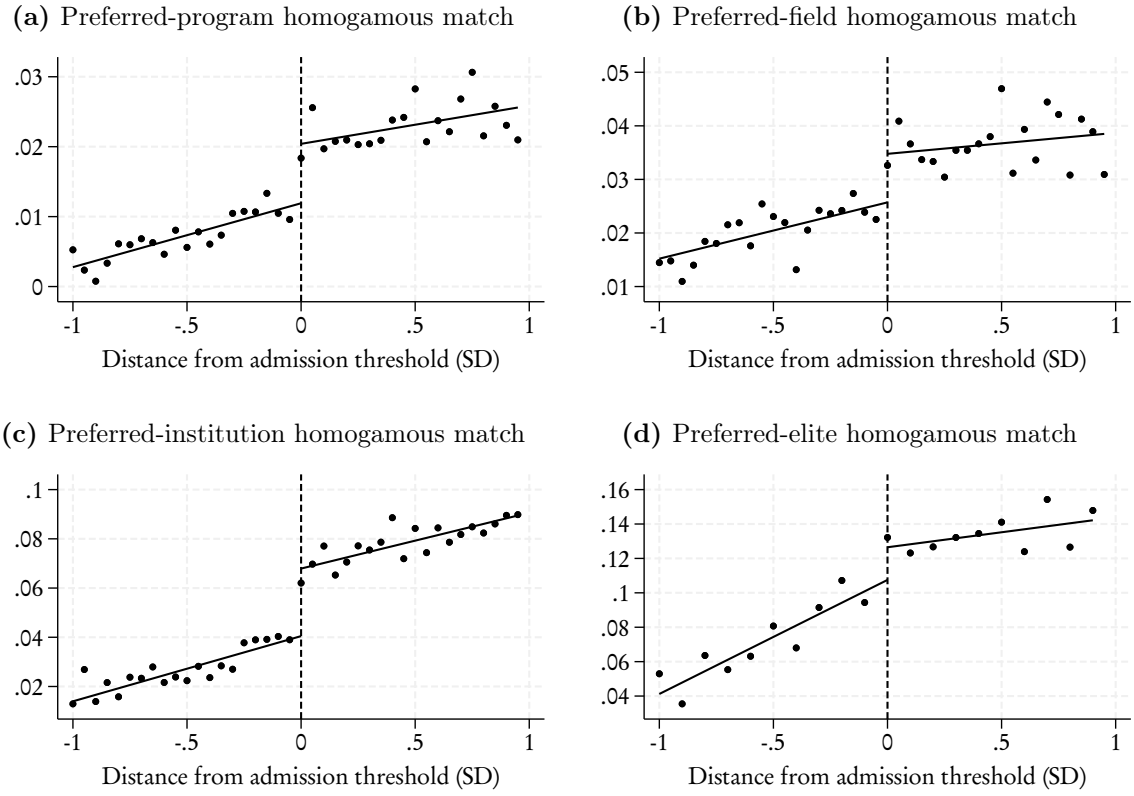


Note: Admission offers in panel (a) are measured in the applicant's first observed application cycle. Immediate enrollment in panel (b) is defined as enrolling in the academic year immediately following the initial application. The outcomes in the remaining panels are all measured cumulatively as of 13 years after the initial application.

Figure 6. What does change across the admission thresholds: education type

leading to a substantial effect of threshold crossing on the probability of ever earning a degree from the preferred program (panel d). Appendix Figures A2, A3, and A4 show analogous discontinuities when crossing the admission threshold into a preferred field, preferred institution, and any elite program, respectively.

Figure 7 shows the reduced-form effects of threshold crossing on homogamous matching. Panel (a) shows that crossing the admission threshold into the applicant's preferred program increases the probability of a homogamous match with respect to that program, i.e. both the applicant and the partner share degrees from that program. Panels (b) and (c) show analogous increases in homogamy with respect to the applicant's preferred field and institution, respectively, and panel (d) shows an increase in the probability of becoming a dual-elite couple as a result of crossing the



Note: A preferred-option homogamous match means that the applicant and their partner completed degrees from the applicant's locally preferred option.

Figure 7. What does change across the admission thresholds: partner type

admission threshold into an elite program.

At first glance, the magnitudes in Figure 7 may seem small, but these are simply the reduced-form effects of threshold crossing among applicants who do not necessarily end up enrolling in the preferred option, or completing a degree from it, or matching with a college-graduate partner. To make these magnitudes comparable to the population homogamy rates in Section 3, which conditioned on college-graduate couples, we calculate that 20 percent of applicants on the program margin who fall just above the admission threshold end up completing the preferred program and matching with any college graduate partner; thus, the preferred-program homogamous match rate of 2 percent just above the threshold in panel (a) implies a $2/20 = 10\%$ program homogamy rate among the marginally admitted, comparable to the 12.3% population rate in Table 1. Analogous calculations imply a $3.5/23.2 = 15\%$ field homogamy rate among the marginally admitted, comparable to 16.5% in the

population; a 6.7/24.9 = 27.1% institution homogamy rate among the marginally admitted, comparable to 34.3% in the population; and a 12.5/37.9 = 33% elite homogamy rate among marginally admitted elite applicants. Only 1 in 5 applicants are on the elite margin; 33/5 = 6.6% is therefore comparable to the 4.9% elite homogamy rate in the population of college graduates in Table 1.

4.5 Instrumental variables model

The discontinuities that arise from the college admission process allow us to identify the intention-to-treat effects of crossing the admission threshold into a preferred option. To estimate effects of actually enrolling in the preferred option, we employ a fuzzy regression discontinuity design that uses threshold crossing as an instrument for an applicant enrolling in their locally preferred option. We estimate the following first stage equation:

$$d_i = \pi z_i + x_i'\gamma + u_i \tag{4}$$

where the dependent variable d_i equals 1 if the applicant ever enrolls in the preferred option. The instrument z_i is the predicted offer for the preferred option, equal to one if the individual's application score exceeds the relevant admission cutoff (and zero otherwise). $x_i'\gamma$ denotes controls, including locally linear functions of the running variable (application score) that are allowed to differ arbitrarily on each side of the threshold, using observations within a two-standard-deviation bandwidth. To reduce residual variance, we also include a set of pre-determined controls for gender, application year, and preferred program. Since 15–20 percent of applicants are observed at two margins, we cluster the standard errors at the applicant level.

The corresponding second stage equation is

$$y_i = \delta d_i + x_i'\beta + e_i \tag{5}$$

where y_i is an outcome of interest of individual i. One target of our estimation is the average of δ among the compliers: applicants who enroll in their preferred option because their application score fell just above the admission cutoff and would not have enrolled otherwise. We use 2SLS with first and second stage equations given by (4) and (5) to estimate δ.

To decompose homogamy into selection versus treatment, we estimate the complier average potential outcomes with and without treatment. Here we follow Abadie

(2003), who shows that with a binary treatment d, binary instrument z, and scalar outcome y, the compliers' mean potential outcome with treatment, y^1, is identified by a 2SLS regression of $d \cdot y$ on d instrumented with z. Similarly, the compliers' mean potential outcome without treatment, y^0, is identified by a 2SLS regression of $(1-d)\cdot y$ on $(1-d)$ instrumented with z. We explicitly measure selection into preferred options as the difference between the mean y^0 among compliers, who preferred their local option but were exogenously denied access from it, versus the random matching benchmark y^R, which is the probability that a complier would randomly match with a partner who completed their preferred option. Identifying y^R is analogous to identifying y^0, but replaces applicant i's actual match outcome y_i with the share of all applicants' partners (separately by applicant gender, and including no partners) who completed applicant i's preferred option.

To identify our parameters of interest, we make three assumptions. The first is that applicants are not able to perfectly sort themselves around the admission threshold in order to receive an offer from their preferred option. As shown above, the data support this assumption; it is also consistent with key features of the admission process. First, the exact admission cutoffs are unknown to applicants when they take their high school exams and when they submit their application. Second, the cutoffs vary considerably over time, in part because of changes in demand, but also because of funding changes that affect the supply of slots. Third, there is limited scope for students to finely manipulate their application scores, as they depend on academic performance over all three years of high school, unlike admission systems based on final-year exams or college entrance tests. The second identifying assumption is that crossing the threshold makes an applicant weakly more likely to enroll in the preferred option. This monotonicity assumption seems plausible in our setting.

The third assumption is that threshold crossing affects the outcomes of interest exclusively through the treatment variable. To evaluate this, note that we specify d_i as an indicator for whether an applicant *enrolls* in the preferred option, regardless of whether they complete it. This specification alleviates concerns about exclusion violations that could arise if one instead specified *completing* the preferred option as the treatment variable. Individuals who enroll in a given option may be more likely to meet and match with other people in that option even if they end up not completing it, whereas simply getting admitted into an option is unlikely to affect outcomes independently of actually enrolling in it.

In Section 5.2 we present results from several specification checks, including successively smaller bandwidths of contributing observations around the cutoff, all of which support our main findings.

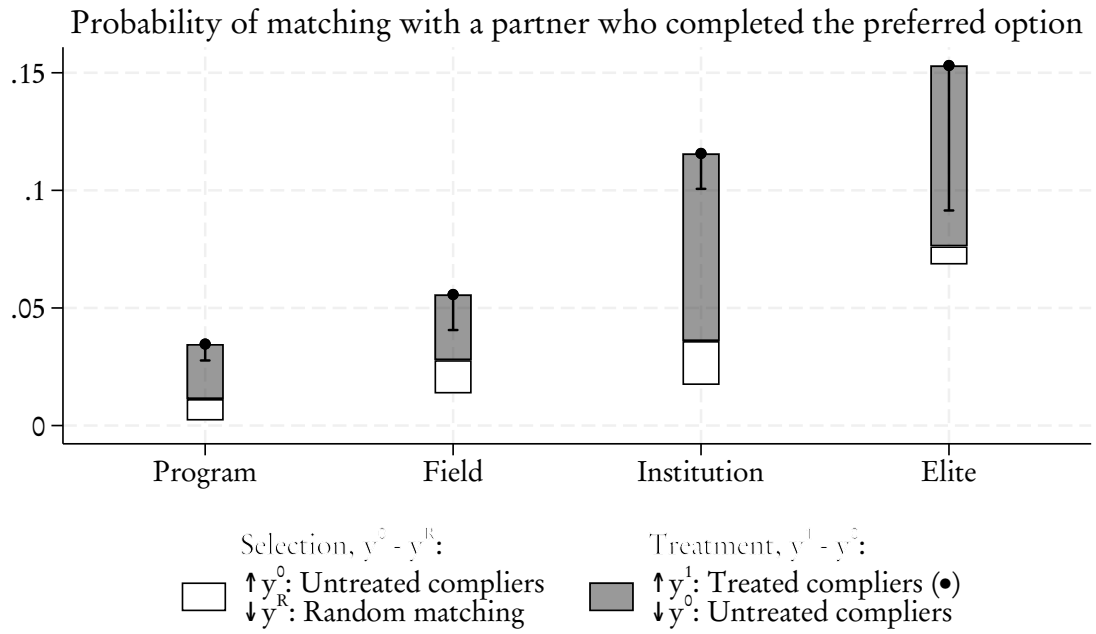
5 Main results

5.1 Decomposition of college homogamy into selection versus treatment

We now have all the tools in place to decompose the strong patterns of college homogamy documented in Section 3 into causal college match-making versus selection. Figure 8 presents the results. Each decomposition involves comparing three probabilities, where the outcome of interest is matching with a partner who completed the applicant's locally preferred option around their application score. y^1 is the probability of such a match among treated compliers, i.e. applicants who actually enroll in their preferred option as a result of narrowly crossing its admission threshold. y^0 is the corresponding probability among untreated compliers, who are identical to the treated compliers but for their exogenous exclusion from the preferred option. $y^1 - y^0$ is thus causal college match-making: the average treatment effect of actually enrolling a complier in the preferred option on the probability that they match with a partner who completed that option. This treatment component is visualized by the shaded vertical bar in Figure 8. The figure also shows the 95 percent confidence half-interval of the treatment effect extended below y^1. Comparing the lower bound of the CI to y^0 is equivalent to comparing the lower-bound of the treatment effect CI to zero, and therefore informative about the statistical precision of the treatment component.

To quantify selection, visualized by the hollow vertical bar, we compare y^0 to the random matching benchmark y^R: if compliers matched with partners randomly, how often would they match with someone who completed their locally preferred option? The difference $y^0 - y^R$ therefore reveals any inevitable matching proclivities among those who self-select into the preferred option, independent of actually getting treated by it.

The results in Figure 8 reveal a dominant role for causal college match-making over selection in the production of homogamous couples. Treated compliers who prefer a given option become much more likely to match with a partner who completed that option, as a direct result of actually enrolling in that option. Untreated compliers, who identically prefer the same option but are exogenously denied access from it, are



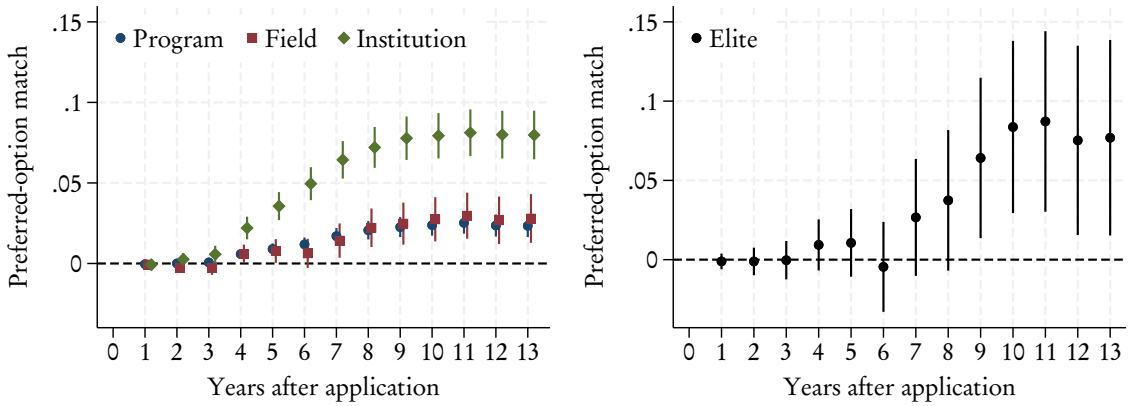
Note: This figure decomposes the probability of matching with a preferred-option partner into selection ($y^0 - y^R$) versus treatment ($y^1 - y^0$). The 95 percent confidence half-interval of the treatment effect is extended below y^1. Comparing this lower bound to y^0 is equivalent to comparing the lower-bound of the treatment effect CI to zero, and therefore informative about the statistical precision of the treatment component.

Figure 8. Decomposition of college homogamy into selection versus treatment

only slightly more likely than random to match with a partner from that option. This pattern holds true even for elite applicants. Measuring total homogamy as the treated complier mean minus the random matching benchmark $y^1 - y^R$, and the causal share as $\frac{y^1 - y^0}{y^1 - y^R}$, the decomposition implies that program homogamy is 74 percent causal; field homogamy is 70 percent causal; institution homogamy is 85 percent causal; and elite homogamy is 92 percent causal. Thus, the matches we observe among the college-educated are not inevitable, even among the elite; rather, they are a direct consequence of the education itself.

5.2 Specification checks

Before moving on, we probe the robustness of these results with a range of specification checks. In Figure 9, we first examine whether estimating effects on match outcomes 13 years after the initial application allows sufficient time for matching dynamics to

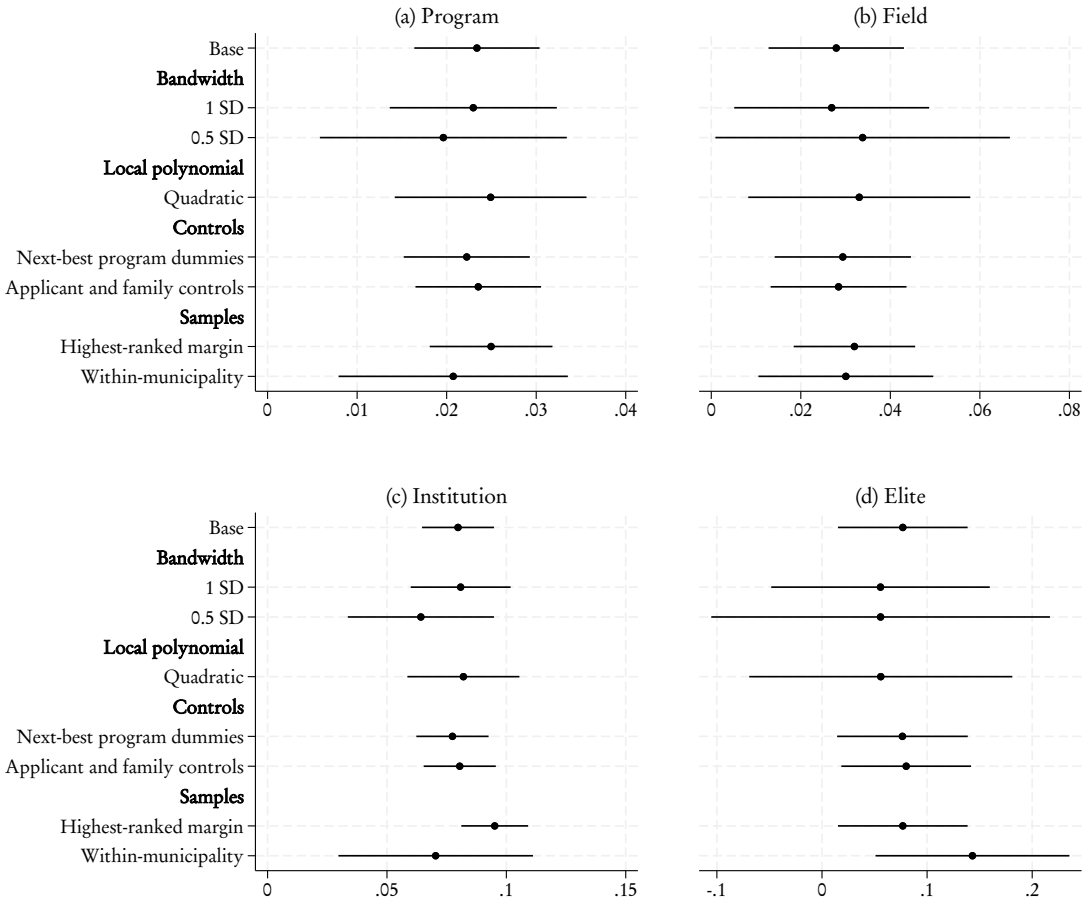


Note: This figure plots 2SLS effects of enrolling in the preferred program/field/institution (left panel) or enrolling in elite education (right) on the probability of matching with a partner who completed that option. Each estimate comes from a separate regression using the main specification described in Section 4.5.

Figure 9. Dynamic effects of enrollment on matching with a preferred-option partner

play out. Indeed, the estimated treatment effects naturally increase over the first decade but have plateaued by the time of our measurement at year 13.

Next, we report treatment effect estimates across a range of alternative specifications. In each panel of Figure 10, the top line reproduces our baseline estimate, corresponding to the treatment component $y^1 - y^0$ in Figure 8. In the first two specification checks, we show that our estimates are robust to the choice of bandwidth used in the locally linear regression discontinuity design. Narrowing the bandwidth from two standard deviations of the running variable to one standard deviation, and further to just half a standard deviation, yields less precision but similar point estimates to our baseline. The next alternative specification replaces the baseline's local linear slopes in the running variable with quadratic functions of the running variable that vary arbitrarily on each side of the cutoff. The next two specification checks investigate robustness to additional covariate controls, including indicator variables for each applicant's next-best program, as well as additional pre-determined measures of applicant and family background: indicators for age at application, municipality of residence at age 16, parental immigration, and parental education. As some applicants in our sample are observed at two margins (see footnote 13), we also include a specification check where we only keep the most-preferred margin of each applicant.



Note: Each figure shows 2SLS estimates of the effects of enrolling in the preferred option on the probability of matching with a partner who completed that option. Each estimate comes from a separate regression using the labeled specification, with details of each described in the main text. Error bars indicate 95% CI (standard errors clustered at the applicant level).

Figure 10. Specification checks

The estimates do not materially change across these alternative specifications.

Finally, to investigate how much of the college marriage-market effects are related to geography, the last row of estimates in Figure 10 restricts the sample to applicants who are on the margin between options within the same municipality. This ensures that the reported effects only capture the consequences of within-municipality treatment assignments. The results show only a modest attenuation of the institution effect and a modest increase in the elite effect, indicating that our results are not driven by geography alone.

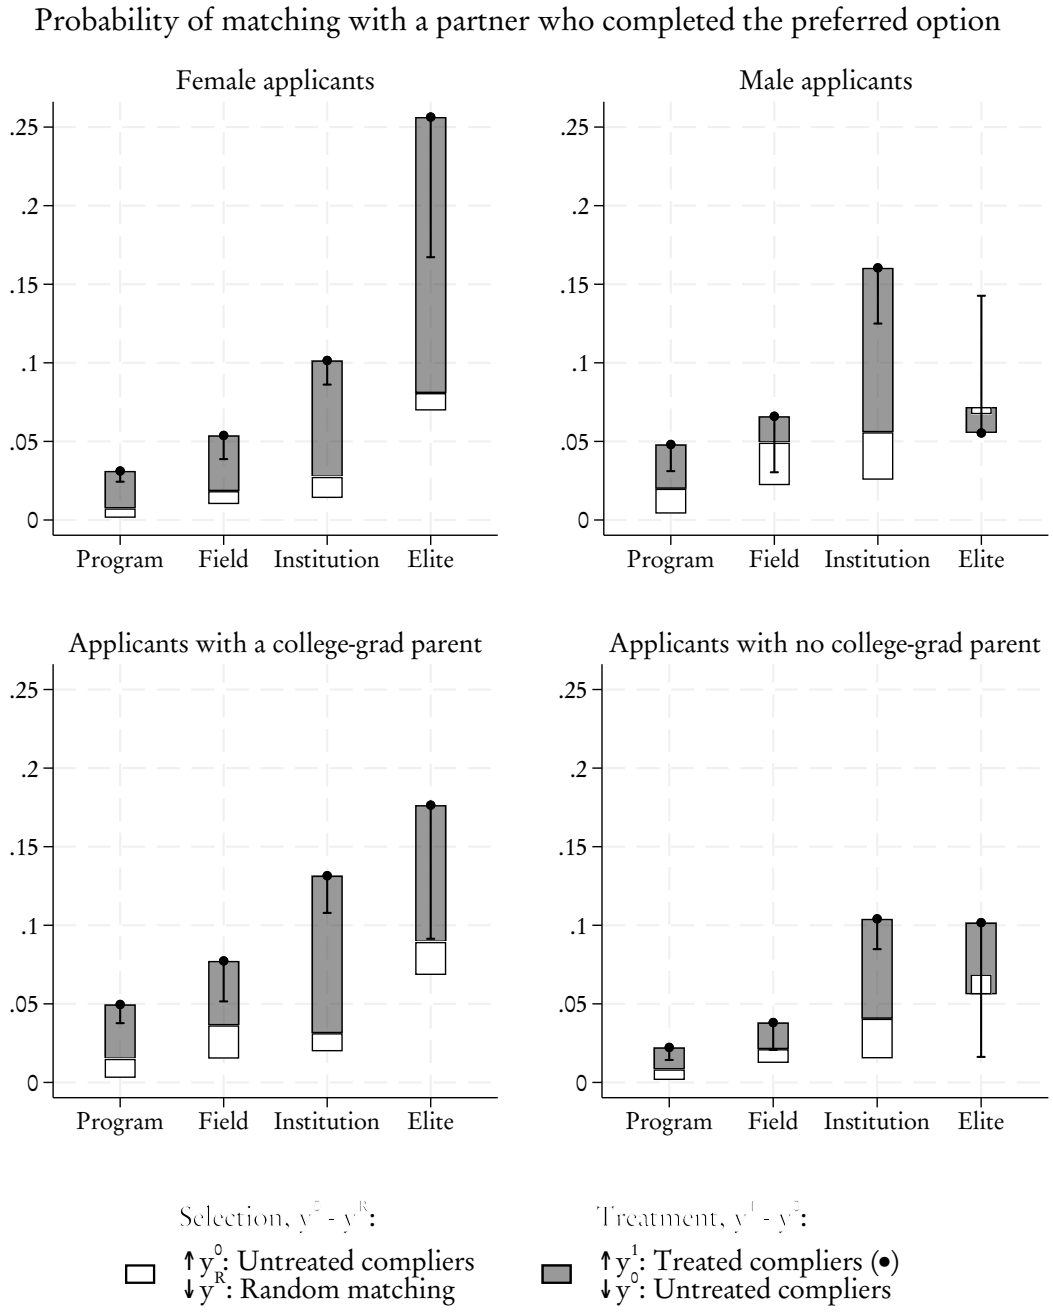
5.3 Heterogeneity by gender and family background

The top panels of Figure 11 examine heterogeneity in selection and treatment across female versus male applicants. For women, along all margins, college homogamy is almost entirely causal: treated female compliers become substantially more likely to match with a partner who completed the preferred option, while untreated compliers do so at rates nearly equivalent to random matching. Treatment effects for women are especially large on the elite margin, with enrollment causing an 18 percentage point increase in the probability of matching with an elite partner. Male elite applicants, in stark contrast, have roughly the same probability of matching with an elite partner regardless of whether they enroll in elite education or not, and that probability is also roughly equivalent to what would be expected if they matched with partners randomly.

The bottom panels of Figure 11 examine heterogeneity by parental education. The effects of enrollment on matching are consistently larger for applicants with at least one college-educated parent. Untreated elite applicants with no college-educated parent match with elite partners at a rate slightly lower than random matching, and they experience a smaller increase in that match rate when treated by elite enrollment compared to their peers from more educated families.

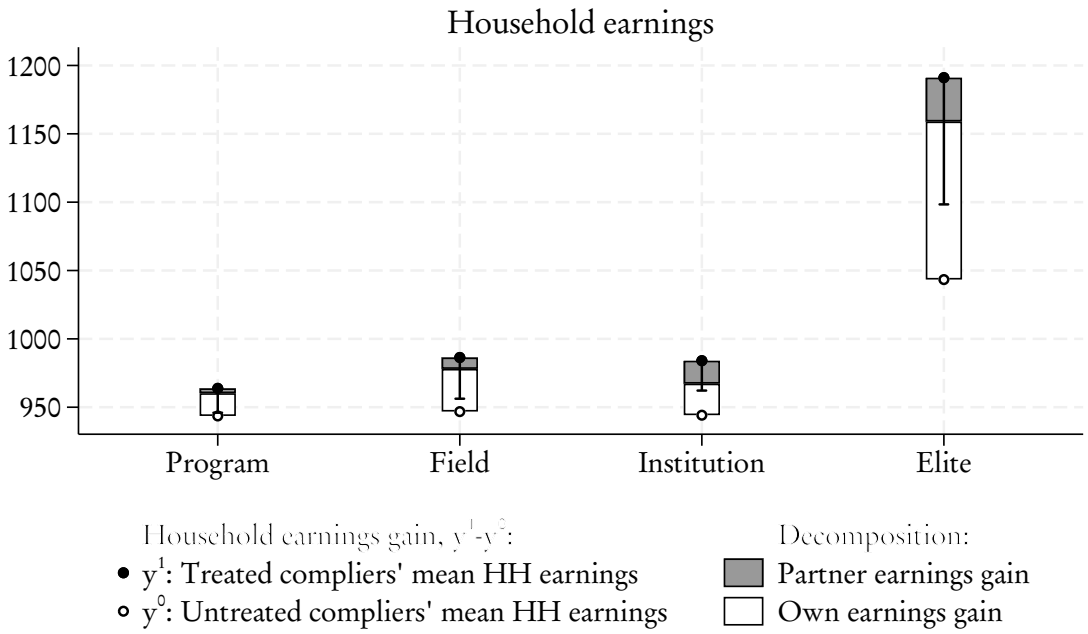
5.4 The role of higher education in the production of elite households

The preceding results have shown that college enrollment itself, rather than selection into it, primarily drives homogamous matching among the college-educated, with especially large effects of elite enrollment for women and applicants from more educated families. But are these college match-making effects economically consequential? We



Note: This figure conducts the same exercise as Figure 8, but separately by gender (top panels) and parental education (bottom panels). See the notes to Figure 8.

Figure 11. Decomposition of homogamy into selection vs. treatment, by applicant gender and parental education

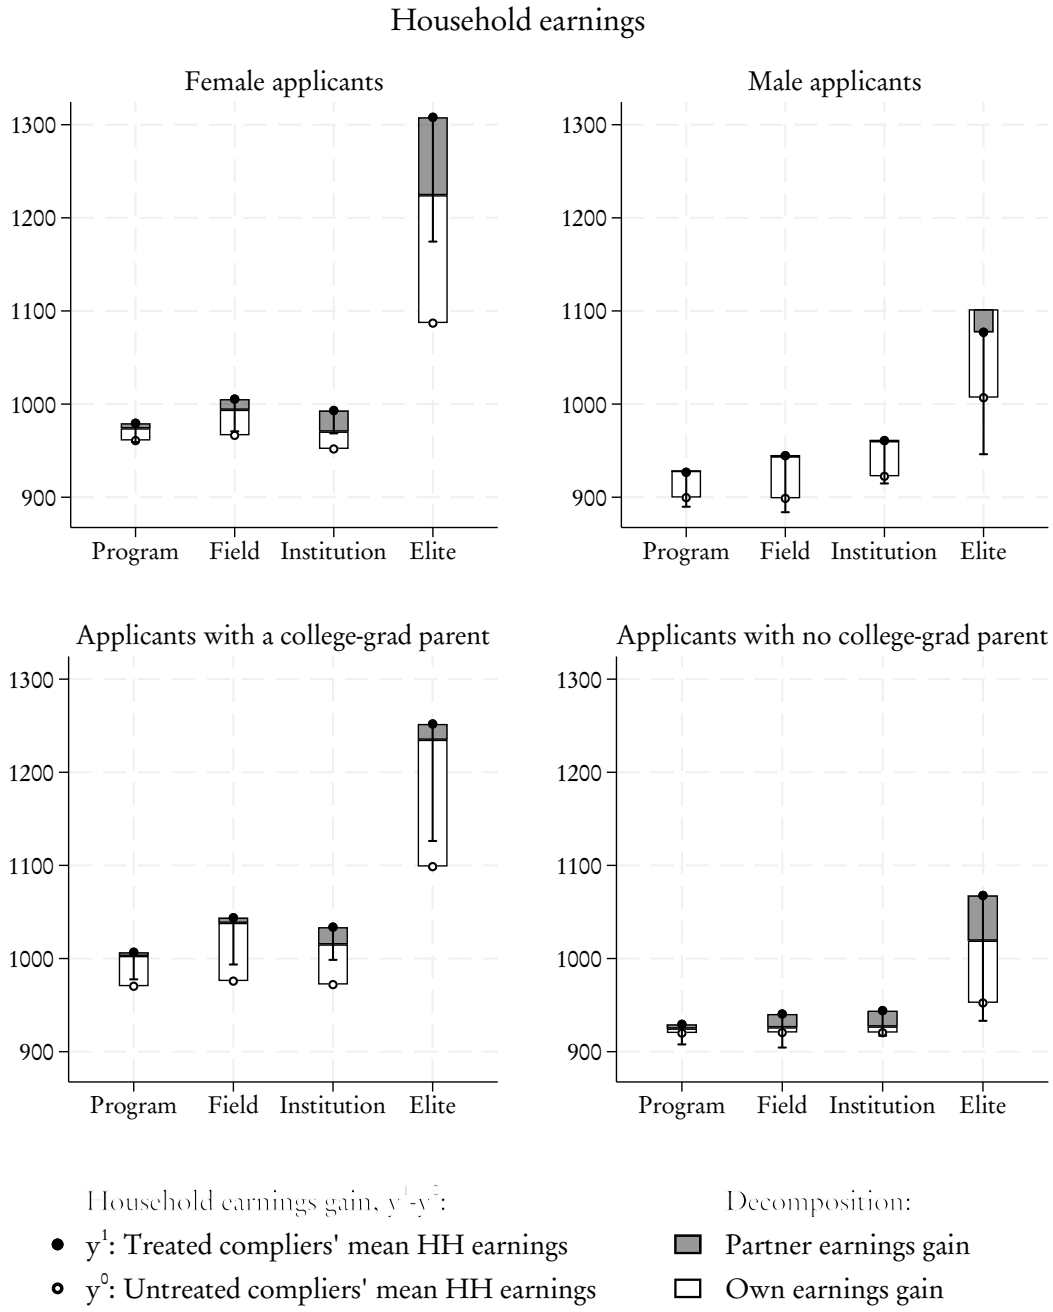


Note: Household earnings are the sum of the applicant's own earnings and their partner's earnings, including zeros and no partners. Earnings are measured 13 years after application in thousands of Norwegian kroner. The 95 percent confidence half-interval of the treatment effect on total household earnings is extended below y^1. Comparing this lower bound to y^0 is equivalent to comparing the lower-bound of the treatment effect CI to zero, and therefore informative about the statistical precision of the treatment effect on total household earnings $y^1 - y^0$.

Figure 12. Household earnings effects of enrolling in the preferred option

use our main specification to estimate effects of enrolling in the preferred option on the applicant's household earnings, measured 13 years after the initial application. We also decompose the total effect on household earnings into the effect on the applicant's own earnings versus the effect on matching with a higher (or lower) earning partner (including no partner).

Figure 12 presents the results, with Figure 13 breaking them down by applicant gender and parental education. On average, enrolling in the preferred option increases household earnings, modestly so along the program, field, and institution margins, and substantially so along the elite margin. This is mostly due to gains in the applicant's own earnings; on average, the smaller partner earnings effects imply that applicants match with only modestly higher-earning partners than they would have if rejected from the preferred option.



Note: This figure conducts the same exercise as Figure 12, but separately by gender (top panels) and parental education (bottom panels). See the notes to Figure 12.

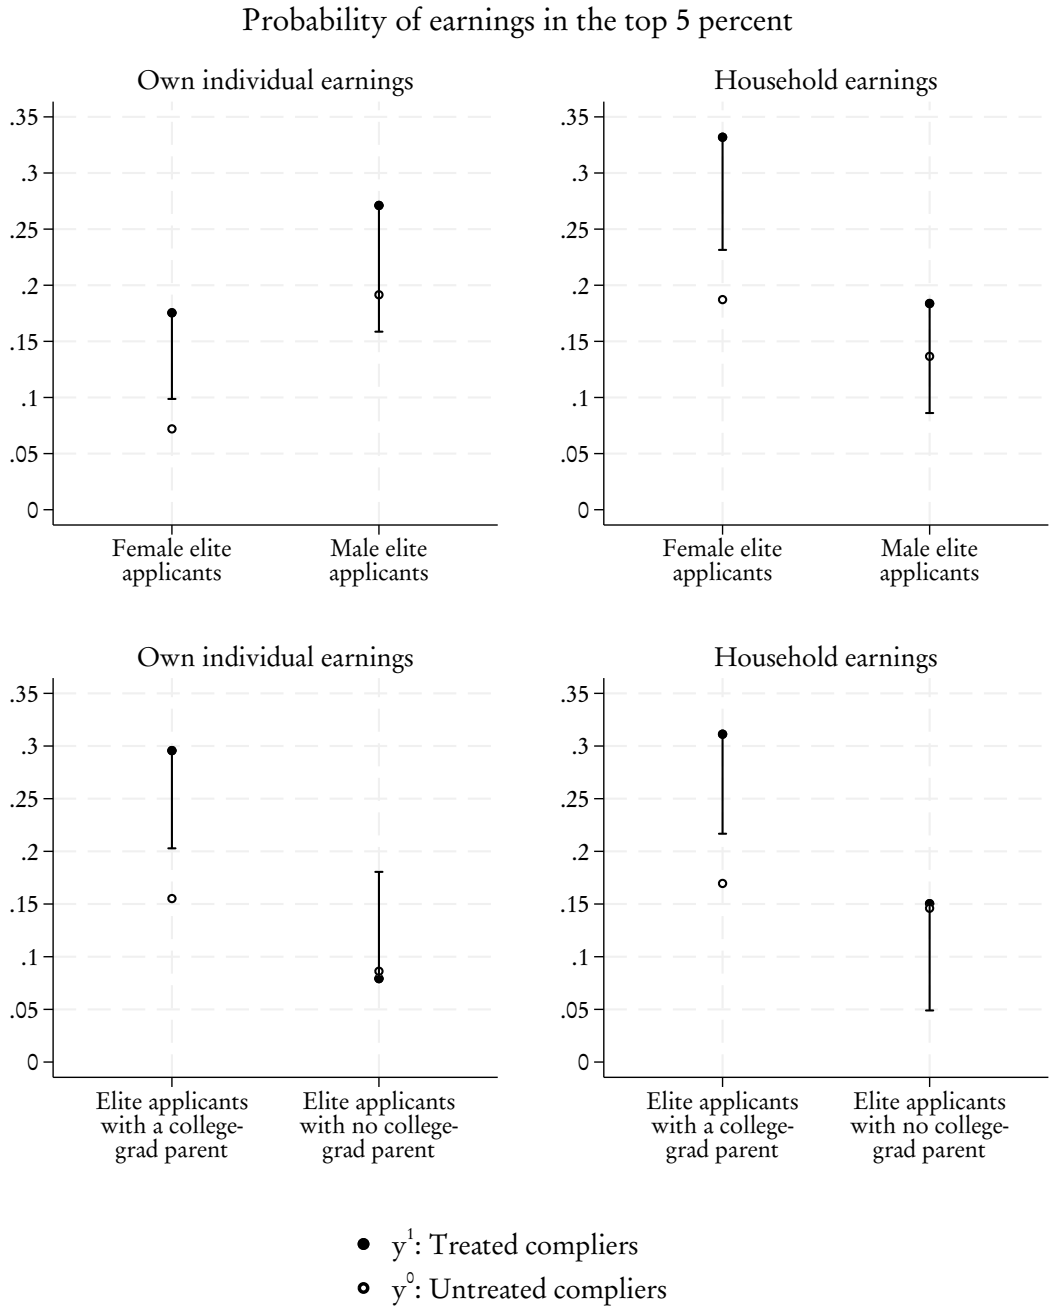
Figure 13. Household earnings effects by gender and parental education

The top panels of Figure 13, however, shows divergent patterns for female versus male applicants. Along all margins, female applicants augment their earnings gains by matching with higher-earning partners, while male applicants' household earnings gains are entirely driven by their own earnings gains. Elite enrollment, in particular, propels marginally admitted women into elite household formation: they earn substantially more themselves and match with substantially higher-earning partners. Marginally admitted elite men, on the other hand, experience a smaller, statistically insignificant increase in household earnings, with no apparent partner earnings contribution (if anything, the point estimate on partner earnings is slightly negative).

The bottom panels of Figure 13 also reveals divergent patterns by parental education. Applicants with a college-educated parent experience larger gains in household income along all margins, almost entirely driven by larger gains in their own earnings. For applicants with no college-educated parent, enrollment in their locally preferred program, field, and institution generally yields small earnings gains that are indistinguishable from zero. The estimated gain from elite enrollment is meaningfully positive, though statistically imprecise, with a relatively larger share of it coming from matching with a higher-earning partner compared to applicants with a college-educated parent.

Figure 14 further examines the role of elite education in the production of elite households by plotting treatment effects of elite enrollment on the probability of having earnings in the top 5 percent of the applicant's entire birth cohort, measured separately in the individual versus household earnings distributions. If we only considered individual earnings in the top left panel, we would conclude that both women and men increase their chances of joining the top earnings ranks (with a statistically imprecise estimate for men) by enrolling in elite education, but with a persistent male-dominated gender gap. Female elite applicants boost their chance of top-5-percent individual earnings from 8 percent to 18 percent by enrolling, but that still leaves them with lower chance than *untreated* male applicants, who have a 20 percent chance of top-5-percent earnings that increases to 28 percent by enrolling.

However, when we consider household earnings in the top right panel of Figure 14, these patterns completely reverse. Treated elite male applicants rise to the level of untreated female applicants at a roughly 18 percent probability of joining the top 5 percent of household income, while fully one third of female applicants who are treated by elite enrollment form households in the top 5 percent.



Note: Earnings are measured 13 years after application. The probability of earnings in the top 5 percent is measured separately for the own (individual) earnings distribution and the household earnings distribution, both of which are conditional on birth cohort and age at measurement. The 95 percent confidence half-interval of the treatment effect is extended below y^1. Comparing this lower bound to y^0 is equivalent to comparing the lower-bound of the treatment effect CI to zero, and therefore informative about the statistical precision of the treatment effect $y^1 - y^0$.

Figure 14. The production of top incomes, by gender and parental education

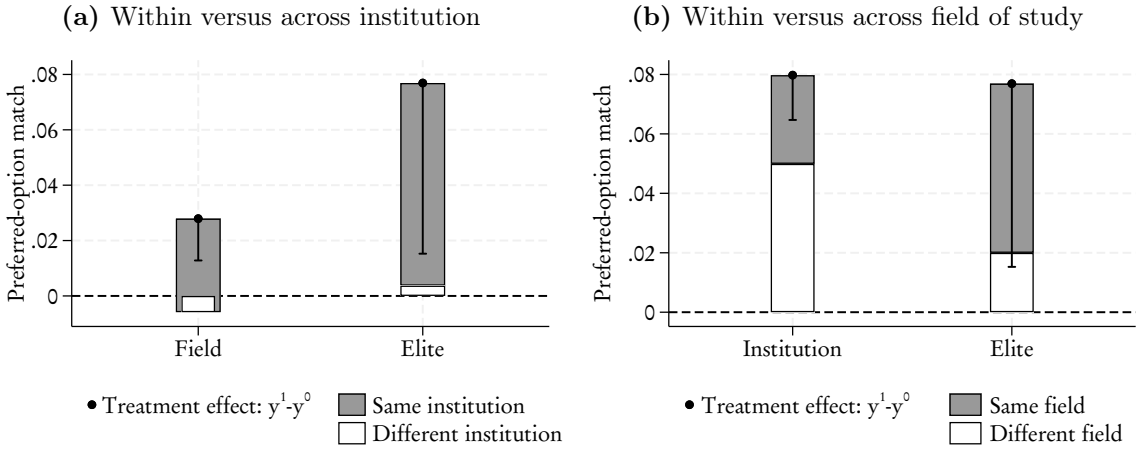
The bottom panels of Figure 14 repeat this exercise by parental education. With respect to both own and household earnings, elite education amplifies pre-existing inequality. For applicants with a college-educated parent, elite enrollment substantially increases their likelihood of joining the top earnings ranks, while applicants with no college-educated parent see no such change. Thus, even if elite enrollment does yield some earnings gains for applicants with no college-educated parent (Figure 13), those gains are not enough to propel them into the top ranks of the earnings distribution either at the individual or household level, in contrast to their peers from more educated families.

6 Mechanisms

We have documented that college education itself, rather than selection into it, drives matching patterns among the college-educated, and that these matches can be economically consequential. Multiple mechanisms could underlie these results. One is that enrolling in a particular option changes an individual's type in the matching market, including their demand for partner characteristics and their supply of own characteristics. Another mechanism is that the matching effects are a direct consequence of the meeting opportunities generated by enrolling at a particular time and place. We now examine the mechanisms through which enrollment affects matching, and implications for the functioning of the marriage market.

6.1 What are the key mechanisms of the enrollment treatment?

The results above showed that the treatment effect of enrollment can account for most of the educational homogamy we observe among the college-educated. But what are the key mechanisms through which enrollment affects matching? If enrolling induces type changes in a matching market where the cost of finding partners is low, then we would expect some of the enrollment effects to be driven by matches across different institutions and cohorts, since the particular location and timing of the education should not matter much for matching on the new type. To help understand the roles of institutions, fields, cohorts, and employers, we now perform a series of decompositions to determine the extent to which college match-making effects are driven within versus across these potential mechanisms.



Note: Figure (a) decomposes the 2SLS effect of enrolling in the preferred option on the probability of matching with a partner who completed the preferred option into within-institution versus across-institution matches. Figure (b) similarly decomposes the effect of enrolling into within-field versus across-field matches. Error bars indicate 95% CI (standard errors clustered at the applicant level).

Figure 15. Match-making mechanisms: within versus across institution and field

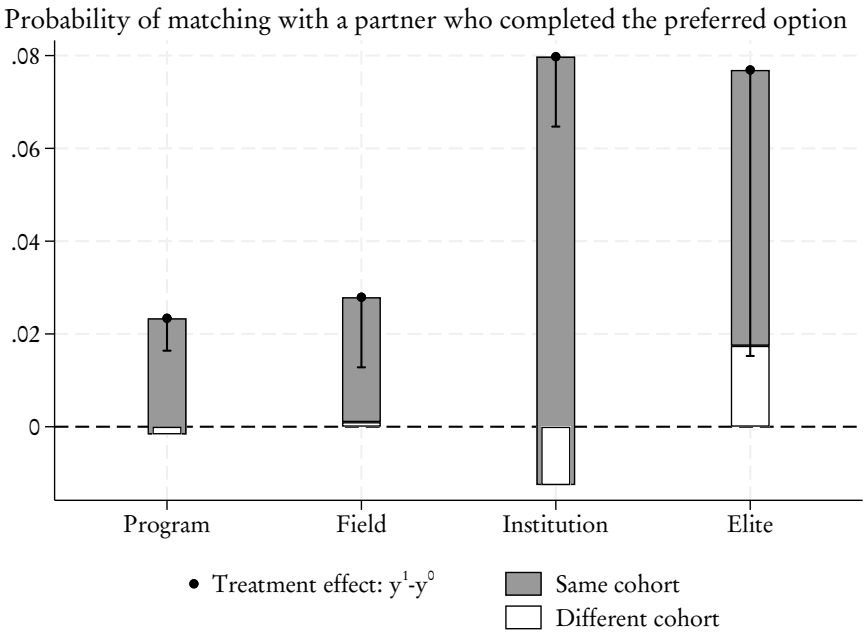
Within versus across institution and field Matches with a partner who completed the applicant's preferred option (M) are either within the same institution ($M \times I$) or across different institutions ($M \times !I$),

$$\Pr(M) = \Pr(M \times I) + \Pr(M \times !I).$$

By estimating the effects of enrollment on each component of this identity, we can decompose the effect on preferred-option matching into within-institution versus across-institution components. Since programs and institutions are automatically within-institution, we consider the field and elite treatments. The left panel of Figure 15 reports the results. We find that the effects of enrollment on both preferred-field matching and preferred-elite matching are completely explained by within-institution matches.

By a similar argument, matches with a partner who completed the applicant's preferred option are either within the same field ($M \times F$) or across different fields ($M \times !F$),

$$\Pr(M) = \Pr(M \times F) + \Pr(M \times !F),$$

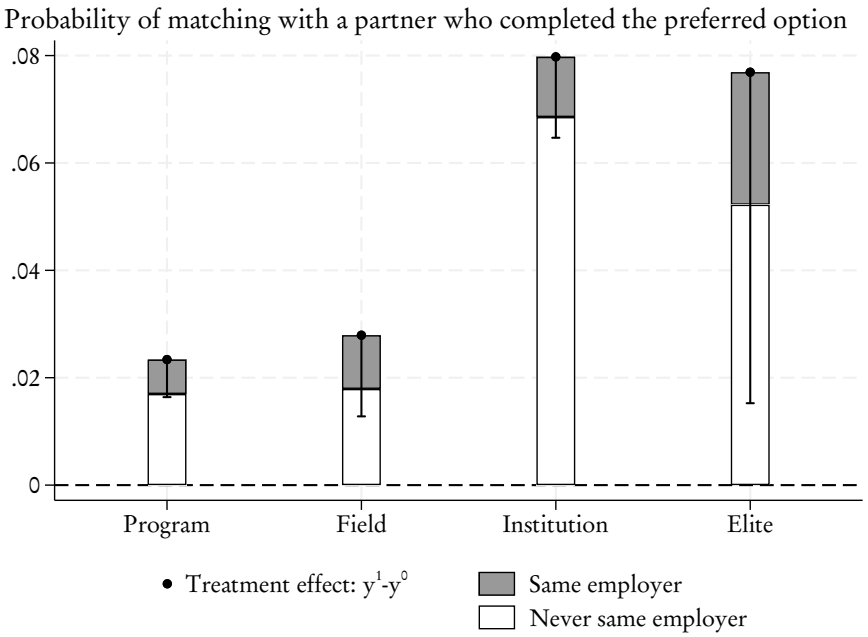


Note: This figure decomposes the 2SLS effects of enrollment on the probability of matching with a partner who completed the preferred option into within-cohort versus across-cohort matches, where same cohort is defined as the applicant and partner attending the same institution in the same year. Error bars indicate 95% CI (standard errors clustered at the applicant level).

Figure 16. Match-making mechanisms: within versus across cohort

enabling a decomposition of the effect of enrollment on preferred-option matching into within-field versus across-field components. Since programs and fields are automatically within-field, we consider the institution and elite treatments. These results are reported in the right panel of Figure 15. About one third of the effect of enrollment on preferred-institution homogamy is explained by within-field matches, while the remainder is due to across-field matches. For elite enrollment, on the other hand, the large majority of the match-making effect is among students in the same field, echoing the descriptive results on elite household homogamy in Figure 3.

Within versus across cohort If search costs are at play in the matching market, students may be more likely to match with partners who not only attend the same institution but do so at the same time, which we denote as a shared cohort. To test this, we use the identity that preferred-option matches are either with a partner from



Note: This figure decomposes the 2SLS effects of enrollment on the probability of matching with a partner who completed the preferred option into matches with shared versus non-shared employer histories after college enrollment and before matching. Error bars indicate 95% CI (standard errors clustered at the applicant level).

Figure 17. Match-making mechanisms: within versus across employers

the same cohort as the applicant (C) or not ($!C$):

$$\Pr(M) = \Pr(M \times C) + \Pr(M \times !C).$$

Figure 16 reports the results. We find that the vast majority of match-making effects are within-cohort, with a bit more across-cohort matches created by elite education.

Within versus across employers While educational choices matter for whom you meet in college, they may also affect matching after graduation through work. Thanks to our employer-employee data, we can not only track people through the education system, but also into the firms that employ them. We can therefore measure whether partners ever overlapped at the same workplace (W) after college before matching:

$$\Pr(M) = \Pr(M \times W) + \Pr(M \times !W)$$

The results in Figure 17 show that only a small part of college match-making effects can be attributed to sharing the same workplace after college. This is also likely an overestimate of the importance of shared workplace as a separate mechanism from college enrollment, since sharing a workplace might itself be a downstream consequence of first meeting during college.

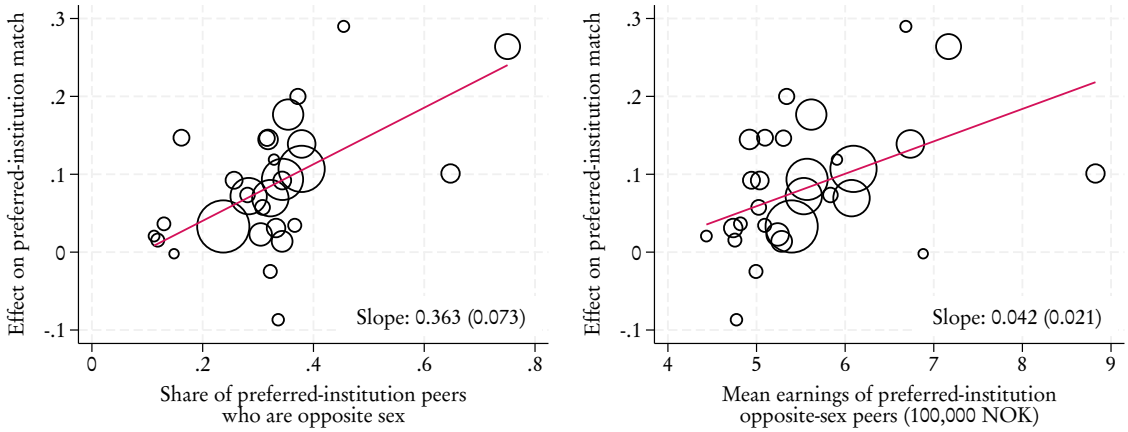
6.2 Variation in potential partner pools across institutions

The previous results revealed the central importance of institutions as match-makers. A natural question, then, is whether different institutions have different effects on matching, and whether such effects can be explained by variation in the types of potential partners available as peers. To investigate this, we estimate distinct effects of enrolling in each preferred institution on the probability of matching with a partner from that institution. We use our main specification but estimate it separately for each locally preferred institution in our applicant sample and separately for female versus male applicants. In the left panels of Figure 18, we plot these gender-specific institution effects against the share of peers at that institution who would be of the opposite sex from the applicant. In the right panels, we plot the institution effects against the mean earnings of that institution's opposite-sex peers.

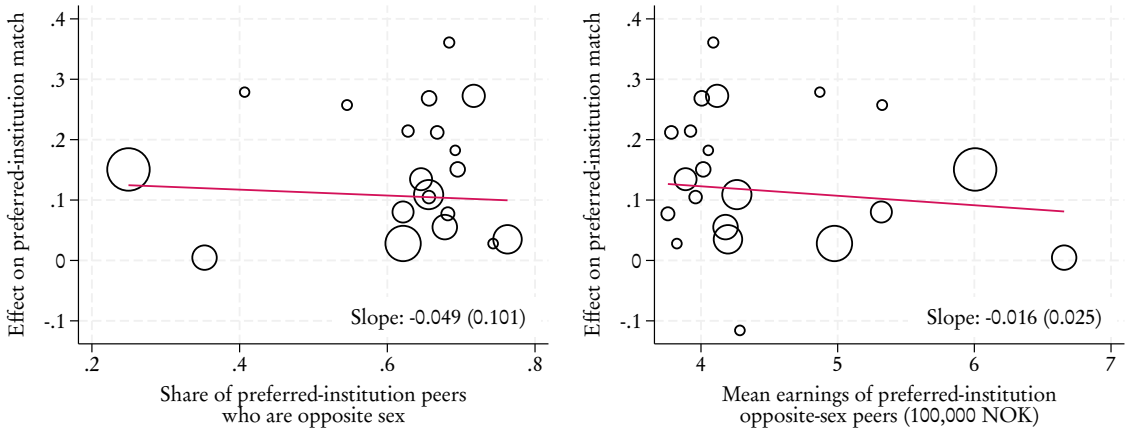
The results show strikingly different patterns across female versus male applicants. For women, the effect of enrolling in a particular institution on the probability of matching with a partner from that institution is systematically higher at institutions in which male peers are relatively more abundant, and conversely trend towards zero as male peers become scarce. Institutional effects on matching are also systematically larger for female applicants when their male peers are higher earners, consistent with the especially large match-making effects of elite programs on female applicants in Section 5.[17] For men, on the other hand, there is no systematic relationship between the effect of an institution on matching and the abundance or earnings of their female peers at that institution.

[17]In Appendix Figure A5, we conduct the same exercise but replace the outcome with ever matching with any partner, not specifically a partner from the same institution. The results show institution-specific effects that are clustered around zero with no systematic trends in the share of opposite sex peers or their mean earnings.

(a) Female applicants



(b) Male applicants



Note: These figures plot separate treatment effects for each preferred institution in our applicant sample, estimated using our main specification but separately for each locally preferred institution and applicant gender. Institutions are weighted by their number of applicants. Mean earnings of opposite-sex peers are measured in hundred thousands of Norwegian kroner when those peers are 13 years out from their initial application.

Figure 18. Variation across institutions in match effects and peers

6.3 Implications for the functioning of the marriage market

Two important features of college match-making effects have emerged from the preceding results. First, these effects are concentrated among students who not only attend the same institution, but do so at exactly the same time. Second, these effects are larger when opposite-sex peers are more abundant, at least for female applicants.

Together, these results suggest the presence of substantial search costs in the marriage market. If enrolling in a particular option changed an individual's type in a broad matching market with low cost of finding partners, then we would expect a larger share of the enrollment effects to be driven by matches across different institutions and cohorts, since the particular location and timing of the education should not matter much for matching on the new type. Furthermore, we would not expect the density of opposite-sex peers within a particular institution to matter much for matching in a broader market. Instead, virtually none of the effects of enrollment are explained by cross-institution or cross-cohort matches, and they increase when women find themselves on campuses with more men, suggesting that search costs are an empirically important feature of the marriage market.

These results do not necessarily imply that type changes are an unimportant consequence of enrollment; the earnings effects for elite women in Section 5.4 indeed show transformations in labor market outcomes from enrolling in the preferred option. But our results do suggest that acting on such type changes in the marriage market will be more difficult outside of the particular pool of peers who find themselves on campus together at the same time, and even inside that pool when opposite-sex peers are more scarce.

7 Conclusion

In this paper, we studied college as a marriage market. We first documented strong assortativity by institution and field of study among the college-educated, especially among high earners from elite programs. We then exploited admission discontinuities to decompose this observed assortativity into causal college match-making versus inevitable selection. Our decomposition results showed that assortative matches are predominantly *caused* by actually enrolling in a particular college treatment, rather than simply reflective of the latent proclivities of those who select into it.

We also found that these college match-making effects can be economically con-

sequential, with the types of applicants experiencing larger match-making effects of enrolling in their preferred option also tending to experience larger gains in household income. Elite professional programs, in particular, propel marginally admitted women into elite household formation: they earn substantially more themselves and match with substantially higher-earning elite-educated partners, becoming much more likely to join the top few percentiles of the household earnings distribution. We do not observe this for marginally admitted elite men. Results also diverge by parental education. For applicants with at least one college-educated parent, elite enrollment substantially increases their earnings and their likelihood of entering the top ranks of the household income distribution. For applicants with no college-educated parents, elite enrollment yields smaller earnings gains and no greater chance of entering the top earnings percentiles.

Finally, we investigated the mechanisms behind these college match-making effects to shed new light on how the marriage market works. We found that match-making effects are concentrated among students who attend the same institution at exactly the same time, and are larger at institutions where opposite-sex peers are more abundant, suggesting that search costs are an empirically important feature of the marriage market.

References

Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics*, 113(2):231–263.

Artmann, E., Ketel, N., Oosterbeek, H., and van der Klaauw, B. (2021). Field of study and partner choice. *Economics of Education Review*, 84:102–149.

Atakan, A. E. (2006). Assortative matching with explicit search costs. *Econometrica*, 74(3):667–680.

Banerjee, A., Duflo, E., Ghatak, M., and Lafortune, J. (2013). Marry for what? Caste and mate selection in modern India. *American Economic Journal: Microeconomics*, 5(2):33–72.

Barrios-Fernandez, A., Neilson, C. A., and Zimmerman, S. D. (2024). Elite universities and the intergenerational transmission of human and social capital. IZA Discussion Paper No. 17252.

Belot, M. and Francesconi, M. (2013). Dating preferences and meeting opportunities in mate choice decisions. *The Journal of Human Resources*, 48(2):474–508.

Bičáková, A. and Jurajda, Š. (2017). Field-of-study homogamy. GLO Discussion Paper.

Bleemer, Z. (2024). Top percent policies and the return to postsecondary selectivity. Working Paper.

Bloch, F. and Ryder, H. (2000). Two-sided search, marriages, and matchmakers. *International Economic Review*, 41(1):93–116.

Blossfeld, H.-P. (2009). Educational assortative marriage in comparative perspective. *Annual Review of Sociology*, 35:513–530.

Britton, J., Drayton, E., and van der Erve, L. (2021). Which university degrees are best for intergenerational mobility. IFS Report.

Bruze, G. (2011). Marriage choices of movie stars: Does spouse's education matter? *Journal of Human Capital*, 5(1):1–28.

Burdett, K. and Coles, M. G. (1997). Marriage and class. *The Quarterly Journal of Economics*, 112(1):141–168.

Bütikofer, A., Jensen, S., and Salvanes, K. G. (2018). The role of parenthood on the gender gap among top earners. *European Economic Review*, 109:103–123.

Cattan, S., Salvanes, K. G., and Tominey, E. (2022). First generation elite: The role of school networks. IZA Discussion Paper No. 15560.

Chetty, R., Deming, D. J., and Friedman, J. N. (2023). Diversifying society's leaders? The determinants and causal effects of admission to highly selective private colleges. National Bureau of Economic Research Working Paper No. 31492.

Chiappori, P.-A. (2020). The theory and empirics of the marriage market. *Annual Review of Economics*, 12:547–578.

Chiappori, P.-A., Costa Dias, M., Meghir, C., and Zhang, H. (2025). Changes in marital sorting: Theory and evidence from the US. National Bureau of Economic Research Working Paper No. 33354.

Dale, S. B. and Krueger, A. B. (2002). Estimating the payoff to attending a more selective college: An application of selection on observables and unobservables. *The Quarterly Journal of Economics*, 117(4):1491–1527.

Eeckhout, J. (1999). Bilateral search and vertical heterogeneity. *International Economic Review*, 40(4):869–887.

Eika, L., Mogstad, M., and Zafar, B. (2019). Educational assortative mating and household income inequality. *Journal of Political Economy*, 127(6):2795–2835.

Fisman, R., Iyengar, S. S., Kamenica, E., and Simonson, I. (2006). Gender differences in mate selection: Evidence from a speed dating experiment. *The Quarterly Journal of Economics*, 121(2):673–697.

Fisman, R., Iyengar, S. S., Kamenica, E., and Simonson, I. (2008). Racial preferences in dating. *The Review of Economic Studies*, 75(1):117–132.

Ford, K. S. (2020). Marrying within the alma mater: Understanding the role of same-university marriages in educational homogamy. *Sociological Research Online*, 25(2):254–272.

Ge, S., Isaac, E., and Miller, A. (2022). Elite schools and opting in: Effects of college selectivity on career and family outcomes. *Journal of Labor Economics*, 40(S1):S383–S427.

Goldin, C. (1992). The meaning of college in the lives of American women: The past one-hundred years. National Bureau of Economic Research Working Paper No. 4099.

Han, S. and Qian, Y. (2024). Concentration and dispersion: School-to-work linkages and their impact on occupational assortative mating. *The Social Science Journal*, 61(4):755–772.

Hitsch, G. J., Hortaçsu, A., and Ariely, D. (2010). Matching and sorting in online dating. *American Economic Review*, 100(1):130–163.

Hoekstra, M. (2009). The effect of attending the flagship state university on earnings: A discontinuity-based approach. *The Review of Economics and Statistics*, 91(4):717–724.

Jacquet, N. L. and Tan, S. (2007). On the segmentation of markets. *Journal of Political Economy*, 115(4):639–664.

Juhn, C. and McCue, K. (2017). Specialization then and now: Marriage, children, and the gender earnings gap across cohorts. *Journal of Economic Perspectives*, 31(1):183–204.

Kaufmann, K. M., Messner, M., and Solis, A. (2021). Elite higher education, the marriage market and the intergenerational transmission of human capital. CRC TR 224 Discussion Paper Series.

Kirkebøen, L. J. (2010). Forskjeller i livsløpsinntekt mellom utdanningsgrupper. Statistics Norway Report.

Kirkeboen, L. J., Leuven, E., and Mogstad, M. (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics*, 131(3):1057–1111.

Landersø, R. and Heckman, J. J. (2017). The scandinavian fantasy: The sources of intergenerational mobility in Denmark and the US. *The Scandinavian Journal of Economics*, 119(1):178–230.

Lee, S. (2016). Effect of online dating on assortative mating: Evidence from South Korea. *Journal of Applied Econometrics*, 31(6):1120–1139.

Liu, H. and Lu, J. (2006). Measuring the degree of assortative mating. *Economics Letters*, 92(3):317–322.

Lovenheim, M. and Smith, J. (2023). Returns to different postsecondary investments: Institution type, academic programs, and credentials. In *Handbook of the Economics of Education*, volume 6, pages 187–318. Elsevier.

Mansour, H. and McKinnish, T. (2018). Same-occupation spouses: Preferences or search costs? *Journal of Population Economics*, 31(4):1005–1033.

Mountjoy, J. (2024). Marginal returns to public universities. National Bureau of Economic Research Working Paper No. 32296.

Mountjoy, J. and Hickman, B. R. (2021). The returns to college(s): Relative value-added and match effects in higher education. National Bureau of Economic Research Working Paper No. 29276.

Nielsen, H. S. and Svarer, M. (2009). Educational homogamy: How much is opportunities? *The Journal of Human Resources*, 44(4):1066–1086.

Oreopoulos, P. and Salvanes, K. G. (2011). Priceless: The nonpecuniary benefits of schooling. *Journal of Economic Perspectives*, 25(1):159–184.

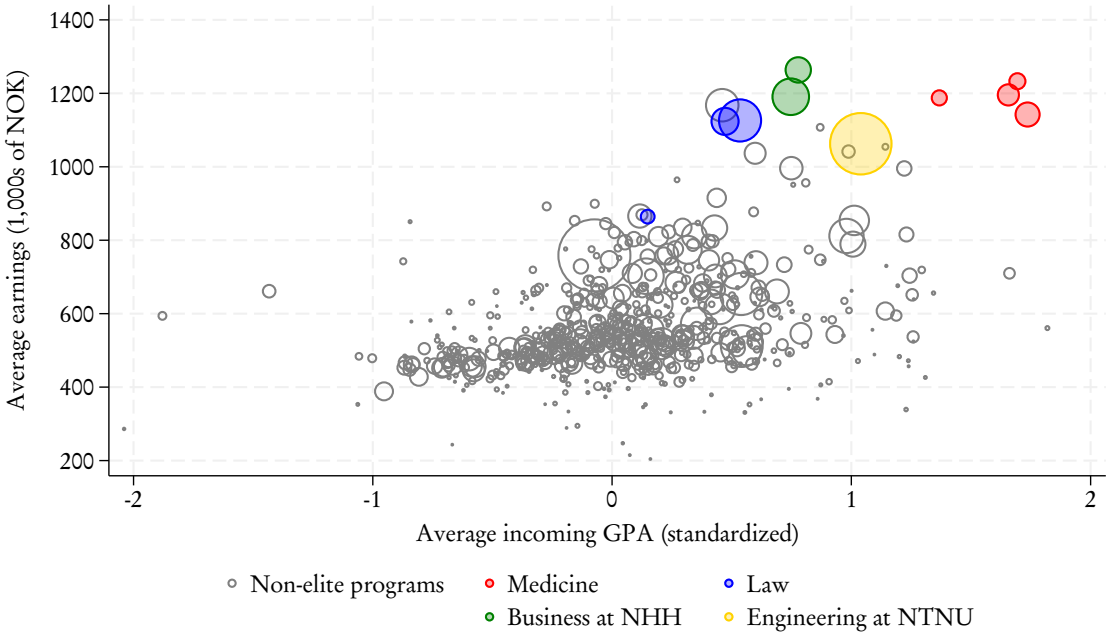
Schwartz, C. R. (2013). Trends and variation in assortative mating: Causes and consequences. *Annual Review of Sociology*, 39:451–470.

Shimer, R. and Smith, L. (2000). Assortative matching and search. *Econometrica*, 68(2):343–369.

Smith, J., Goodman, J., and Hurwitz, M. (2025). The economic impact of access to public four-year colleges. *Journal of Human Resources*, Forthcoming.

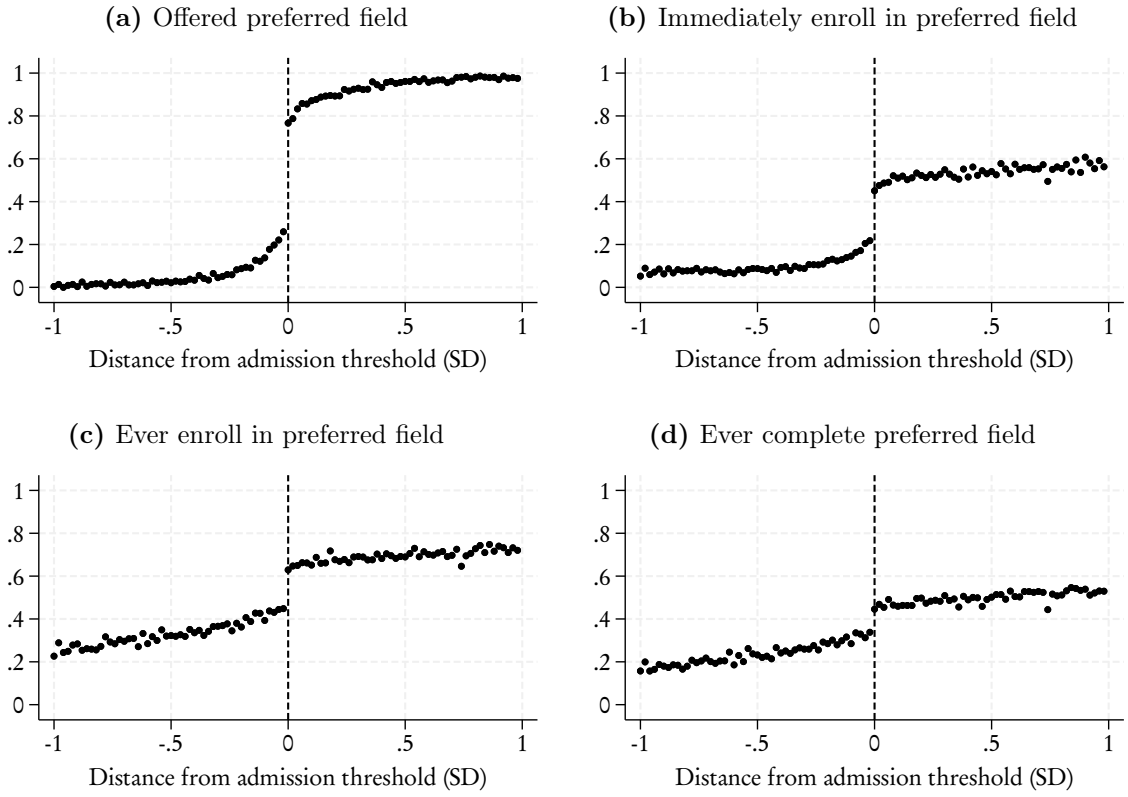
Zimmerman, S. D. (2019). Elite colleges and upward mobility to top jobs and top incomes. *American Economic Review*, 109(1):1–47.

Appendix



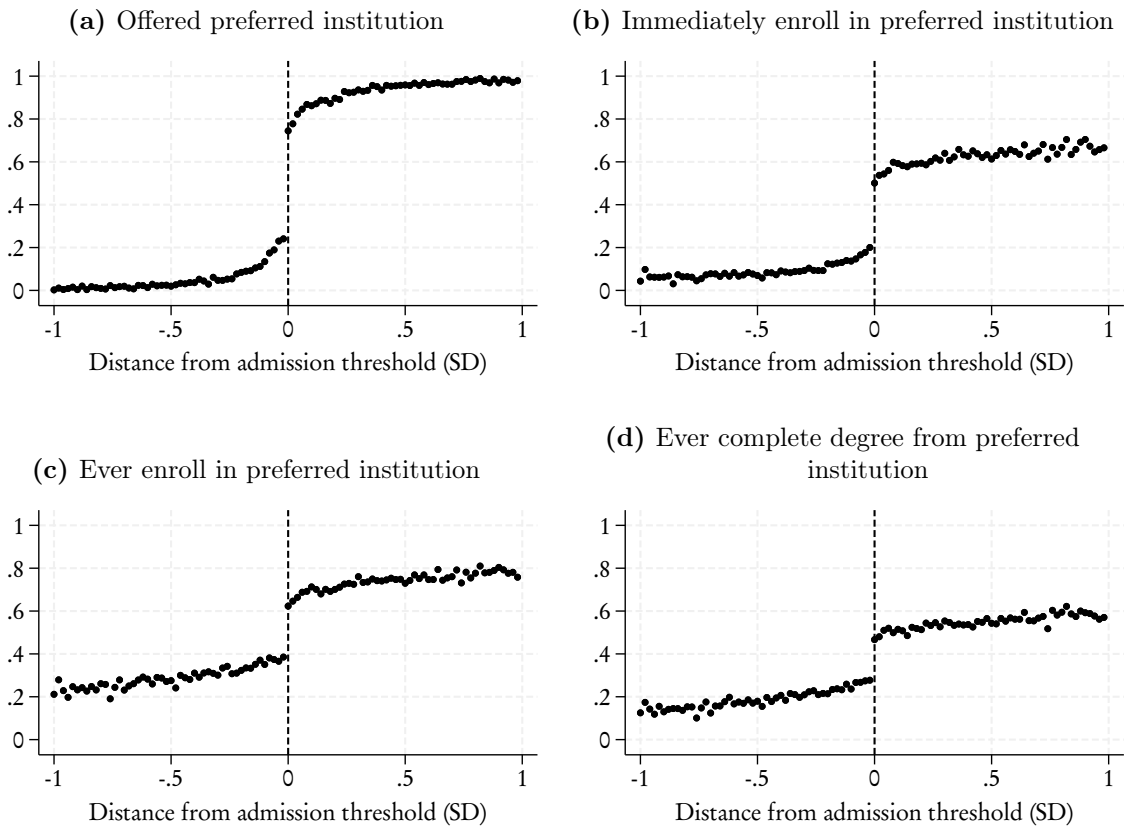
Note: Following Cattan et al. (2022) and Bütikofer et al. (2018), elite programs include Medicine or Law at any institution, Business at the Norwegian School of Economics (NHH), and Engineering at the Norwegian University of Science and Technology (NTNU). The earnings measure begins with all individuals in Norway aged 40-45 in 2018. We average their earnings over the years in which they are 40-45. We then plot the average earnings and average high school grade point average (GPA), standardized to have mean zero and standard deviation one, among individuals who completed each program, with the circle size reflecting market share.

Figure A1. Elite vs. non-elite programs: average earnings and GPA



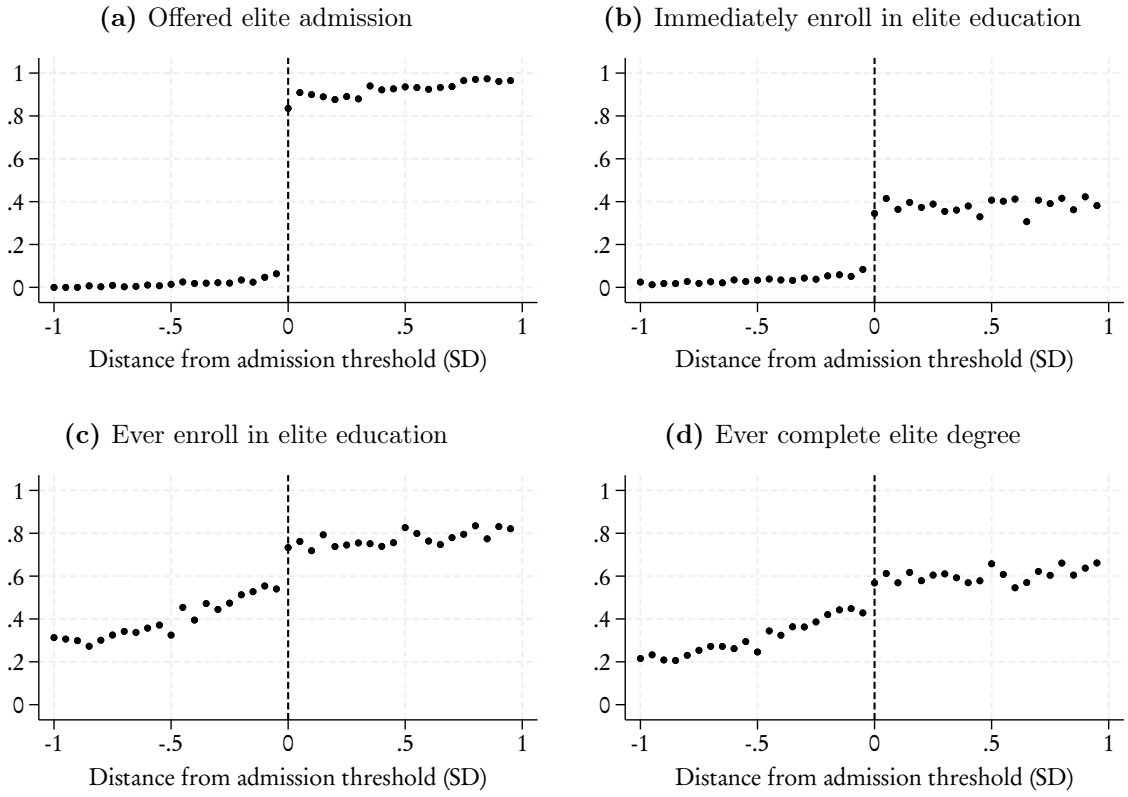
Note: Admission offers in panel (a) are measured in the applicant's first observed application cycle. Immediate enrollment in panel (b) is defined as enrolling in the academic year immediately following the initial application. The outcomes in the remaining panels are all measured cumulatively as of 13 years after the initial application.

Figure A2. What does change across the admission thresholds: field



Note: Admission offers in panel (a) are measured in the applicant's first observed application cycle. Immediate enrollment in panel (b) is defined as enrolling in the academic year immediately following the initial application. The outcomes in the remaining panels are all measured cumulatively as of 13 years after the initial application.

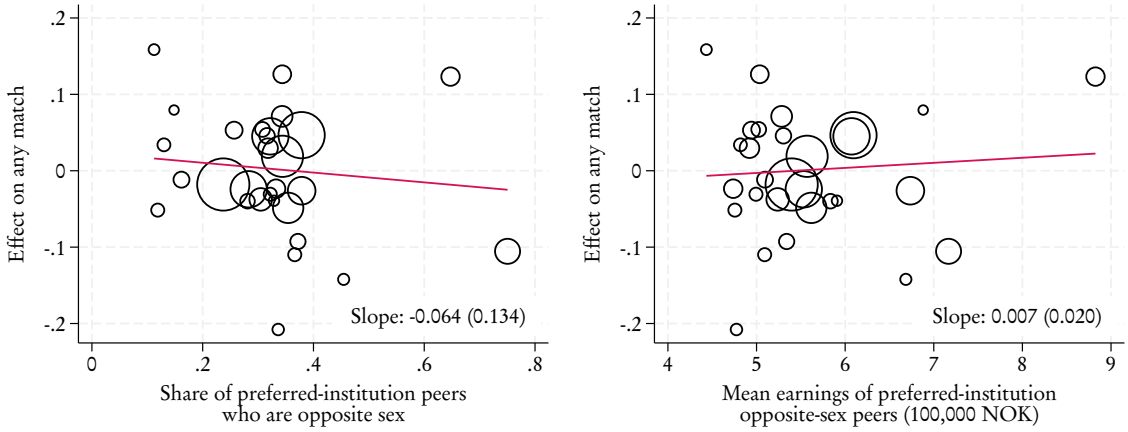
Figure A3. What does change across the admission thresholds: institution



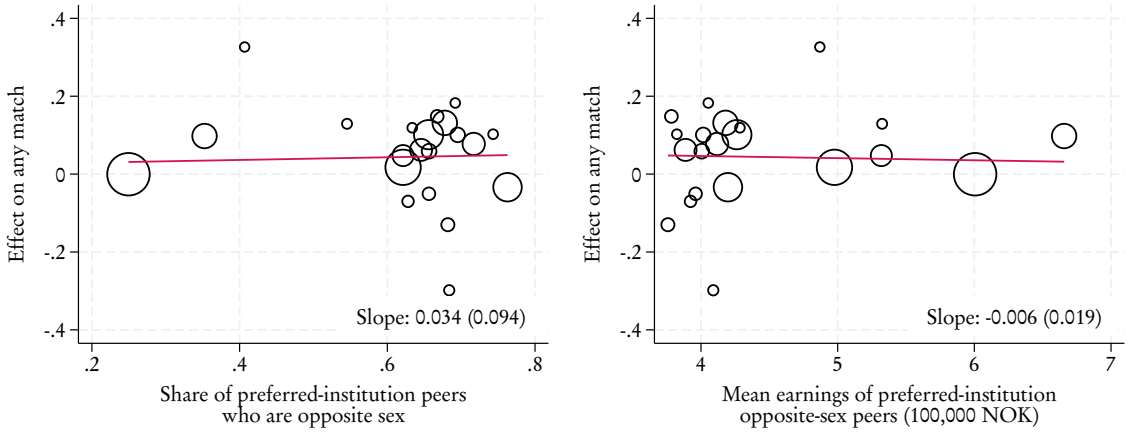
Note: Admission offers in panel (a) are measured in the applicant's first observed application cycle. Immediate enrollment in panel (b) is defined as enrolling in the academic year immediately following the initial application. The outcomes in the remaining panels are all measured cumulatively as of 13 years after the initial application.

Figure A4. What does change across the admission thresholds: elite education

(a) Female applicants



(b) Male applicants



Note: These figures conduct the same exercise as Figure 18 but replace the outcome with matching with any partner, not specifically a partner from the same institution. We estimate separate treatment effects for each preferred institution in our applicant sample using our main specification but separately for each locally preferred institution and applicant gender. Institutions are weighted by their number of applicants. Mean earnings of opposite-sex peers are measured in hundred thousands of Norwegian kroner when those peers are 13 years out from their initial application.

Figure A5. Variation across institutions in match effects and peers: matching with any partner

Table A1. College homogamy and assortativity, assuming county-level matching markets

	Homogamy			Assortativity	
	Observed h	Random h^r	Maximal h^m	Absolute $h - h^r$	Rescaled $\frac{h-h^r}{h^m-h^r}$
Same program	.123	.012	.596	.110	.189
Same field	.165	.039	.659	.126	.204
Same institution	.343	.100	.827	.243	.334
Both elite	.049	.018	.108	.030	.338

Note: This table conducts the same exercise as Table 1 but assumes that the relevant matching market is the county, such that potential matches happen within but not across the twenty regional counties of Norway.