THÈSE de DOCTORAT

de l’UNIVERSITÉ TOULOUSE CAPITOLE

Présentée et soutenue par

Madame Anaïs FABRE

Le 12 juin 2024

Essais en Economie de l'Education et Econométrie

École doctorale : **TSE Toulouse Sciences Economiques**

Spécialité : **Sciences Economiques - Toulouse**

Unité de recherche : **TSE-R - Toulouse School of Economics - Recherche**

Thèse dirigée par *Monsieur Thierry MAGNAC et Olivier DE GROOTE*

Composition du jury

Rapporteure : Mme Michela TINCANI
Rapporteur : M. Peter ARCIDIACONO
Examinateur : M. Christian BELZIL
Directeur de thèse : M. Thierry MAGNAC
Co-encadrant de thèse : M. Olivier DE GROOTE

Essays on Higher Education Access and Econometrics

Anaïs Fabre

June 12, 2024

Acknowledgement

I am extremely grateful for the invaluable mentorship provided by my advisors, without whom this thesis would not have been possible. I would like to thank Olivier De Groote for his unwavering support over the last five years. Having you as an advisor was a great privilege of my PhD journey. I thank you for always being very generous with your time, and for the unfailing guidance you provided. You will be an inspiration for me in the future, and I aspire to transmit the same healthy vision of academia that you always put forward.

I also want to thank Thierry Magnac, for his involvement in my thesis. From your insightful classes in Master 1 to your thorough feedback on my job market paper, you influenced all my early steps as an economist. Thank you for your invaluable advice.

Matteo Bobba has also been a constant presence in my journey at TSE, as the supervisor of my very first research project in Master 1 and as a professor in the Master Public Policy & Development. Your enthusiasm when discussing research in class was a key piece in my decision to become a researcher. Thank you for your incredible energy and for always pushing me, I deeply enjoyed our numerous discussions, and I hope they will continue in the future.

This thesis has been influenced by many other professors in TSE, which has been the perfect place to grow as a researcher. In particular, I want to thank Christian Hellwig, for his insightful advice on my job market paper and for always pushing me to do better, Ana Gazmuri, for supervising my Master's thesis and making me discover the interaction of empirical industrial organization and the economics of education, and Nour Meddahi, for his dedication to improving the PhD program in TSE. I have greatly benefited from your efforts.

My graduate studies were also enriched by people beyond TSE. I would like to thank Arnaud Maurel, for his invaluable comments and advice. I have learned a lot from being your co-author. I also want to thank Margaux Luflade and Hanming Fang for welcoming me during my visit in Penn, where I had the chance to spend an incredibly stimulating semester.

I also want to thank the friends I met in Philly, who made my visit in Penn unforgettable. Thank you Natacha, Laure, Kathleen, Wei, Javiera, and Rita for many great moments, often shared over nice food. Special thanks to Laure, Natacha, and Wei for their early morning cheers during the Philly half-marathon! Thank you so much to Andrea and See Yu, for being 'not so bad' office mates. I hope to share many more cups of (drip) coffee with you.

The past five years would have never been the same without the wonderful friends I made in Toulouse. Thank you to my cohort of adoption, Alfonso, Lisa, Tuuli, Max, Hippolyte,

Gökçe, Moritz. I cherish all the hikes, coffees, drinks we shared, and the trips to Aveyron (I had to find a way to acknowledge where I am from, you know that). A special thank you to Léa, for being always up for anything and for asking the DJ to play *La Tribu de Dana*. And thank you to all the friends I have made along the way, with whom this journey would not have been the same — Guillem, Juan, Tanja, Vale, Gyung Mo, Gilles, Maria, Lorenzo. I hope we will dance over *Baby, one more time* many more times. Alice et Laura, merci d'avoir été là pour m'écouter, mais aussi pour me faire oublier le stress du doctorat quand j'en avais besoin.

Je tiens également à remercier ma famille, de leur soutien sans faille et d'être mes points d'ancrage. A mes parents, en particulier, car ma thèse vous doit énormément. Bien que très loin du monde universitaire, votre ouverture d'esprit et votre intelligence m'ont appris à me poser beaucoup de questions sur le monde et à être curieuse. Je vous admire pour ces qualités, et je vous remercie de m'avoir laissé la liberté totale de choisir ma voie, bien que si différente de ce que vous connaissiez. Ma recherche est profondément enrichie par l'éducation que vous m'avez offerte et par le milieu dans lequel j'ai eu la chance de grandir. Des inégalités spatiales dans l'accès à l'éducation au manque d'information dans le choix des études supérieures, tous mes papiers portent la marque de là d'où je viens. (A l'exception, peut-être, de celui sur la double-différence, mais je cherche encore un lien caché.) J'espère continuer à puiser dans cette richesse, et c'est toujours avec fierté que je parlerai de mes racines. Je vous remercie pour tout cela, et pour votre soutien inestimable tout au long de mon parcours.

Enfin, bien sûr, je souhaite remercier Tim. Quelle chance de partager l'aventure du doctorat avec non seulement un économiste aussi brillant, mais, surtout, une personne aussi merveilleuse. Sans toi, je n'aurais tout simplement pas commencé cette thèse. Et sans ton soutien constant, j'aurais souvent baissé les bras. Merci de croire en moi plus que je ne le fais moi-même. Merci aussi pour ton enthousiasme à discuter de recherche, ces longues conversations ont largement contribué à cette thèse et m'ont tellement appris. Tu as aussi toujours su me faire prendre du recul et me rappeler que la recherche, ce n'est pas si sérieux, après tout. Mais, par-dessus tout, merci de m'apporter tant de bonheur, merci pour ces kilomètres de course à pied et de randonnée partagés, merci pour ces éclats de rire au quotidien, merci pour ces voyages incroyables, et merci pour ces kilomètres de route au son de la même playlist. J'espère continuer à m'émerveiller à tes côtés.

Overview

This thesis is composed of four chapters that have two main goals: (i) broaden our understanding of the determinants of inequalities in access to higher education, and (ii) expand and deepen our understanding of available econometric tools.

Chapter 1, 'The Geography of Higher Education and Spatial Inequalities', uses rich administrative data to document that the spatial distribution of higher education options is highly uneven, while students' demand for programs is elastic to their geographic proximity. Leveraging policy changes, I show that mitigating barriers to mobility would allow students to apply farther away from their home location, but could generate a long-lasting brain drain toward higher education hubs. I build and estimate a dynamic model encompassing these channels, linking equilibrium sorting on the higher education market and location choices of entry-level workers. I find that the interaction of the uneven distribution of colleges and mobility frictions accounts for one-third of regional gaps in educational attainment. Eliminating mobility frictions, however, generates a trade-off, as it benefits students from low-opportunity areas but accelerates their migration to higher education hubs, magnifying regional inequalities. Policies tying mobility scholarships and incentives to return would resolve this tension.

Chapters 2 and 3 explore additional drivers of inequalities in access to higher education across students. Chapter 2, 'Opportunity Costs of Time and the Design of College Admission Mechanisms', co-authored with Olivier De Groote, Margaux Luflade and Arnaud Maurel, documents how the design of college admission platforms impacts students' enrollment outcomes. We explore the welfare and distributional effects of using a sequential assignment procedure to match students to programs. Such mechanisms, commonly used around the world, create a dynamic trade-off for students: they can choose to delay their enrollment decision to receive a better offer later, at the cost of waiting before knowing their final admission outcome. Building and estimating a dynamic model of application and enrollment, we find that waiting costs are a key determinant of the timing of students' acceptance decisions and of their final assignment. We find substantial, but unequal, welfare gains from using a multi-round system.

Chapter 3, 'Incomplete Information and the Complexity of Centralized College Application Processes', focuses on information frictions in college admission procedures. Using data from the Chilean centralized admission mechanism, I document that 26% of college applicants submit an invalid application, suggesting that misinformation about the rules of the system is widespread. To study the impact of being misinformed on students' admission outcomes, I build and estimate a model of application decisions under incomplete information. 80% of unassigned uninformed students would have been assigned in a counterfactual scenario where information frictions are eliminated. Results indicate that the complexity of the application process disproportionately harms disadvantaged students.

Finally, Chapter 4, 'Robustness of Two-Way Fixed Effects Estimators to Heterogeneous Treatment Effects', aims at deepening economists' understanding of the widely used Two-Way Fixed Effects (TWFE) estimator. It provides necessary and sufficient conditions for this estimator to be robust to heterogeneous treatment effects. I decompose the estimator to show that it is a weighted sum of five types of difference-in-differences, with positive weights. Parallel trends assumptions on either the untreated or treated potential outcomes must hold for each comparison to identify the Average Treatment Effect (ATE) of the group switching treatment status, when the effect of the treatment is contemporaneous. Both parallel trends assumptions are thus necessary and sufficient for the TWFE estimator to weigh each ATE positively.

Contents

Chapter 1

The Geography of Higher Education and Spatial Inequalities

Abstract. This paper shows that the within-country spatial distribution of colleges largely contributes to spatial inequalities. Using data on the universe of college applicants and programs in France, I document that higher education options are unevenly distributed across space while students' demand is highly sensitive to geographic proximity. This creates inequalities in access to higher education across space, feeding gaps in educational attainment and spatial skill sorting. To quantify these effects, I build a dynamic model linking equilibrium sorting on the higher education market and location choices of entry-level workers. I show that students' and programs' preferences can be identified and estimated from data on choices and equilibrium outcomes. One-third of regional gaps in educational attainment are explained by the interaction of the uneven distribution of colleges and mobility frictions. Eliminating the latter, however, generates a trade-off, as it benefits students from low-opportunity areas but accelerates their migration to higher education hubs, magnifying regional inequalities. Low-opportunity areas could halt this brain drain and outsource the education of their local labor force by tying mobility scholarships with incentives to return.

1.1 Introduction

While human capital is a fundamental economic input, individuals' educational attainment widely differs across birthplaces. The gap in college graduation rates between commuting zones in the United States reaches up to 50 percentage points (Chetty et al., 2018b). These disparities reveal the existence of local barriers to investment in human capital, in turn hindering growth (Hsieh et al., 2019). Spatial differences in educational attainment may also drive the observed concentration of skills across local labor markets, which impacts regional inequalities and efficiency (Diamond and Gaubert, 2022). Understanding the sources and welfare consequences of these spatial gaps is thus key.

Inequalities in economic opportunities across birthplaces and earnings differences across regions have been studied separately. A first strand of the literature documents the causal effect of individuals' childhood location on their lifetime outcomes (Chetty and Hendren, 2018; Chyn and Katz, 2021). Its underlying mechanisms, however, remain unclear. A distinct literature identifies worker sorting as a main cause of earnings disparities across local labor markets (Combes et al., 2008; Dauth et al., 2022; Card et al., 2023). While sorting is often understood as the result of individuals' migration decisions (Diamond, 2016), Combes et al. (2012) suggest that another important factor may have been overlooked: 'sorting at birth', whereby individuals attain different skill levels due to disparities in the fundamental characteristics of their home region. This paper bridges these two strands of the literature by identifying a key driver of inequalities across both birthplaces and local labor markets: the geography of the higher education market. It uses rich administrative data, quasi-experimental policy variations, and a dynamic structural model to show how the within-country distribution of higher education programs interacts with mobility frictions to generate spatial gaps in students' welfare and educational attainment, in turn contributing to spatial skill sorting and regional inequalities.

I investigate this link in four steps providing several empirical and methodological contributions. First, I use data on the universe of college applicants and programs in France to empirically document that (i) higher education options are highly unevenly distributed across space, while (ii) students' demand for programs is elastic to geographic proximity. Second, I leverage local policy variations to show that mitigating mobility frictions within the higher education market allows students from educational deserts to apply farther away from home but could trigger a long-lasting brain drain toward higher education hubs. Third, I build a dynamic model combining a two-sided matching model of the higher education market with a model of location decisions at entry on the labor market. I show that students' and programs' preferences are identified from data on choices and equilibrium outcomes, without parametric assumptions on the deterministic part of their payoffs. Fourth, I take the model to the data and investigate the effects of eliminating mobility frictions — encompassing pecuniary and non-pecuniary costs of studying away from home — within the French higher education market, while holding constant the spatial distribution of programs. I find that spatial gaps in educational attainment would decrease by one-third. However, mitigating mobility frictions would generate a trade-off, as it would benefit students from low-opportunity areas but accelerate their migration to higher education hubs, increasing regional divergence. I show how policies tying mobility scholarships to incentives to return would allow low-opportunity areas to outsource the education of their local labor force to higher education hubs.

Specifically, I begin by using rich administrative data to empirically document two facts linking the geography of the higher education market, individuals' educational choices, and spatial inequalities. First, I show that higher education options are unevenly distributed across space. Out of the ninety-six provinces (*départements*) of metropolitan France, fifty gather less than 5% of the total supply of seats in academic programs, leading to a bachelor's degree. This does not reflect the spatial distribution of the population, as these areas concentrate 25% of high school students. In contrast, 50% of the academic seats supply is agglomerated in only twelve provinces.

Second, I provide evidence suggesting that students' demand for higher education programs is highly sensitive to proximity. More than 50% of students enroll in a program located less than 25 kilometers away from where they live. This pattern reveals students' preferences rather than colleges' admission decisions, as the median offer received by students is located twenty kilometers farther away than the accepted one. Overall, by estimating a gravity model, I find that the elasticity of students' enrollment flows to distance is -2.5, which is substantial and larger than that found for labor migration in France (Zerecero, 2021). College choices are thus highly sensitive to distance from home. As these mobility frictions interact with the uneven spatial distribution of colleges, I find that students' probability to enroll in an academic program increases with their proximity to such institutions. This may further impact local labor markets, as I document a positive relationship between the share of a region's labor force with at least a three-year degree and the local availability of programs delivering such diplomas.

In this context, breaking the dependence of students' application decisions on the local supply of higher education options seems key to tackling spatial inequalities. I leverage local policy changes to explore the direct and distributional effects of mitigating mobility frictions. I first evaluate a program guaranteeing moving subsidies to low-income students conditional on enrolling out of their home region. I find that it allowed students to apply to programs farther away from their home location. To document the distributional effects of removing barriers to mobility, I then zoom in on the Greater Paris area and exploit a change in geographic priorities determining students' university admission. The removal of the priority of students from Paris to enroll in their local universities led to a 50% increase in the share of high-ability students from peripheral locations enrolled in programs in Paris. These students left their local universities and crowded out low-ability applicants from Paris. Survey data on education and work migration flows suggest that, if happening at the scale of the country, such a brain drain could be long-lasting. Indeed, more than half of the students graduating out of their home location do not come back at entry on the labor market.

I build a dynamic model encompassing these mechanisms. It flexibly accounts for differences in preferences and admission chances across students, and captures potential long-run distributional effects of students' spatial mobility. The framework combines a two-sided matching model of the higher education market, where students are forward-looking, with a model of optimal location choices at entry on the labor market. In the first period, students and programs meet on a centralized platform and express their preferences over one another. Each student's payoff from enrolling in a given program depends on its associated flow and continuation value and mobility frictions, capturing pecuniary and non-pecuniary costs associated with studying away from home. I assume that the resulting equilibrium is stable with respect to students' lifetime utility and programs' preferences, and that programs rank applicants truthfully. The second stage of the model takes place after graduation, when individuals choose the location in which to enter the labor market, taking into account migration costs, a potential home bias, and location-specific wages and amenities.

I describe how to identify the primitives of this model without imposing parametric assumptions on the deterministic part of agents' payoffs. I first show how to identify preferences over local labor markets from data collecting individuals' home, study, and work locations.

I then develop a novel procedure combining data on programs' ranking of applicants and on realized matches to identify programs' and students' preferences in a two-sided matching model. This strategy circumvents the need for exclusion restrictions in disentangling preferences of agents on both sides of the market highlighted in Agarwal and Somaini (2022) and He et al. (2023). I propose a tractable estimation procedure. I use a Maximum Likelihood strategy to estimate individuals' preferences over locations and programs' preferences over applicants, before estimating students' preferences over programs using a Simulated Method of Moments strategy with optimal instruments.

Parameter estimates show that students' demand is highly elastic to distance. The average student would need to receive a monthly scholarship of 1,102 euros to agree to enroll 100 kilometers away from her home location, which corresponds to 75% of the average entry-level wage. Mobility frictions, rather than programs' characteristics, are the main component of the value associated with a program for students from higher education deserts. Regarding local labor market decisions, I find that individuals display a large preference for their home location. Yet, the magnitude of moving costs for individuals graduating out of their home location on average outweighs their preference to return home. Overall, model predictions closely match both enrollment and work flows between education hubs and deserts.

I use the estimated model to quantify how mobility frictions interact with the geography of the higher education market to generate spatial inequalities. Holding constant the distribution of programs, I simulate individuals' education and labor market choices in a counterfactual scenario where mobility frictions are eliminated. This breaks the link between their demand and the local supply they face. On the higher education market, predicted sorting depends both on how students shift their demand, as well as on how programs select who to admit given the new pool of applicants. The dynamic nature of the model also enables to simulate counterfactual location choices at entry on the labor market. Overall, this allows us to quantify changes in educational attainment and welfare between students who grew up in different locations, as well as in skills and economic activity across regions.

I find that the uneven distribution of programs across space severely constrains students' human capital investments and welfare. One-third of the spatial gaps in the share of individuals graduating with at least a bachelor's degree is explained by mobility frictions within the higher education market. Eliminating these frictions, however, generates a trade-off. It would lead to a 38% decrease in welfare inequalities between students from higher education hubs and deserts, but would magnify regional inequalities by 10%. This is driven by a brain drain, characterized by a 14% increase in the number of individuals from deserts working in hubs.

I perform additional quantification exercises. I show how modeling both programs' admission decisions and entry workers' location choices is key in highlighting this trade-off. Furthermore, I eliminate both mobility frictions within the higher education market and workers' moving costs between their study and home locations. As individuals can now return home at no cost, inequalities between both students and local labor markets decrease. Yet, students from low-opportunity areas accumulate less human capital as they anticipate that they will likely enter their home labor market, where returns to education are lower.

Finally, I explore the role of policies in circumventing the above trade-off. I show how low-opportunity areas can leverage regional trade in education services to invest in their local human capital. I design contracts tying a mobility scholarship to incentives to return home after graduation, whereby students failing to come back have to repay the full scholarship. For a cost equivalent to 6% of entry-level workers' total wages, such a policy would decrease inequalities across local labor markets and between students by 17% and 10%, respectively.

Related Literature This paper contributes to several strands of the literature. First, it relates to the literature studying the geography of opportunities (Chetty et al., 2014; Chetty et al., 2018b) by underlining the role of differences in the local supply of higher education options. The existence of educational deserts has been documented in the United States (Hillman, 2016). Studies in Belgium (Kelchtermans and Verboven, 2010), England (Gibbons and Vignoles, 2012), and the United States (Fu et al., 2022; Ishimaru, 2020) show that students' demand for higher education programs is sensitive to proximity. Fu et al. (2022) study how the spatial distribution of four-year colleges impacts students' welfare, using survey data on applications and enrollment decisions. Ishimaru (2020) uses NSLY data to additionally model migration decisions on the labor market, and explore the mechanisms linking individuals' childhood location and economic outcomes. I contribute to this literature by linking inequalities in opportunities to spatial gaps in educational attainment and spatial skill sorting. I use data on the universe of students and programs to model the two-sided nature of the higher education market and account for potential equilibrium effects triggered by counterfactual policies.

Second, this paper contributes to our understanding of regional inequalities. Spatial skill sorting is a key driver of these disparities (Combes et al., 2008; Dauth et al., 2022; Card et al., 2023). The urban economics literature primarily focuses on studying sorting through workers' migration decisions (Diamond, 2016). This paper highlights the importance of an overlooked factor, 'sorting at birth' (Combes et al., 2012). In particular, it underlines that spatial inequalities in higher education opportunities are key in generating the latter.

Third, this paper relates to the literature evaluating place-based policies designed to mitigate inequalities across local labor markets. It typically investigates the role of firms' subsidies (Busso et al., 2013; Austin et al., 2018; Ku et al., 2020) and local infrastructure investments (Kline and Moretti, 2014). As human capital is portable, it remains unclear whether regions would benefit from investing directly in the production of local skills (Bound et al., 2004). The construction of higher education institutions is associated with increases in education and local economic development (Nimier-David, 2022; Howard et al., 2022). In contrast, Hsiao (2023) uses a spatial equilibrium model to show that a school construction program in Indonesia increased inequalities across urban and rural regions. In this paper, I instead take the spatial distribution of education options as fixed, given that the construction of colleges is limited in practice and could lead to inefficiencies due to potential agglomeration and peer effects, and study the potential for moving subsidies to increase local human capital. Kennan (2015) similarly takes as given this distribution to evaluate how tuition reductions would affect the supply of college graduates across states in the United States by extending the migration model of Kennan and Walker (2011). In this paper, as I find that mobility frictions are a key driver of students' human capital investments, I quantify the effects of subsidizing students' moving while taking into account the local supply they face at a very granular level. I show that mobility subsidies would increase regional trade in tertiary education services. When tied to incentives to return home, this would enable low-opportunity areas to outsource the education of their local labor force to higher education hubs.

Finally, I contribute to the literature on empirical models of two-sided matching. A couple of recent papers show that exclusion restrictions are necessary to identify preferences of agents on both sides of the market from data on realized matches only (Agarwal and Somaini, 2022; He et al., 2023). I show how to avoid such assumptions by exploiting additional data on programs' rankings of applicants and provide a tractable estimation procedure.

Overview The rest of the paper is organized as follows. Section 1.2 introduces the context and data. Section 1.3 describes the geography of the higher education market and explores the effects of mitigating barriers to students' mobility. Section 1.4 presents the model. Section 1.5 discusses the identification and estimation of agents' preferences. Estimation results and model fit are presented in Section 1.6. Section 1.7 describes counterfactual results. Finally, Section 1.8 concludes.

1.2 Context and Data

This section provides an overview of the French higher education market, before detailing the college admission procedure. Finally, it describes the data sources used throughout the paper.

1.2.1 The French Higher Education Market

Education in France is organized in a centralized manner, with the higher education system being managed by the Ministry of Higher Education and Research. It is administrated through thirty catchment areas (Figure 1.A.1), implementing the directives of the Ministry within their area.

To access the higher education system, students have to take a national exam in their last year of high school.[1] After this exam, and conditional on receiving a passing grade, students can choose to pursue their education by enrolling in a higher education program. There are six main types of programs available, half of them targeting students willing to acquire a master's degree — they are the equivalent of four-year colleges in the United States. First, students can apply to three-year university programs, with the aim of obtaining a bachelor's degree — after three years — and a master's degree, after five years of education. Second, they can apply to two-year elite programs. These programs are highly selective and prepare students to take competitive exams for admission to the most prestigious public and private higher education institutions. Once enrolled in these elite schools, students will then typically continue their education and earn a master's degree. Third, there are also programs offered by engineering and business schools.

[1] Appendix 1.B.1 provides more details about the high school system in France.

Additionally, there are three main types of short-term programs, providing more technical and vocational training. First, students can enroll in technology-oriented programs tailored to train mid-level technical professionals. Second, there are two-year vocational programs, with a stronger emphasis on internships. Finally, students can also apply to programs offered by social work and nursing schools, which train workers in two to three years. Details can be found in Appendix 1.B.2.

1.2.2 College Admission Procedure

Since 2009, the higher education admission procedure is centralized and operates on an online platform.[2] Every year, students in their final year of high school register on this platform. An online search engine is available: students can specifically search for programs in a given major or location. A map allows students to locate each institution offering a certain major they may be interested in. A general presentation of each program can be found on the platform, as well as a description of the admission criteria and information regarding the number of seats and the number of applicants in the year before. Figures 1.A.2 to 1.A.7 present this interface, as well as the type of information students have access to.

After searching for programs, students submit their applications on the online platform. Programs consist of an institution-type-major combination. For instance, a student wishing to be considered for admission to a university program in either economics or mathematics at the Sorbonne University needs to submit two separate applications. In theory, the number of applications is limited to ten. However, the effective number of applications allowed is substantially larger in practice as, in many cases, several applications are processed as a single one. For example, listing up to ten different institutions offering a two-year elite program in physics would only count as one application. In the selected sample, which I define below, the mean number of applications is 10.6, and 25% of students submit 13 applications or more (Table 1.A.3).

Once applications are submitted, each program observes its set of applicants and has to rank them. Applicants are ranked based on the information they and their high school submit on the platform. In particular, programs can access students' high school grades. They can also require students to submit a resume or a cover letter on the platform. Programs do not have to disclose how the different characteristics they observe are weighted to determine their

[2] See Appendix 1.B.3 for a history of the admission system.

ranking: programs' preferences are thus unobserved. Interviews of members of admission committees, however, reveal that programs typically first choose the weights to assign to each observed characteristic.[3] These weights are then fed into an algorithm generating the ranking of applicants. The use of cover letters and resumes seems less common.

The ranking provided by programs is then modified to take into account some equity-based admission policies chosen by policymakers ex-ante. In particular, a minimum share of low-income students — defined by their scholarship eligibility — is targeted for each program. There are also place-based quotas for university programs, which specify the maximum share of out-of-catchment area students who can be admitted.

Once programs have finalized their ranking of applicants, the admission stage starts. Offers are released to students at the top of the programs' ranked-ordered list, up to the number of seats available. Students can thus receive several offers. When this is the case, they can accept one offer and decline the others. Rejected offers are sent dynamically to applicants remaining unmatched, following the initial ranking of each program. Students can also decide to delay their decision by holding on to their current offers, waiting for an offer from a program they prefer. Students observe both their rank on the waitlist and the rank of the last admitted student the year before, helping them predict whether they will receive an offer.

1.2.3 Data

I assemble a unique dataset, gathering the actions of the universe of applicants and programs on the college admission platform in France and tracking students until they exit the higher education system. I complement this information with a survey of individuals at entry on the labor market to understand individuals' location decisions. This section describes these different datasets.

College Admission System: First, I use administrative data from the college admission platform in France. It gathers information on the universe of college applicants each year. It includes a very rich set of students' characteristics, including their gender, the socio-economic status of their parents, their scholarship status, as well as their test scores. Importantly, I

[3]See 'Parcoursup : enquête dans les coulisses des commissions d'examen des vœux', Le Monde, Séverin Graveleau, Soazig Le Nevé et Eric Nunès. Published on May 23^{rd}, 2023. Last accessed: August 18^{th}, 2023.

also observe the precise location of residence of each student, at the ZIP code level — there are 37,901 ZIP codes in France. Table 1.A.1 presents the average characteristics of the main sample of students considered throughout the paper. I focus on students living in metropolitan France and who are in high school at the time they apply. This represents a sample of 505,069 students in 2019.

This administrative dataset also gathers the universe of programs available on the platform. I observe the characteristics of each program, including its type, major, the number of available seats, and its GPS coordinates such that the distance between each program and student on the platform can be computed. I restrict attention to the 10,917 programs located in metropolitan France, excluding online programs. Table 1.A.2 describes the average characteristics of this sample of programs.

Finally, this dataset covers in detail the application and admission steps. It gathers all applications submitted by each student. It also records the ranking of applicants by programs. It tracks the offers sent, students' acceptance and rejection decisions, and their final program of enrollment. Tables 1.A.3-1.A.4 describe applications, offers and enrollment patterns.

Higher Education Census: Second, I use an administrative dataset recording the path of each student enrolled in the French higher education system until their exit. As this dataset can be merged with the one described above, I can track students' path from their entry in the higher education system to the time they graduate or drop out. I observe whether students switch programs, the highest degree they obtain, and their province of graduation.

Location Choices and Earnings: I complement these datasets with a survey of individuals, following students after they exit the French education system. I restrict attention to students who graduated from high school. This dataset records the home, graduation, and work locations of each individual, such that their migration decisions are observed. The survey also reports individuals' wages within the three years after they left the education system. Table 1.A.5 describes the work flows, study flows, and earnings of these individuals.

1.3 Descriptive Evidence and Policy Evaluation

In this section, I first provide descriptive evidence linking the geography of the higher education market to spatial inequalities in educational attainment and spatial skill sorting.

Figure 1.1: Distribution of the Supply of Academic Seats and High School Student Population

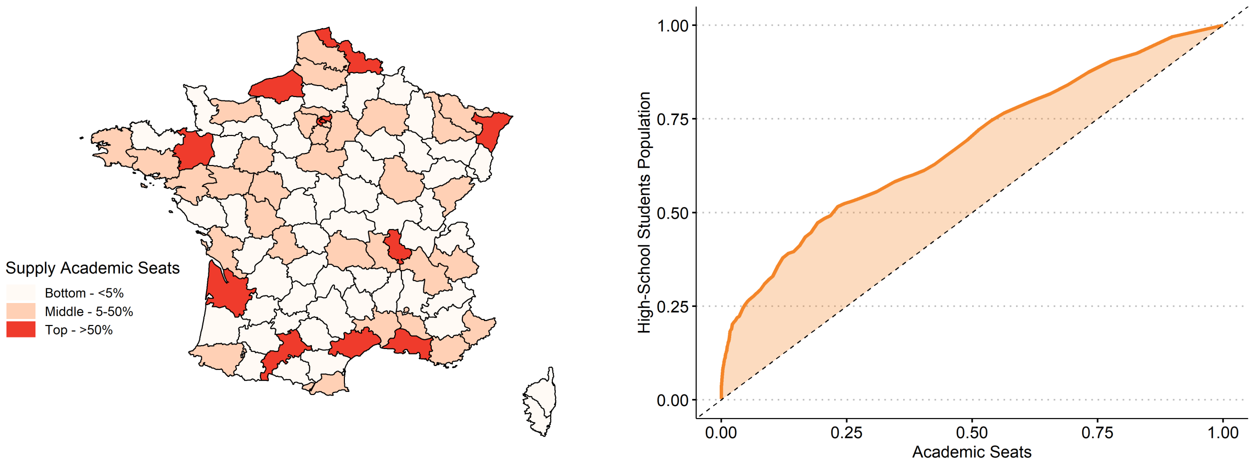

a) Spatial Distribution of Academic Seats **b) Student-Seat Mismatch**

NOTES. Panel a) presents a map of the distribution of academic seats in metropolitan France. Provinces are ranked according to their position in the distribution of academic seats. Provinces in white are at the bottom of the academic seats supply distribution: they jointly account for less than 5% of the total supply. In contrast, provinces in red are at the top of this distribution: half of the academic seats supply is concentrated there. Panel b) plots the relationship between the spatial distributions of academic seats and of high school students. It depicts the cumulative distribution of the supply of academic seats across provinces on the x-axis, and the corresponding cumulative distribution of the high school student population on the y-axis. Provinces are ranked according to the share of academic seats they offer. For example, provinces gathering 25% of the total supply of academic seats concentrate 50% of the population of high school students.

Second, I use local policies mitigating mobility frictions to explore their impact on students' demand, as well as their distributional effects.

1.3.1 Spatial Distribution of Colleges and Inequalities

I first document that the distribution of programs across space is uneven, and does not match the spatial distribution of the population of students. Second, I show that — consistently with the presence of (monetary and non-monetary) mobility costs — students' demand for higher education programs is elastic to geographic proximity. Overall, these patterns correlate with observed spatial gaps in enrollment choices and skill sorting.

Fact 1: Higher Education Programs are Unevenly Spatially Distributed

The spatial distribution of higher education programs is highly uneven. I focus on academic programs, combining university and two-year elite programs, which are the main paths to a bachelor's degree. Panel a) of Figure 1.1 maps the spatial distribution of first-year academic seats at the province level. Of the 96 provinces of metropolitan France, ten do not offer a

Table 1.1: Number of Programs Around Students' Home Location

	Quantile			
	10^{th}	25^{th}	50^{th}	75^{th}
	(1)	(2)	(3)	(4)
Panel A: 10 kilometer-radius				
Three-Year University Program	0	0	0	29
Two-Year Elite Program	0	0	0	10
Engineering & Business Schools	0	0	0	6
Two-year Technical Programs	0	0	1	10
Two-year Vocational Programs	0	2	17	63
Nursing Schools	0	0	1	8
Panel B: 40 kilometer-radius				
Three-Year University Program	0	7	35	115
Two-Year Elite Program	0	3	11	34
Engineering & Business Schools	0	2	7	25
Two-year Technical Programs	2	6	13	29
Two-year Vocational Programs	30	53	93	268
Nursing Schools	1	4	9	24

NOTES. Number of programs, by program types, located within a given radius around students' home location. Panel A considers a ten-kilometer radius, while Panel B presents results for a forty-kilometer radius. I consider 40 kilometers as the threshold defining students' commuting zone, to the extent that 96% of the daily commuting trips in France are 40 kilometers or less. This table shows that 10% of the students have no three-year university programs located within a forty-kilometer radius around their home location.

single academic seat.[4] Overall, half of the provinces jointly account for less than 5% of the total supply of academic seats. I refer to these areas as higher education deserts. In contrast, 50% of the total supply is agglomerated in only 12 provinces. Overall, the supply of academic seats is thus highly concentrated.

This geography of academic seats could, however, merely reflect the population distribution of students over space. Yet, Panel b) of Figure 1.1 illustrates that the spatial distribution of academic programs is far from mirroring that of high school students. Provinces concentrating less than 5% of the total supply of academic seats gather 25% of the high school student population. In contrast, only 28% of students live in one of the twelve provinces accounting for more than 50% of the total supply. There is thus an imbalance between the spatial distributions of academic seats and students.[5]

[4] In total, 17 provinces offer no university programs, while 18 provinces do not have any two-year elite programs.

[5] This imbalance contrasts greatly with the number of seats in two-year technological and vocational programs, which is considerably more evenly distributed across space (Figure 1.A.9) and thus better matches the spatial distribution of high school students (Figure 1.A.10).

As a result, a substantial share of students face very a low supply of academic programs close to their home location. Table 1.1 presents the distribution of the number of programs, by type, located within a 10- and 40-kilometer radius around students' home. I consider 40 kilometers as the threshold defining students' commuting zone, to the extent that 96% of the daily commuting trips in France are 40 kilometers or less.[6] I find that 50% of students do not have a single university, two-year elite, engineering or business program within ten kilometers from their home. Enlarging the radius to include any programs located within commuting distance, we observe that more than 10% of applicants cannot find any university or two-year elite programs within this boundary. A quarter of the students find at most seven university programs within commuting distance.[7] As each program corresponds to a specific major, this implies that these students have very few different majors available close to their home location, as illustrated by Table 1.A.6. For example, 10% of the students have to travel at least 92 kilometers to reach a medicine program.

Fact 2: Students' Demand is Elastic to Geographic Proximity

The concentration of academic programs across space would have no consequences on students' college choices unless they face some monetary or non-monetary moving costs. Such costs could include expenses for traveling to the study location, rent, as well as the emotional toll of being far away from family and friends. Students' enrollment patterns suggest that such costs are present and play a significant role in their higher education choices. As shown in Figure 1.2, 25% of students enroll less than 8.5 kilometers away from their home. The median distance of enrollment is only 25 kilometers. This pattern is not generated by the fact that students only receive admission offers from programs located very close to their home location. On the contrary, the distribution of the distance to the offers received (purple line) first-order stochastically dominates that of the distance to accepted offers (orange line). The median offer is located 20 kilometers farther away than the accepted one. This gap increases to 50 kilometers at the third quartile of the distribution.

[6] According to the same study, in France, the average individual only makes 6.3 trips at more than 80 kilometers from their home each year. Source: SDES, Insee – Enquête Mobilité des Personnes 2018-2019. Data: 'La mobilité locale et longue distance des Français - Enquête nationale sur la mobilité des personnes en 2019', accessed at statistiques.developpement-durable.gouv.fr. Last accessed: October 9^{th}, 2023.

[7] Figure 1.A.8 provides a map of the mean number of academic programs located within commuting distance of students' home location at the province level. The mean number of academic programs located within a 40-kilometer radius around students' home location is less than ten in 24 provinces.

Figure 1.2: Cumulative Distribution - Distance of Enrollment & Offers

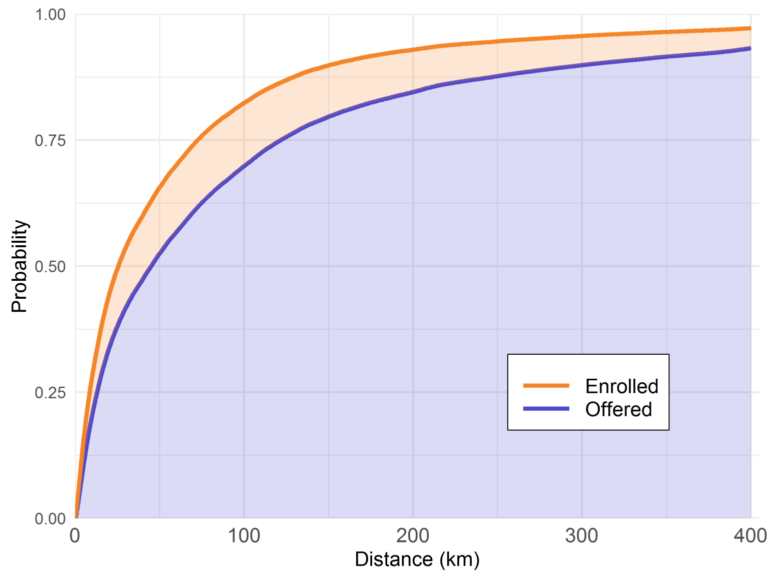

NOTES. This figure shows the cumulative distribution function of the distance (in kilometers) between students and the program where they enroll (orange) and between students and the programs for which they receive offers (purple).

This suggests that students favor programs closer to their home location, consistently with the existence of mobility frictions. To measure the elasticity to distance of students' enrollment patterns, I borrow from the migration and trade literature and estimate a gravity model. I regress the flow of students from a given province enrolling in programs located in each destination province, controlling for the distance between the centroids of the two locations and including origin and destination fixed effects, capturing differences across locations, such as local college supply.[8] I find an elasticity of enrollment to distance equal to -2.5 (Table 1.A.7), which suggests that college choices are highly sensitive to distance from home.[9]

Implications: Spatial Inequalities

Overall, the interaction of mobility frictions with the uneven distribution of higher education programs is likely to generate important inequalities between students from higher education hubs and deserts. Panel a) of Figure 1.3 highlights the existence of a correlation between the share of students from a given province enrolled in an academic program and the local

[8]I estimate this model using a Pseudo-Poisson Maximum-Likelihood Estimation strategy. See Silva and Tenreyro (2006) on the advantages of using this approach to estimate gravity models.

[9]This elasticity is about 30% larger than that found for labor migration in France (Zerecero, 2021).

Figure 1.3: Higher Education Market Geography and Spatial Inequalities

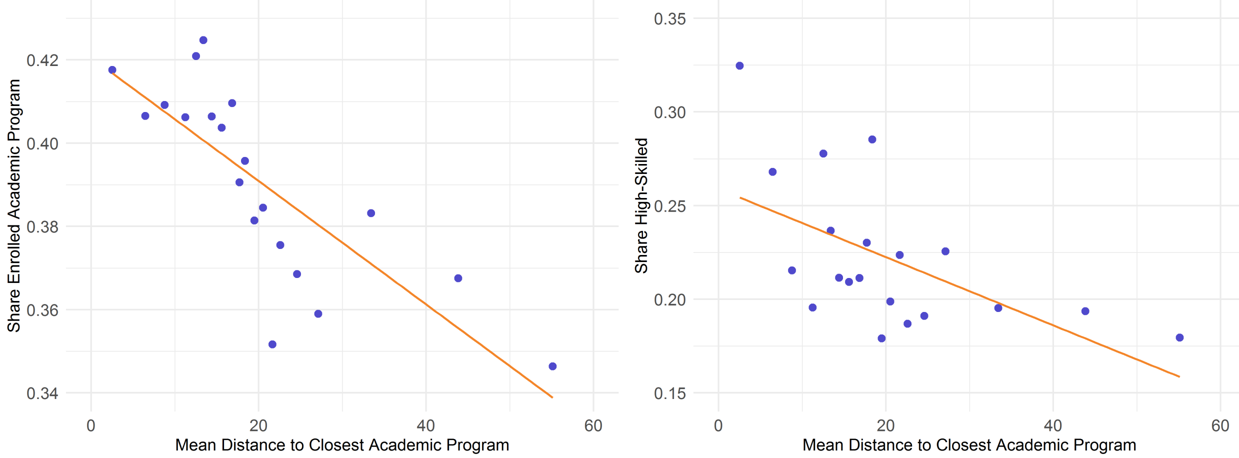

a) Spatial Gaps in Academic Enrollment **b) Spatial Skill Sorting**

Notes. Panel a) is a binned scatter plot, at the province level, of the relationship between the share of students from a province enrolled in an academic program and the mean distance (in kilometers) to the closest academic program for the average student of the province. Figure 1.A.11 shows that a similar relationship holds for specific majors, such as medicine programs. Panel b) is a binned scatter plot, at the province level, of the relationship between the share of the total labor force of a province holding a three-year degree or more (share high-skilled) and the mean distance (in kilometers) to the closest academic program for the average student of the province.

availability of such programs.

Additionally, as individuals often choose to work in their home region — given moving costs and home bias (Kennan and Walker, 2011) — the previous correlation is likely to contribute to the skill sorting observed across local labor markets. We indeed observe a similar relationship between the share of a province's labor force with at least a three-year degree and the local availability of academic programs (Panel b) of Figure 1.3).

Both these correlations could be driven by other factors than mobility frictions — such as differences in students' pre-college human capital, preferences or returns to education across provinces. In Section 1.4, I build a model accounting for these other potential channels, allowing us to isolate how mobility frictions interact with the geography of the higher education market to generate spatial inequalities between students and across local labor markets.

1.3.2 Mitigating Mobility Frictions: Changes in Demand and Distributional Consequences

The above patterns suggest that mobility frictions within the higher education market generate spatial inequalities. To mitigate the latter, it would be key to break the dependence between students' application decisions and their local supply of programs. In this section, I first show that moving subsidies allow students to apply farther away from their home location. Second, I explore the distributional effects of alleviating barriers to mobility, both within the higher education market and across local labor markets.

Mobility Scholarship and Changes in Students' Demand

Studying away from home entails direct expenses — such as travel costs and rents — which are likely to be an important component of mobility frictions. Could moving subsidies allow students to apply farther away from their home location? To answer this question, I evaluate a moving scholarship introduced in 2019, guaranteeing 500 euros to low-income students conditional on enrolling in a program located outside their catchment area. This amount covers a fraction of the costs associated with such a move, corresponding to approximately 9% of the annual budget of the average student.[10]

This program is only available to recipients of a means-based scholarship in high school, as a way to target low-income students. Using applications data from 2015 to 2019, I can thus compare the change in the application distance of eligible and non-eligible students, before and after the implementation of the policy. As several reforms of the admission system have been implemented over the years, potentially affecting differently the application behavior of students with different eligibility status, I use an additional comparison group. In particular, I compare the change in outcomes, before and after the introduction of the policy, of eligible (E) and non-eligible (NE) students living in provinces with low (L) and high (H) supply of higher education programs.[11] As they face many options close to their home location, the latter group is unlikely to be impacted by the reform. I estimate the following model:

$$d_{i,e,s,t} = \alpha_{s,t} + \alpha_{e,t} + \alpha_{s,e} + \sum_{\tau=-3,\tau\neq 0}^{1} \beta_{2018+\tau} D_{e,s,2018+\tau} + \epsilon_{i,s,e,t} \tag{1.1}$$

where $d_{i,e,s,t}$ is the distance, in kilometers, between a student and the program to which she submits application i, in year t, while being of eligibility status $e \in \{E, NE\}$ and living in a province with supply level $s \in \{H, L\}$. $\alpha_{s,t}$ are supply level-year fixed effects controlling for the fact that changes in outcomes may have been different for students living in high-

[10] Source: Observatoire National de la Vie Etudiante.

[11] I define as high-supply areas the ten provinces at the top of the supply distribution. Low-supply areas are defined as the bottom five provinces of the supply distribution.

Figure 1.4: Effect of the Spatial Mobility Scholarship on Application Distance

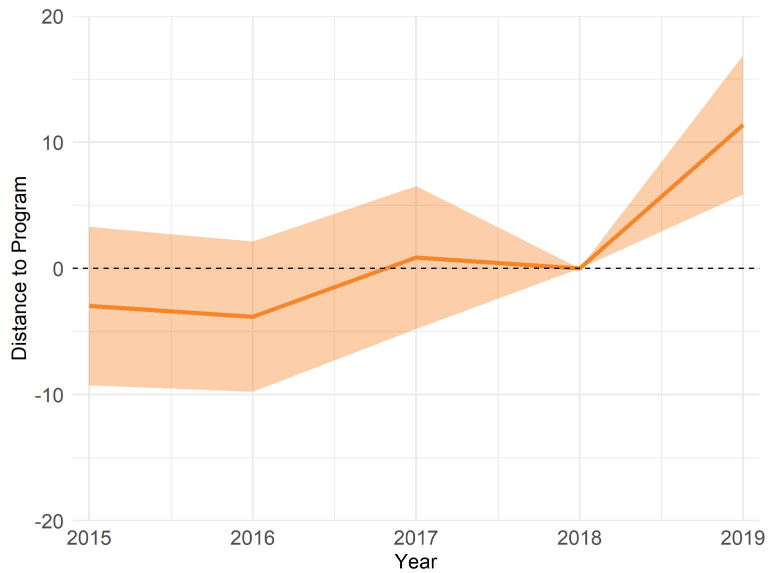

NOTES. This figure presents parameter estimates of $\beta_{2018+\tau}$ from Equation (1.1). It shows the effect of the policy offering moving subsidies, implemented in 2019, on the mean application distance, in kilometers. 95% confidence intervals are shown, computed using robust standard errors. These estimates are displayed in Table 1.A.8.

or low-supply areas, regardless of their eligibility status. $\alpha_{e,t}$ are eligibility status-year fixed effects. They control for differences in trends between eligible and non-eligible students that are common across areas. Finally, $\alpha_{s,e}$ capture application behaviors for students with a given eligibility-location status that are constant over time. $D_{e,s,2018+\tau}$ is a dummy equal to one for students eligible for the scholarship and living in a low-supply area in year $2018+\tau$, for $\tau \in \{-3, 1\} \setminus \{0\}$. I let its associated parameter vary across years. $\beta_{2018+\tau}$ is thus a triple-difference estimator, as it estimates the relative treatment effects for eligible students in the low- versus high-supply areas.

The identifying assumption is that, in the absence of the treatment, the relative evolution of outcomes for eligible and non-eligible students would be the same in high- and low-supply areas (Olden and Møen, 2022). Details are provided in Appendix 1.C.1. I test this common trends assumption through placebo tests involving years prior to the policy implementation.

Figure 1.4 and Table 1.A.8 present estimates of $\beta_{2018+\tau}$ for different years before and after the implementation of the policy, in 2019. I find that the moving subsidy allowed eligible students from low-supply areas to apply to programs located farther away from their home location. The reform was associated with an 11-kilometer increase in the distance to programs to which students apply. I find no pre-trends in the years prior to the implementation of the reform. Table 1.A.8 shows that the impact of the policy is driven by an increase in the share of out-of-catchment area applications. However, this did not translate into a significant increase in the share of offers received from out-of-catchment area programs, such that enrollment patterns were not affected.

Overall, these results suggest that students are willing to apply farther away when mobility frictions are softened, for example through moving subsidies. While the latter may allow students to apply to programs they prefer, the policy evaluation above does not allow us to quantify this effect. Furthermore, applications may fail to translate into admissions, as the market is two-sided: programs choose which applicants to admit.

Distributional Effects of Alleviating Barriers to Mobility

Alleviating mobility frictions alters students' demand. Such changes may entail substantial distributional effects, especially given the two-sided nature of the market and capacity limits: students can only enroll in a program admitting them. In this section, I investigate these distributional effects. First, I use a local policy change to show that eliminating barriers to students' mobility generates a brain drain toward higher education hubs in the short run. I then use survey data to illustrate how, if happening at the scale of the country, such a brain drain could be long-lasting as there exists stickiness in location choices.

Distributional Effects — Higher Education Market I first focus on the short run and present evidence suggesting that removing barriers to mobility is associated with a brain drain of students toward regions concentrating high-quality programs. To show this, I zoom in on the Greater Paris area, which is composed of three catchment areas: Paris, Versailles and Créteil.[12] While several universities are present in the region, Paris concentrates high-quality institutions, attracting students from the entire area. Geographic priorities — akin to quotas used by public universities in the United States, capping the share of out-of-state admitted students — were, however, limiting the admission chances of students living outside of Paris, up to 2018.[13] Table 1.A.10 shows how these priorities affected the final ranking of students: on average, a student from Versailles applying to a program in Paris was down-ranked by

[12] Figure 1.A.13 presents a map of the area and Table 1.A.9 shows summary statistics of programs located in the region.

[13] In the United States, in 2017, the University of California for example implemented an 18% cap on out-of-state enrollment. See 'A bold plan for UC: Cut share of out-of-state students by half amid huge California demand', Los Angeles Times. Published on May 25^{th}, 2021. Last accessed: October 10^{th}, 2023.

483 positions after their implementation, while the average rank of the last-admitted student was 1,000. In 2019, to equalize the playing field, these geographic priorities were removed.

I evaluate how this reform impacts student sorting within the higher education market. In particular, I consider the following model:

$$y_{j,\ell,t} = \alpha_j + \sum_{\ell} \beta_\ell D_{\ell,t} + \epsilon_{j,\ell,t} \tag{1.2}$$

where $y_{j,\ell,t} = \frac{n^H_{j,\ell,t}}{n_{j,\ell,t}}$ is the share of high-ability students admitted by program j from location $\ell \in \{\text{Paris, Versailles, Créteil}\}$ in year t, with $n_{j,\ell,t}$ ($n^H_{j,\ell,t}$, respectively) the number of students (high-ability students, respectively) admitted by j in period t. High-ability students are those graduating from high school with honors or with the highest honors. I also consider the share of high-ability applicants to program j to investigate direct changes in demand. α_j are program fixed effects, capturing constant differences in $y_{j,\ell,t}$ across programs. $D_{\ell,t}$ is a dummy variable equal to one for university programs in location ℓ and in $t = 2019$, i.e. after the reform. I allow the reform to impact differently programs according to their location.

Table 1.2 presents parameter estimates from Equation (1.2). I find that this reform generated a brain drain of students from peripheral regions toward programs located in Paris. In particular, the share of high-ability students from Versailles enrolling in programs in Paris increased by 50% after the reform (Column (1)). In contrast, the share of high-ability students from Versailles enrolling in their local universities decreased by 25%. These students thus enrolled in universities located in Paris instead of their local institutions, revealing a brain drain.[14]

As the number of seats is limited, this brain drain had further consequences for overall sorting. Indeed, students from the peripheral areas crowded out low-ability students from Paris from their local universities: their share decreased by 31%, from a baseline of 22% when they were receiving priority (Column (3)). In turn, the share of low-ability students from Paris enrolled in universities located in the peripheral areas significantly increased.

Table 1.A.11 shows that these enrollment effects were largely driven by programs being able to express their preferences without being subject to geographic priorities anymore, rather than by changes in demand spurred by the reform. The change in the share of

[14] We do not observe such movements for students from Créteil. This is consistent with the presence of geographic variation in students' schooling achievement: the share of students classified as high-ability is 6.5 points larger in Versailles than in Créteil.

Table 1.2: Removing Barriers to Mobility: Sorting Effects

	Share High-Ability		Share Low-Ability
	from Versailles	from Créteil	from Paris
	(1)	(2)	(3)
β_{Paris}	0.025***	0.009	-0.069***
	(0.007)	(0.007)	(0.015)
$\beta_{\text{Versailles}}$	-0.029***	-0.000	0.022***
	(0.010)	(0.003)	(0.001)
$\beta_{\text{Créteil}}$	0.003	-0.006	0.024***
	(0.002)	(0.004)	(0.006)
Baseline Paris	0.049	0.048	0.220
Baseline Versailles	0.114	0.013	0.030
Baseline Créteil	0.009	0.053	0.047
Program FE	Yes	Yes	Yes
Observations	680	680	680

NOTES. This table presents estimates from Equation (1.2). I study the effect of removing barriers to mobility on the share of high-ability students from Versailles (Column (1)), the share of high-ability students from Créteil (Column (2)) and the share of low-ability students from Paris (Column (3)) enrolled in programs located in Paris, Versailles and Créteil. Standard errors clustered at the program level are shown in parentheses. *p<0.1; **p<0.05; ***p<0.01.

applications from high-ability students from Versailles to Paris is much smaller than the observed enrollment increase. This indicates that programs' admission decisions are key in shaping the equilibrium of the higher education market.

These results are robust to accounting for other factors that may have affected the higher education market at the same time (Tables 1.A.12-1.A.13). In particular, I use non-university programs, which never faced geographical restrictions, as a control group (see Appendix 1.C.2). This is akin to implementing a difference-in-differences where the groups compared share the same treatment status at endline, also known as reverse difference-in-differences (Kim and Lee, 2019). Fabre (2023) discusses the identification of treatment effects in such alternative difference-in-differences designs. Overall, these results suggest that alleviating mobility frictions may have important effects on students sorting within the higher education market.

Distributional Effects — Local Labor Markets If happening at the scale of the country, such a brain drain within the higher education market may, in turn, have long-lasting consequences for regional inequalities. Using data recording individuals' migration flows between their home, study and work locations, I indeed find that a substantial share of stu-

Figure 1.5: Contribution of the Student Population to the Non-Local Labor Force

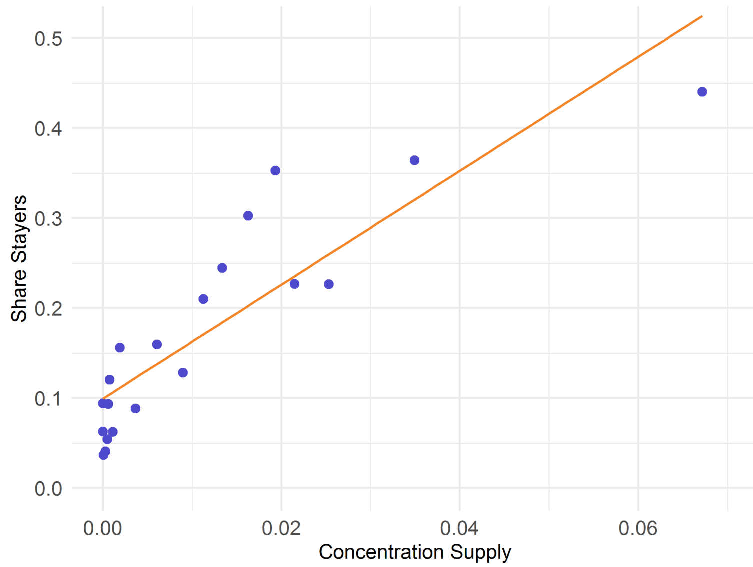

Notes. This binned scatter plot highlights the correlation between the concentration of the supply of academic programs and the capacity of provinces to retain students upon graduation. The x-axis measures the concentration of the supply in a given province through the share of academic seats located there. The y-axis displays the share of the non-local labor force of a province who graduated there.

dents graduating outside their location of origin enter the labor market elsewhere. Table 1.A.5 shows that, among the 51% of students graduating outside their home location, more than half do not return at entry on the labor market. In contrast, only 18% of individuals studying in their home province enter the labor market in a different location.

Conditional on studying outside their home location, 17% of individuals enter the labor market in their study location. Location choices thus seem to display stickiness, with individuals having a high probability of contributing to the labor force of their study location. Higher education hubs benefit from this, as they can retain this population, as shown in Figure 1.5.

Overall, the extent to which mobility frictions generate inequalities across students living in different locations depends on two elements. First, it hinges upon students' preferences. The scholarship evaluated above hints at the fact that subsidizing moving costs could change students' applications, but does not allow us to pin down how students trade-off proximity and other programs' characteristics. Second, evidence shows that programs' admission decisions also matter. Initial spatial inequalities, such as differences in school quality and test scores, could then limit admission chances of students from low-opportunity areas, regardless of the geography of the higher education market.

Regarding local labor markets, mobility frictions may act as a constraint on the human capital investment of the local population of higher education deserts. Yet, the above evidence suggests that such frictions could in fact limit the brain drain of workers from low- to high-opportunity areas, as students who leave to study do not necessarily return home upon graduation. This choice ultimately depends on preferences over locations' characteristics — such as wages and amenities — and moving costs.

In what follows, I build a framework flexibly incorporating the above channels. It includes a rich model of students' preferences over programs, accounting for mobility frictions. Importantly, students are forward-looking: they take into account how their enrollment choice may affect their location decision at entry on the labor market, through moving costs and location-specific amenities and returns to education. The higher education market is at equilibrium and takes into account programs' preferences over students. Overall, the model allows us to isolate how mobility frictions interact with the geography of the higher education market to generate spatial inequalities, both between students and across local labor markets.

1.4 Model

In this section, I build a dynamic model combining an equilibrium model of the higher education market, where students are forward-looking, with a model of location choices at entry on the labor market. I consider a two-stage framework, where the timing is as follows. In the first period, a set of high school students $\mathcal{I}$ and a set of higher education programs $\mathcal{J}$ meet on a centralized platform, where they can express their preferences over one another. A match between students and programs is formed on the platform. In the second stage, individuals exit the education system with a certain degree and choose the location in which to enter the labor market, from a set of locations $\mathcal{L}$.

In what follows, I first describe agents' preferences at each stage of the model. I then discuss how they map into choices and outcomes by detailing the equilibrium assumptions.

1.4.1 Preferences

I present the model backward, starting with payoffs from individuals' location decisions at entry on the labor market. I then describe students' and programs' payoffs within the higher education market. Finally, I discuss the assumptions I impose on preferences.

Labor Market Location Decision: Payoffs

The second stage of the model takes place after graduation: individual $i \in \mathcal{I}$ chooses which local labor market $\ell \in \mathcal{L}$ to enter. The utility of individual i from choosing a given local labor market ℓ depends on the wage she would receive in this location, as well as other location attributes. This utility varies according to two types of individual's characteristics, h_i and y_i. The first vector collects variables determined before the individual enters the higher education system. The second object is a vector of state variables, which are influenced by the enrollment decisions of the student in the first stage, such as the degree obtained and the location at graduation. The utility from entering the labor market in location ℓ for individual i then writes:

$$\mathcal{W}_{i\ell} = w_{i\ell} + \omega(h_i, y_i, m_\ell) + \frac{1}{\theta}\nu_{i\ell} \tag{1.3}$$

where $w_{i\ell}$ is the expected monthly wage of individual i in location ℓ, and $\omega(h_i, y_i, m_\ell)$ is the value that an individual with characteristics (h_i, y_i) would receive from location ℓ's attributes, m_ℓ, such as its amenities. $\nu_{i\ell}$ is the idiosyncratic shock of individual i for location ℓ. θ controls the relative importance of the deterministic component of the utility relative to the shock.

Higher Education Market: Students' and Programs' Payoffs

The first stage of the framework is an equilibrium model of the higher education market, taking into account its two-sided nature. Indeed, students cannot simply choose in which program to enroll, as programs have capacity constraints. I consider a market where students and programs meet and express their preferences over one another on a centralized market. Before defining the equilibrium concept, I first formalize the payoffs derived by students when enrolling in a given program, and those of programs when admitting a given student.

Students' Preferences Over Programs On the platform, students observe program j's characteristics, $z_j, \forall j$. The payoff derived by student i when enrolling in program j is assumed to write as follows:

$$\mathcal{U}_{ij} = u(x_i, z_j) - c(x_i, d_{ij}) + \sigma\epsilon_{ij} + \bar{U}(x_i, h_i, z_j) \tag{1.4}$$

This payoff has two parts. First, students derive a flow utility from enrolling in a given program. $u(x_i, z_j)$ is the value from enrolling in program j, and is a function of student i's characteristics, x_i, and program j's characteristics, z_j. I detail in Section 1.5.1 how x_i and h_i may differ. $c(x_i, d_{ij})$ captures mobility frictions associated with enrolling to program j, as a function of student i's characteristics and the distance between i and j. I denote the outside option $0_{\mathcal{J}}$, and normalize to zero $u(x_i, z_{0_{\mathcal{J}}})$ and $c(x_i, d_{i0_{\mathcal{J}}})$.[15] ϵ_{ij} is an idiosyncratic preference shock. The parameter σ controls the scale of this shock.

Second, students are assumed to be forward-looking: they take into account that their choice on the higher education market will influence their outcomes in stage two, i.e. when they choose which local labor market to enter. This captures the fact that students anticipate that moving in the future may be costly, which makes labor market opportunities in their study location, as well as amenities, key when choosing where to enroll. Overall, $\bar{U}(x_i, h_i, z_j)$ captures the discounted expected value of choosing the optimal location in stage two for a student with characteristics (x_i, h_i), given the characteristics of program j. In particular:

$$\bar{U}(x_i, h_i, z_j) = \int \beta^{t(y_i)} \mathbb{E}\left[\max_{\ell \in \mathcal{L}} \left(w_{i\ell} + \omega(h_i, y_i, m_\ell) + \frac{1}{\theta} \nu_{i\ell} \right) |y_i \right] f\left(y_i | x_i, h_i, z_j\right) dy_i \tag{1.5}$$

where $f\left(y_i | x_i, h_i, z_j\right)$ captures the transition probability to graduate with state variables y_i, given student i's initial characteristics, (x_i, h_i), and the characteristics of program j, z_j. It will for example capture the probability that a student graduates with a certain type of degree, given both her characteristics and the characteristics of the chosen program. This is as in Kennan (2015), which also models graduation as the outcome of a stochastic process. Students thus know that each program is associated with a different probability of graduating with a given degree, depending on both x_i and z_j, as will be described below. β is the discount factor, and its exponent depends on the state variables of student i at graduation. It takes into account that individuals enter the labor market at different stages depending on the length of their studies. The scale of this continuation value is normalized by the relative importance of the deterministic part of the utility to the idiosyncratic shock in stage two.

As the scale of $\bar{U}(x_i, h_i, z_j)$ is fixed, due to the normalization of the wage parameter in $\mathcal{W}_{i\ell}$, I let the scale of the idiosyncratic part of the flow utility associated with program j,

[15] The outside option consists in not matching to any program available on the platform. As the latter gathers the quasi-universe of higher education programs in France, the outside option can be thought of as dropping out of the higher education system.

σ, free. Normalizing this scale to one, as is typically done in practice, would be a strong assumption in this model as it is a non-stationary framework. In particular, the choices made in the first and second stages are very different, while normalizing the scale of ϵ_{ij} to one would be equivalent to fixing the relative scale of utility in both stages one and two. I show how to identify this parameter thanks to an exclusion restriction in Section 1.5.1.

Supply Side: Program-Specific Score On the platform, programs observe applicants' characteristics, (x_i, h_i). While programs can express their preferences over applicants, they have to take into account some equity objectives imposed by the government. For example, scholarship-eligible applicants benefit from an affirmative action policy. Admission decisions are thus based on programs' preferences after they internalized the preferences of the government. These preferences can be captured by an applicant-program specific score, depending both on the student's and program's characteristics. The final score of applicant $i \in \mathcal{I}$ for program $j \in \mathcal{J}$ is defined as:

$$\mathcal{V}_{ji} = v(x_i, h_i, z_j) + \eta_{ji} \tag{1.6}$$

where $v(x_i, h_i, z_j)$ is the deterministic part of this score, depending on both students' and programs' characteristics. I normalize to zero the flow utility associated with programs' outside option, denoted $0_{\mathcal{I}}$, which captures that programs may decide to leave a seat empty. η_{ji} is an idiosyncratic preference shock.

Assumptions

I define here the regularity conditions and distributional assumptions I impose on agents' preferences. In particular:

Assumption 1

(i) The functions $\omega(.)$, $u(.)$, $c(.)$ and $v(.)$ are bounded.

(ii) $\nu_{i\ell}$, ϵ_{ij} and η_{ji} are independently distributed from a Type I Extreme Value distribution.

Assumption 1.(i) is satisfied in most settings. Assumption 1.(ii) is standard in the literature using dynamic discrete choice models, and stems from the fact that the distribution of unobserved shocks cannot be identified in these models (Magnac and Thesmar, 2002a). It however implies the independence of irrelevant alternatives property. In Section 1.5.2, I explore how to relax this assumption by adding unobserved heterogeneity.

1.4.2 Equilibrium

In this section, I characterize the equilibrium of the model. I start by describing individuals' optimal behavior when they enter the labor market. I then detail programs' optimal behavior on the platform, before defining the higher education market equilibrium.

Individuals' Location Decisions on the Labor Market

In the second stage of the model, taking place upon graduation, students choose in which local labor market to enter. In particular, they choose the location maximizing their second-stage payoff. Individual i will choose to enter location ℓ^* such that:

$$\ell^* = \arg\max_{\ell \in \mathcal{L}} \mathcal{W}_{i\ell}$$

When choosing which location to enter, individuals take into account the expected monthly wage they can receive in each local labor market ℓ, $w_{i\ell}$. In particular, I define $w_{i\ell} = E[w^*_{i\ell}|h_i, y_i]$ where $w^*_{i\ell}$ is a latent variable such that:

$$w^*_{i\ell} = w_\ell(h_i, y_i) + \tilde{\epsilon}_{i\ell}$$

where $\tilde{\epsilon}_{i\ell}$ is an idiosyncratic shock with zero-mean, independently distributed across i and ℓ, that the individual does not observe prior to choosing location ℓ, as in Kennan and Walker (2011). $w_\ell(h_i, y_i)$ is a function defined in Section 1.5.

This wage equation is very flexible, as it allows for differential returns to education across local labor markets through y_i, which includes the degree at graduation. This is thus consistent with a labor demand model allowing for degree-specific and location-specific exogenous productivity, as in Diamond (2016).[16] Note, however, that it does not allow for individuals' location choices to influence wages in equilibrium, which rules out the existence of within-location spillovers across workers and imperfect substitutability of workers with different

[16]The current wage equation is consistent with perfectly competitive local labor markets and a production function of the form $Y_\ell = \sum_s A_{\ell s} L_{\ell s}$, where Y_ℓ is the output in location ℓ, $A_{\ell s}$ is the location-specific productivity of workers of skill level s and $L_{\ell s}$ is the number of workers with skill level s in ℓ.

skills (Katz and Murphy, 1992). These effects are likely to be second-order in the short run, as I consider the location decisions of *entry-level workers* only, which represents a small fraction of the overall labor supply.[17]

Higher Education Market: Programs' Ranking of Applicants

On the platform, programs provide a ranking of applicants based on the student-program specific score defined in Equation (1.6), reflecting programs' preferences over applicants after having internalized the equity objectives of the government. This implies that the student ranked r^{th} by program j, denoted $i^*_{j,r}$, has the highest program j-specific score among applicants who have not been ranked higher, that is:

$$i^*_{j,r} = \underset{i \in \mathcal{A}_j \setminus \{i^*_{j,k}\}_{k=1}^{r-1}}{\arg\max} \mathcal{V}_{ji} \tag{1.7}$$

where $\mathcal{A}_j$ is the set of applicants to program j, including the outside option $0_{\mathcal{I}}$.

The above equation is implied by the assumption that rankings of applicants are a truthful reflection of programs' preferences, after taking into account policymakers' objectives. While programs do not have to disclose how their ranking decisions are made, this assumption is consistent with interviews of members of admission committees covering a broad range of programs, as discussed in Section 1.2. Admission officers state that they typically use an algorithm to rank students, where they first choose the weights they assign to each observed characteristic to reflect its importance in the admission decision. This behavior is made easier as the platform provides an online tool to generate a truthful ranking based on the program's chosen weights. This ranking is then modified to take into account the equity objectives of the government if not already met throughout the initial ranking provided.[18]

Higher Education Market: Equilibrium

Finally, I define the higher education market equilibrium. On this market, each student i is matched to a program $j \in \mathcal{J} \bigcup \{0_{\mathcal{J}}\}$. I assume that this final match is stable. A matching

[17] In future work, I plan to extend the model by including a location-specific labor demand allowing for spillovers across workers and imperfect substitutability across skill types, following Diamond (2016).

[18] For example, if most scholarship-eligible students have been ranked by the program at the bottom, the initial ranking would be modified to up-rank some of these students. These changes preserve 1) the initial ranking between scholarship-eligible students and 2) the initial ranking between non scholarship-eligible students.

between students and programs is stable if there exists no student-program pair that would like to break from their current match to re-match to each other (Gale and Shapley, 1962). Formally, this means that there is no (i, j)-pair, with $i \in \mathcal{I} \bigcup \{0_{\mathcal{I}}\}$ and $j \in \mathcal{J} \bigcup \{0_{\mathcal{J}}\}$, such that:

$$\mathcal{U}_{ij} > \mathcal{U}_{i\mu(i)} \text{ and } \mathcal{V}_{ji} > \mathcal{V}_{j\mu(j,\mathcal{R}_j)} \tag{1.8}$$

where $\mu(i)$ is the program to which student i is currently matched, $\mathcal{R}_j$ is the rank of the last-admitted student by program j and $\mu(j, \mathcal{R}_j)$ is the identity of this student. Such an equilibrium exists as long as there are no complementarities between admitted students. This assumption is met mechanically by the behavior of admission officers: they rank applicants individually, examining their characteristics separately, while admission offers are sent following this ranking. Regarding uniqueness, Azevedo and Leshno (2016) prove that a continuum model, where a finite number of programs are matched to a continuum of students, typically has a unique stable matching. They further prove that all of the stable matchings of a discrete economy converge to this unique stable matching of the continuum economy.[19] With its very large number of applicants, the French college admission system qualifies as a large discrete economy.

The definition of stability implies that, on the higher education market, each student is matched to the program they prefer among the set of programs willing to admit them. We can thus write this match as the solution of a discrete choice model with individual-specific choice sets (Fack et al., 2019):

$$\mu(i) = \underset{j \in \mathcal{F}_i \bigcup 0_{\mathcal{J}}}{\arg\max} \, \mathcal{U}_{ij}$$

where $\mathcal{F}_i$ is the set of programs which would admit student i, i.e. student i's *feasible set*. It is defined as the set of programs preferring to admit student i over their last-admitted applicant:

$$\mathcal{F}_i = \{j \in \mathcal{J} : \mathcal{V}_{ji} > \mathcal{V}_{j\mu(j,\mathcal{R}_j)}\} \tag{1.9}$$

Note that if a program leaves at least one seat empty, the identity of the last-admitted student will be the outside option, with associated utility $\mathcal{V}_{j0_{\mathcal{I}}}$. Such a situation can happen in two scenarios. First, a program may deem some applicants undesirable and prefer to keep some seats unfilled rather than admitting them. In other words, the program chooses the outside option over these applicants. Second, programs may receive fewer applications than seats.

[19] A discrete economy is an economy with a finite number of students.

In both cases, the program would then be feasible for a student i only if the latter is more desirable than the outside option.

Overall, assuming that the equilibrium is stable allows us to write the observed match as a function of the primitives of the model without having to take a stand on students' application behavior. Stability only requires that each student applies to their preferred feasible program. This is consistent with several plausible application strategies, making it a good approximation of what happens on the platform. For example, the equilibrium would be stable as long as students submit applications to the programs they prefer to the outside option and accept their favorite offer among the ones they receive.[20] While the number of applications is limited to ten in theory, which would make such behavior less plausible, the effective number of applications allowed is much larger in practice: applications to different institutions offering the same program type and major are typically processed as one. As a result, 25% of applicants submit more than 13 applications, while the maximum number of applications observed in the data is 162 (Table 1.A.3).

A second plausible application behavior is that students are strategic: they submit applications maximizing their expected utility. In this case, a stable equilibrium is likely to arise when students are able to anticipate which programs are feasible. This is made easier by the fact that we are in the presence of a very large market, with many students and many seats per program. In such a setting, admission probabilities become close to degenerate, such that it is unlikely that students with rational expectations omit their favorite feasible program from their application portfolio (Fack et al., 2019). In particular, I find that the correlation between the rank of a program's last admitted student in 2018 and 2019 is 0.96, supporting the idea that applicants can anticipate whether a program is feasible or not.

1.5 Identification and Estimation

In this section, I first show how the primitives of the model can be identified from data on individuals' location decisions at entry on the labor market, programs' rankings of applicants, and realized matches on the higher education market. Importantly, identification does not

[20] Note that many studies of school-choice mechanisms assume that students are truthful (Abdulkadiroğlu et al., 2017; Abdulkadiroğlu et al., 2020): they rank programs following their preferences. The behavior described above can thus be thought of as a less stringent version of this assumption, as students do not have to rank programs on the French college admission platform.

rely on any functional form assumptions on the deterministic part of utility functions. I then present the estimation strategy.

1.5.1 Identification

I describe the identification argument, which works by backward induction. I first show that preferences over locations can be identified from data on location choices at entry on the labor market. I then turn to the first stage of the model and describe how to identify programs' preferences over students from observed rankings of applicants by programs. Finally, I show how to identify students' preferences over programs from data on realized matches. An exclusion restriction allows us to identify the scale of students' idiosyncratic preference shocks, avoiding fixing the relative scale of payoffs across stages of the model.

Individuals' Preferences over Locations

After graduation, students choose which local labor market to enter. At this stage, the model becomes a static discrete choice model. Let us denote $\boldsymbol{m}$ the vector collecting locations' characteristics and $\boldsymbol{w}_i$ the vector collecting individual i expected wages across locations. Assumption 1 allows us to directly map the share of individuals with characteristics (h_i, y_i) working in location ℓ, $\forall \ell \in \mathcal{L}$, denoted $\Pr(\ell^* = \ell | h_i, y_i, \boldsymbol{m}, \boldsymbol{w}_i)$, to the deterministic part of individuals' preferences over local labor markets:

$$\Pr(\ell^* = \ell | h_i, y_i, \boldsymbol{m}, \boldsymbol{w}_i) = \frac{\exp(\theta(w_{i\ell} + \omega(h_i, y_i, m_\ell)))}{\sum_{\kappa \in \mathcal{L}} \exp(\theta(w_{i\kappa} + \omega(h_i, y_i, m_\kappa)))} \tag{1.10}$$

The left-hand side of the above equation can be directly identified from data recording the location choices of individuals upon graduation. As this mapping can be inverted, we can thus recover θ and $\omega(h_i, y_i, m_\ell)$. Details of the proof can be found in Appendix 1.C.3.

From this, the continuation value associated with a program with characteristics z_j chosen by a student with characteristics (x_i, h_i) can be identified. It writes:

$$\bar{U}(x_i, h_i, z_j) = \frac{1}{\theta} \int \beta^{t(y_i)} \left(\gamma + \log \left[\sum_{\ell \in \mathcal{L}} \exp(\theta(w_{i\ell} + \omega(h_i, y_i, m_\ell))) \right] \right) f\left(y_i | x_i, h_i, z_j\right) dy_i$$

where γ is the Euler constant and transition probabilities $f\left(y_i | x_i, h_i, z_j\right)$ can directly be identified from data recording each student's path within the higher education system. This captures, in particular, the probability that a student with characteristics (x_i, h_i) enrolling in a program with characteristics z_j graduates with a certain degree.

Program-Specific Student's Score

Programs observe applicants and rank them based on the student's program-specific score defined in Equation (1.6). Let us denote $\boldsymbol{x}$ and $\boldsymbol{h}$ the vectors collecting x_i and h_i, respectively, for all i. Under Assumption 1, the probability that student i is ranked r^{th} by program j, denoted $\Pr(i^*_{j,r} = i|\boldsymbol{x}, \boldsymbol{h}, z_j)$, has a closed-form expression:

$$\Pr(i^*_{j,r} = i|\boldsymbol{x}, \boldsymbol{h}, z_j) = \frac{\exp(v(x_i, h_i, z_j))}{\sum_{a \in \mathcal{A}_j \setminus \{i^*_{j,k}\}_{k=1}^{r-1}} \exp(v(x_a, h_a, z_j))} \tag{1.11}$$

The left-hand side of this equation is directly identified from data on the final rankings of applicants by each program on the platform. Again, this mapping can be inverted, such that we can now recover $v(x_i, h_i, z_j)$. Details are provided in Appendix 1.C.4.

Students' Preferences Over Programs

Finally, given the parameters identified above, I show that one can identify students' preferences from realized matches. Note that when programs' preferences over students, $\mathcal{V}_{ji}$, are perfectly observed, one can directly construct each student's feasible set. This is the case in many school choice systems, where the criteria and weights used to rank students are publicly available. In such settings, identification of students' preferences is direct as the match can be written as the solution of a discrete choice model with *observed* student-specific choice sets (Fack et al., 2019).

Yet, in the French higher education context, students' program-specific scores are not observed, which prevents constructing $\mathcal{F}_i$ directly. This is the case in most college admission processes, such as in the United States and the United Kingdom, where universities base their decisions on a discretionary combination of test scores, GPA, and extra-curricular information. Agarwal and Somaini (2022) and He et al. (2023) conclude that, in such settings, two different types of exclusion restrictions are necessary to identify students' preferences from realized matches: both a variable shifting only students' preferences and one shifting only feasible sets are needed, but may be difficult to find in practice.

I now show how additional data on programs' rankings of applicants allows us to circumvent the need for such exclusion restrictions. In particular, the share of students with observable characteristics (x_i, h_i) who are matched to program j writes:

$$\Pr(\mu(i) = j | x_i, h_i, \boldsymbol{z}) = \left[\sum_{\mathcal{F} \in \mathbf{F}} \Pr(\mathcal{F}_i = \mathcal{F} | x_i, h_i, \boldsymbol{z}) \Pr(\mathcal{U}_{ij} \geq \mathcal{U}_{ik}, \forall k \in \mathcal{F} | x_i, h_i, \boldsymbol{z}) \right] \times \Pr(j \in \mathcal{F}_i | x_i, h_i, \boldsymbol{z}) \tag{1.12}$$

where $\mathbf{F}$ is the set of all potential feasible sets and $\boldsymbol{z}$ is a vector collecting programs' characteristics for all j. The first term of the right-hand side expression is the probability for student i to face feasible set $\mathcal{F}$. Feasible sets depend on programs' preferences, as defined in Equation (1.9). The second term is the probability that the utility associated with program j is greater or equal to the one derived from any programs in $\mathcal{F}$. Under Assumption 1, we can derive a closed-form expression for this second term:

$$\Pr(\mathcal{U}_{ij} \geq \mathcal{U}_{ik}, \forall k \in \mathcal{F} | x_i, h_i, \boldsymbol{z}) = \frac{\exp(\frac{1}{\sigma}(u(x_i, z_j) - c(x_i, d_{ij}) + \bar{U}(x_i, h_i, z_j)))}{\sum_{k \in \mathcal{F} \bigcup 0_{\mathcal{J}}} \exp(\frac{1}{\sigma}(u(x_i, z_k) - c(x_i, d_{ik}) + \bar{U}(x_i, h_i, z_k)))} \tag{1.13}$$

where $\bar{U}(x_i, h_i, z_j)$ has been identified in a first step.

Students' preferences over programs are identified in two steps. First, as $v(x_i, h_i, z_j)$ is identified from data on programs' rankings of applicants, one can construct $\mathcal{F}$ following Equation (1.9), given a draw of programs' preference shocks:

$$\mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta}) = \{ j \in \mathcal{J} : v(x_i, h_i, z_j) + \eta_{ji} > v(x_{\mu(j,\mathcal{R}_j)}, h_{\mu(j,\mathcal{R}_j)}, z_j) + \eta_{j\mu(j,\mathcal{R}_j)} \} \tag{1.14}$$

where $\boldsymbol{\eta}$ is the vector collecting η_{ji} for all programs and students.[21] We can then rewrite Equation (1.12) as an integral over $\boldsymbol{\eta}$:

$$\frac{\Pr(\mu(i) = j | x_i, h_i, \boldsymbol{z})}{\Pr(j \in \mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta}) | x_i, h_i, \boldsymbol{z})} = \int \Pr(\mathcal{U}_{ij} \geq \mathcal{U}_{ik}, \forall k \in \mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta}) | x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta}) g(\boldsymbol{\eta}) d\boldsymbol{\eta}$$

where $g(.)$ is the probability distribution function of $\boldsymbol{\eta}$. The numerator of the left-hand side of this equation can be identified directly from data on realized matches. The denominator can be recovered using Equation (1.14) given that we identified programs' preferences in

[21] As in Agarwal and Somaini (2022), equilibrium cutoffs can be taken as fixed as I consider a large economy (Azevedo and Leshno, 2016). Hence, feasible sets only depend on observable characteristics and $\boldsymbol{\eta}$.

a previous step, while we observe the set of applicants to each program. The interior of the integral on the right-hand side has an analytical expression (Equation (1.13)) and the distribution of $\boldsymbol{\eta}$ is known.

Second, this mapping can be inverted to identify sequentially the scale of the shock, σ, $u(x_i, z_j)$, and $c(x_i, d_{ij})$.

Identification of σ: Note that we typically cannot identify the scale of the utility, such that we normalize the scale of the shock to one. In this model, it would however be a strong restriction because it would fix the relative scale of individuals' utility in stages one and two. I avoid this assumption, and now show how to identify σ.

The identification argument relies on an exclusion restriction: there is at least one variable belonging to the vector of student's characteristics entering in the second stage of the model, h_i, but excluded from the one entering in the flow utility associated with a higher education program, x_i. Under such an exclusion restriction, the data can be mapped into an expression where σ is the only unknown parameter.

Here, I assume that the excluded variable is the home province of the student. It can enter the flow utility associated with a given location, to allow for a potential home bias, but not the one from enrolling in program j. This exclusion restriction is fairly weak: it allows students' demand to depend on various forms of mobility frictions. In particular, the distance between student i and program j, d_{ij}, can be included in $c(x_i, d_{ij})$. The identification argument then relies on considering two students i and o with $x_i = x_o$ but different home locations, living at the same distance from a given program j to ensure that $u(x_i, z_j) - c(x_i, d_{ij}) = u(x_o, z_j) - c(x_o, d_{oj})$.[22] In contrast, the continuation value associated with a program depends on their home location and will thus differ across students i and o, allowing us to identify σ.

Proposition 1 *Suppose that $u(x_i, z_j) - c(x_i, d_{ij}) = u(x_o, z_j) - c(x_o, d_{oj})$ for $j \in \mathcal{J}$ with $x_i = x_o$ and $h_i \neq h_o$. If $\bar{U}(x_i, h_i, z_j) - \bar{U}(x_i, h_i, z_{0_\mathcal{J}}) \neq \bar{U}(x_o, h_o, z_j) - \bar{U}(x_o, h_o, z_{0_\mathcal{J}})$, then σ is identified.*

[22]Note that this would remain true even if we allow for the presence of a home bias in $c(.)$ as well. The identification argument would then rely on a program j located at the same distance from both students i and o, and outside of their respective home location. Similarly, the requirement of finding a program equidistant from both i and o can be relaxed: in fact, we can consider two programs, j' and j'', as long as $u(x_i, z_{j'}) - c(x_i, d_{ij'}) = u(x_o, z_{j''}) - c(x_o, d_{oj''})$, which requires the distance between i and j' and o and j'' to be equal.

The proof consists in taking a double difference in the logarithm of the probability to match to different programs across students i and o, in the spirit of a difference-in-differences. The within-student difference eliminates the denominator of the conditional choice probability, which is student-specific and depends on $\boldsymbol{\eta}$. The flow utility associated with program j, which is not identified at this stage, cancels out through the cross-student difference. Details can be found in Appendix 1.C.5.

Identification of u(.) and c(.): Given the parameters related to students' preferences over locations and the scale parameter previously identified, one can then non-parametrically identify the functions $u(x_i, z_j)$ and $c(x_i, d_{ij})$. This does not rely on exclusion restrictions. The proof is again constructive and can be found in Appendix 1.C.6.

1.5.2 Estimation

This section describes the estimation strategy. As the identification argument, the estimation proceeds backward, starting with the second stage of the model.

Estimation of Preferences over Local Labor Markets

Parametrization I parametrize the flow utility associated with entering the labor market in location ℓ for individual i as follows:

$$\begin{aligned}\omega(h_i, y_i, m_\ell) = &\ \mathbf{Degree}_i'\boldsymbol{\phi_1}\mathbb{1}\{\ell = \ell_i^G, \ell = \ell_i^H\} \\ &+ \mathbf{Degree}_i'\boldsymbol{\phi_2}\mathbb{1}\{\ell \neq \ell_i^G, \ell = \ell_i^H\} + \mathbf{Degree}_i'\boldsymbol{\phi_3}\mathbb{1}\{\ell = \ell_i^G, \ell \neq \ell_i^H\} + \delta_\ell\end{aligned} \tag{1.15}$$

where $\mathbf{Degree}_i$ is the set of dummies indicating the degree upon graduation of individual i. I consider three degree types: high school degree, two-year degree, and bachelor's degree and more. ℓ_i^G is the location of individual i at graduation, while ℓ_i^H is her home location. $\boldsymbol{\phi_1}$ thus captures individuals' taste for staying in their home location, conditional on having graduated there. $\boldsymbol{\phi_2}$ captures individuals' taste for returning to their home location when they did not graduate there. Both vectors reflect individuals' home bias, which may be stronger for those who also studied in their home location. These parameters are allowed to vary across degree types to take into account that the level of education may influence the strength of this home preference. $\boldsymbol{\phi_3}$ reflects individuals' taste for staying in their graduation location, when the latter does not correspond to their home location. It captures the presence of stickiness in location choice at entry on the labor market, or moving costs, and is also allowed to vary by degree type. δ_ℓ are location fixed effects, capturing differences in amenities.

I define $w_\ell(y_i, h_i)$ as follows:

$$w_\ell(y_i, h_i) = \mathbf{Degree}_i'\boldsymbol{\kappa}_\ell + \boldsymbol{\kappa} h_i$$

where $\boldsymbol{\kappa}_\ell$ captures that returns to different degrees may differ across locations. h_i is a vector of individual i's characteristics including ability, measured by the individual's end-of-high school exam grades, gender, and scholarship status in high school.

Estimation The wage equation is estimated via an auxiliary linear regression. The primitive parameters capturing preferences over locations are estimated by Maximum Likelihood, using data on location choices of entry-level workers. I consider that individuals choose which province to enter in metropolitan France, i.e. $\mathcal{L}$ is the set including the 96 French provinces. The estimated parameters are those maximizing the following log-likelihood:

$$L^L(\mathbf{h}, \mathbf{y}, \mathbf{m}, \mathbf{w}; \boldsymbol{\Phi}) = \sum_{i\in\mathcal{I}} \sum_{\ell\in\mathcal{L}} \mathbb{1}\{\ell_i^* = \ell\} \log\left(\frac{\exp(\theta(w_{i\ell} + \omega(h_i, y_i, m_\ell)))}{\sum_{\kappa\in\mathcal{L}} \exp(\theta(w_{i\kappa} + \omega(h_i, y_i, m_\kappa))}\right)$$

where $\mathbf{h}$, $\mathbf{y}$ and $\mathbf{m}$ gather the individuals' and locations' characteristics, respectively, $\mathbf{w}$ is the vector of individuals' expected wages across local labor markets and $\boldsymbol{\Phi} = \{\theta, \boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \boldsymbol{\phi}_3\}$.

Estimation of Programs' Preferences

Parametrization Program-j specific score of applicant i is a flexible function of the program's and student's characteristics. In particular, we have:

$$\mathcal{V}_{ji} = z_j'\boldsymbol{\Lambda} x_i^v + \eta_{ji}$$

where z_j is a vector of dummies corresponding to program j's type. I consider seven program types: three-year university programs, two-year elite programs, engineering and business schools, two-year technology-oriented programs, two-year vocational programs, nursing schools, and others. x_i^v is a vector of students' characteristics including the student's high school GPA, gender, high school track, socio-economic status, and scholarship status.

Estimation The vector of programs' preference parameters $\boldsymbol{\Lambda}$ can be estimated by Maximum Likelihood, as the ranking of applicants is observed. The log-likelihood writes:

$$L^P(\mathbf{x}, \mathbf{z}; \boldsymbol{\Lambda}) = \sum_{j \in \mathcal{J}} \sum_{i \in \mathcal{A}_j} \sum_{r=1}^{\bar{R}_j} \mathbb{1}\{i = i^*_{j,r}\} \log \left(\frac{\exp(z_j' \boldsymbol{\Lambda} x_i^v)}{\sum_{a \in \mathcal{A}_j \setminus \{i^*_{j,k}\}_{k=1}^{r-1}} \exp(z_j' \boldsymbol{\Lambda} x_a^v)} \right)$$

where $\bar{R}_j$ is the number of applicants ranked by j, and $\mathbf{x}$ and $\mathbf{z}$ gather the individuals' and programs' characteristics, respectively.

Estimation of Students' Preferences

Parametrization The value and mobility costs associated with enrolling in program j for student i are parametrized as follows:

$$u(x_i, z_j) = z_j' \boldsymbol{\delta} x_i^u \qquad\qquad c(x_i, d_{ij}) = d_{ij} \boldsymbol{\alpha}_1 x_i^d + d_{ij}^2 \boldsymbol{\alpha}_2 x_i^d$$

where $u(.)$ depends on the interaction of z_j with the vector of students' characteristics x_i^u including the student's gender, high school track, socio-economic status, honors at the end-of-high school exam, and scholarship status. $c(.)$ depends on the distance, in kilometers, between student i's home location — at the ZIP code level — and program j's GPS coordinates, denoted d_{ij}. The square of this distance is also included to allow for non-linearities in these costs. Through interactions with the vector of students' characteristics x_i^d, mobility frictions are allowed to vary across students with different socio-economic groups, scholarship eligibility, and gender.

Estimation As shown in the identification section, the conditional choice probability that a student with observable characteristics (x_i, h_i) is matched to j is an integral over $\boldsymbol{\eta}$, which does not have a closed-form expression. While preferences could be estimated by Simulated Maximum Likelihood (SMLE), the estimates would not be consistent when using a finite number of draws of $\boldsymbol{\eta}$. I avoid this issue by instead using a Simulated Method of Moments (SMM) procedure, which guarantees consistency with a finite number of draws (Gourieroux and Monfort, 1996), with optimal moments — the score of the likelihood. In particular, given the parameters estimated in a first step and for a given guess of students' parameters, I proceed as follows:

(i) Draw a vector of programs' shocks, $\boldsymbol{\eta}^b$.

(ii) Given programs' preference parameters and $\boldsymbol{\eta}^b$, compute $\mathcal{V}_{ji}^b, \ \forall i, j$ and construct $\mathcal{F}_i^b = \mathcal{F}^b(x_i, h_i, \boldsymbol{\eta}^b)$.

(iii) Construct the optimal moment associated with each variable entering in $\mathcal{U}_{ij}$. For example, the moment associated with d_{ij} writes:

$$\sum_i \sum_{j \in \mathcal{F}_i^b} \left[\mathbb{1}\{\mu(i) = j\} - \Pr(\mu(i) = j | \mathcal{F}_i^b, x_i, h_i, \boldsymbol{z}, j \in \mathcal{F}_i^b)\right] d_{ij}$$

where $\mathbb{1}\{\mu(i) = j\}$ is a dummy equal to one if student i matches to program j and $\Pr(\mu(i) = j | \mathcal{F}_i^b, x_i, h_i, \boldsymbol{z}, j \in \mathcal{F}_i^b)$ is the conditional choice probability that student faced with feasible set $\mathcal{F}_i^b$, given $(x_i, h_i, \boldsymbol{z})$, matches to j, defined in Equation (1.13).

For each vector of parameters guess, I repeat (i)-(iii) ten times and average moments across simulations. These averaged moments are then used to compute the Method of Moments objective function to be maximized.

Note that this method provides consistent estimates even for a fixed number of simulations (Gourieroux and Monfort, 1996; Train, 2009b). This is an advantage with respect to using SMLE, which is not consistent with a fixed number of draws. Method of Moments is usually less efficient than Maximum Likelihood estimation. However, it is efficient when the optimal instruments, the score, are used. In this framework, the score has to be simulated, which introduces some noise. Simulation noise decreases as the number of draws increases. The estimation method presented here is thus a computationally tractable way to obtain consistent and efficient estimators of two-sided matching models.

Estimation of Transition Probabilities

Parametrization As in Kennan (2015), I let students' graduation outcome be the result of a stochastic process, depending on their characteristics and on the chosen program's characteristics. Students have thus different probabilities of graduating with a high school degree, a two-year degree or a three-year degree or more depending on their program of enrollment. I also let these characteristics influence the student's graduation location. I thus do not impose that a student entering a program in location ℓ also graduates in location ℓ.

Transition probabilities between the initial program and the location-degree at exit of the higher education system can be recovered directly from the data. In particular, I use the administrative dataset tracking the universe of students enrolled in the French higher education system. From this, I compute the joint probability that a student with characteristics x_i and h_i entering a program with characteristics z_j graduates with a given type of degree and in a given location, i.e. $y_i = (\mathbf{Degree}_i, \ell_i^G)$.

Although transition probabilities are identified non-parametrically, I estimate them parametrically, for tractability. I use a conditional logit model to capture the probability for an individual to be in a given location-degree bundle, across $|\mathcal{D}| \times |\mathcal{L}|$ possibilities, where $\mathcal{L}$ is the set of graduation locations, which is the same as the one of local labor markets. $\mathcal{D}$ is the set of degree types: high school degree, two-year degree, and bachelor's degree or more. I parametrize this joint probability as follows:

$$f(y_i|x_i, h_i, z_j) \equiv \Pr(\mathbf{Degree}_i, \ell_i^G|x_i, h_i, z_j) = \frac{\exp(\rho_1 \mathbb{1}\{\ell_i^G = \ell_j\} + \mathbf{Degree}_i' \boldsymbol{\rho}_2 z_j + \mathbf{Degree}_i' \boldsymbol{\rho}_3 x_i^u)}{\sum_{\ell^G, \mathbf{Degree}} \exp(\rho_1 \mathbb{1}\{\ell^G = \ell_j\} + \mathbf{Degree}' \boldsymbol{\rho}_2 z_j + \mathbf{Degree}' \boldsymbol{\rho}_3 x_i^u)}$$

where ℓ_j is the province where program j is located.

Estimation The set of parameters $\boldsymbol{\rho} = \{\rho_1, \boldsymbol{\rho}_2, \boldsymbol{\rho}_3\}$ is estimated by Maximum Likelihood. Details can be found in Appendix 1.D.2.

Extension: Adding Unobserved Heterogeneity

The model can be extended to capture persistent unobserved preferences over locations. It may indeed be important to account for such unobserved preferences when modeling location choices, as they often display persistence. Allowing for unobserved heterogeneity would help disentangle the stickiness in location choices coming from unobserved preferences from state dependence (Heckman, 1981). We can rewrite the model in the second stage as follows:

$$\begin{aligned} \omega(h_i, \alpha_i, y_i, m_\ell) = \mathbf{Degree}_i' \boldsymbol{\phi_1} \mathbb{1}\{\ell = \ell_i^G, \ell = \ell_i^H\} + \mathbf{Degree}_i' \boldsymbol{\phi_2} \mathbb{1}\{\ell \neq \ell_i^G, \ell = \ell_i^H\} \\ + \mathbf{Degree}_i' \boldsymbol{\phi_3} \mathbb{1}\{\ell = \ell_i^G, \ell \neq \ell_i^H\} + \alpha_i a_\ell \end{aligned} \tag{1.16}$$

where a_ℓ is an amenity index and $\alpha_i \sim \mathcal{N}(0, \zeta)$. α_i is a persistent unobserved idiosyncratic taste for amenities.

Given the parameters of Equation (1.16) and a draw of α_i, $\forall i$, one can compute the expected discounted continuation value associated with choosing the optimal location in the second stage. Further assuming full persistence of the taste for the amenity index, the utility from enrolling in program j writes:

$$\mathcal{U}_{ij} = u(x_i, z_j) - c(x_i, d_{ij}) + \alpha_i a_j + \sigma\epsilon_{ij} + \bar{U}(x_i, h_i, \alpha_i, z_j)$$

where a_j is the amenity index associated with the location of program j. The probability that student i matches with program j simply rewrites:

$$\Pr(\mu(i) = j | x_i, h_i, \boldsymbol{z}) = \int \Pr(\mu(i) = j | x_i, h_i, \alpha_i, z) f(\alpha_i | \zeta) d\alpha_i$$

Details are provided in Appendix 1.D.3. Note that estimating preferences in this framework is more computationally costly as it requires to jointly estimate preferences.

1.6 Estimation Results and Model Fit

In this section, I first present parameter estimates, before discussing the fit of the model.

1.6.1 Estimation Results

Students' Preferences over Programs Estimates of students' preferences over programs highlight that students' demand is highly sensitive to distance (Tables 1.A.14-1.A.16). The average student would need to receive a monthly scholarship of 1,102 euros to accept to enroll 100 kilometers away from her home location. There is also significant heterogeneity across socio-economic groups, as this amount reaches 1,317 euros/month for low SES students, which is 384 euros larger than for high SES students. While they take into account both pecuniary and non-pecuniary costs associated with studying away from home, these figures are in the range of the average monthly budget for students, estimated to be 635 euros.[23]

Panel a) of Figure 1.6 further illustrates the importance of mobility frictions in students' higher education choices. In particular, for each applicant, I rank programs according to their flow utility. I then compute the correlation between this rank and the rank of the

[23]This figure is provided by the *Observatoire National de la Vie Etudiante*, a research institute created by the Ministry of Education. It is based on a survey of 245,000 students in 2020.

Figure 1.6: Determinants of Students' and Programs' Preferences

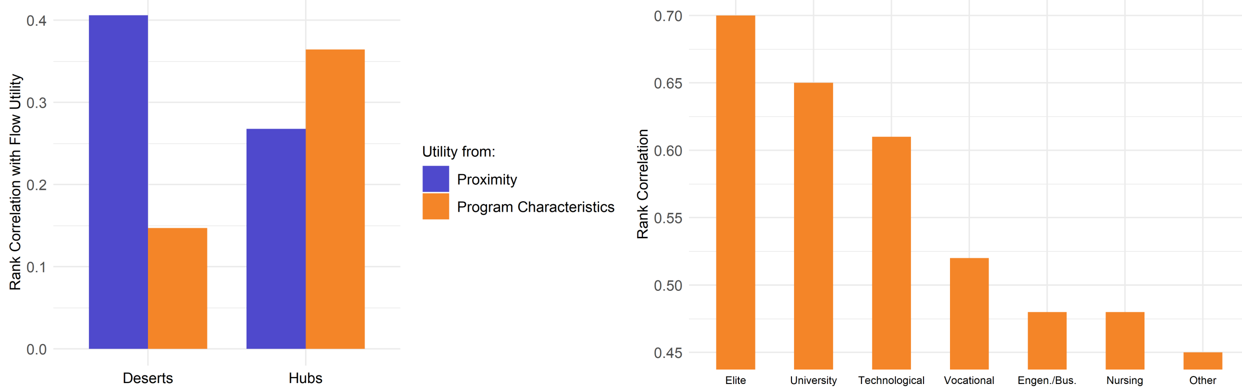

a) Students **b) Programs**

NOTES. Panel a) illustrates the determinants of the payoff students derive from enrolling in a given program. Using data at the student-program level, it plots, in purple, the rank correlation between the flow utility associated with a program and its proximity to the student's home location. It displays, in orange, the rank correlation between the flow utility associated with a program and the consumption value, $u(x_i, z_j)$, associated with the program's characteristics, z_j. It restricts attention to each student's top fifty programs. The figure shows these correlations separately for students from higher education deserts (the 50 provinces gathering less than 5% of the total supply of academic seats) and hubs (the ten provinces at the top of the seat supply distribution). Panel b) illustrates the importance of students' GPA in determining their ranking by different types of programs. Using data at the student-program level, it plots the rank correlation between the program-specific score associated with an applicant, $\mathcal{V}_{ji}$, and the GPA of the student. The figure shows these correlations separately by program type.

program according to its proximity to the student (in purple).[24] I do this separately for applicants from higher education deserts and hubs. I consider as educational deserts the fifty provinces gathering less than 5% of the supply of academic seats, while concentrating 25% of the students. Educational hubs are defined as the top ten provinces in the distribution of academic seats and similarly gather 25% of the student population. While this correlation is large for all students, it is about 60% higher for students from educational deserts. This is driven by the fact that these applicants have few higher education options around them, making distance a prevalent component of the payoff associated with a program. In contrast, with many more programs close to their home location, the ranking of students from higher education hubs is less sensitive to how distant programs are. As a result, their ranking is mostly determined by the utility they derive from the other characteristics of programs. The correlation between the utility rank associated with a program and its rank based only on the utility associated with its non-distance characteristics, $u(x_i, z_j)$, is twice as large for students from hubs (36%) as for students from low-opportunity areas (15%).

[24] A higher rank corresponds to a program closer to the student's home location. Given that distance is associated with a cost for students, ranking programs according to their proximity is the same as ranking them based on $-c(x_i, d_{ij})$.

Figure 1.7: Location Choices - Home Bias

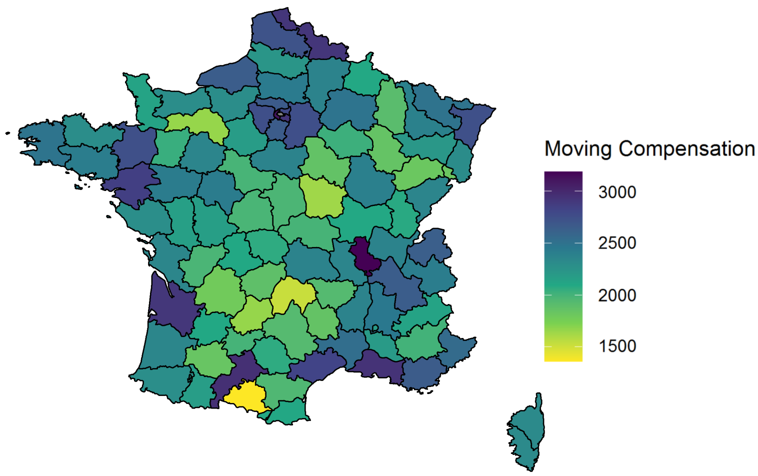

NOTES. This map uses estimated preferences over local labor markets (Table 1.A.18) to compute the compensation, in monthly euros, an individual from a given province would need to receive to be willing to move to Paris, conditional on having studied in their home location and having reached at least a bachelor's degree, abstracting from wage differences. Results can be found in Tables 1.A.20-1.A.21, Column (1).

Program-Specific Score Table 1.A.17 presents estimates of programs' preferences. Panel b) of Figure 1.6 shows how the ranking of students by programs correlates with students' GPA. This correlation is high for all programs, illustrating the importance of academic achievement in their admission decisions. This is particularly true for university and two-year elite programs, which are designed to lead students to graduate with at least a bachelor's degree. Final rankings are also determined by other applicants' characteristics. For example, the effect of being eligible for an income-based scholarship is equivalent to an increase in GPA of 0.27 standard deviations for admission to two-year elite programs, which may at least partly be driven by the equity objectives of the government.

Preferences over Local Labor Markets Table 1.A.18 presents estimates of the parameters associated with the utility of entering the labor market in a given location. Results show that workers display a very large preference for their home location. In Figure 1.7 and corresponding Tables 1.A.20-1.A.25, I use these estimates to compute the compensation an individual would demand to accept to enter the labor market in Paris rather than in her home location. I do this separately by level at graduation and keep wages fixed at their home location level. I first consider individuals graduating in their home location (Column (1)). For example, an individual from Creuse — one of the provinces at the bottom of the higher education supply distribution — who graduated with at least a bachelor's degree in Creuse would need to receive 2,060 euros per month to agree to enter the labor market in Paris rather than in her home location. This reflects a strong home bias, as this amount is more than twice as large as the average wage difference between Paris and Creuse. This home bias is stronger for individuals who graduated with a high school degree: they would need to receive 3,015 euros/month to make such a move. Cross-province variation highlighted in Figure 1.7 reflects differences in amenities: workers from provinces with high amenity levels would need larger compensation to be willing to move out of their home location.

While the preference for home is a key determinant of individuals' location choice, they also face very high moving costs. This can be seen in Column (2) of Table 1.A.20 and corresponding Figure 1.A.14, where I compute the same object as above, but conditional on graduating in Paris instead of home. A student from Creuse who moved to Paris to study and graduated there with at least a bachelor's degree is now better off when entering the labor market in Paris, holding wages across labor markets equal. This individual would actually need a monthly tax rebate of 542 euros to agree to return home upon graduation. As the home bias is stronger for individuals with lower levels of education, an individual who moved to Paris to study but dropped out from higher education would only need to receive an additional 156 euros per month to return home. Overall, the magnitude of these moving costs is likely to allow higher education hubs to retain students graduating there, despite the existence of a large home bias.

1.6.2 Model Fit

I simulate the higher education match and location decisions at entry on the labor market to evaluate the fit of the model.[25] Table 1.A.26 compares these simulations to what is observed in the data. Panel A presents characteristics of the observed higher education match in Column (1) and of the simulated one in Column (2). Panel B shows patterns related to location choices. Overall, the simulated shares of students enrolling in different program types, distance to the matched program, and share of individuals entering the labor market in their home or study location closely match what is observed in the data.

[25] I use both students' (Tables 1.A.14-1.A.15) and programs' (Table 1.A.17) preference parameters to compute their utility. I then feed them into a Deferred Acceptance algorithm to simulate the stable match. I use individuals' preferences (Table 1.A.18) over locations to simulate their choice of local labor market upon graduation.

Figure 1.8: Model Fit

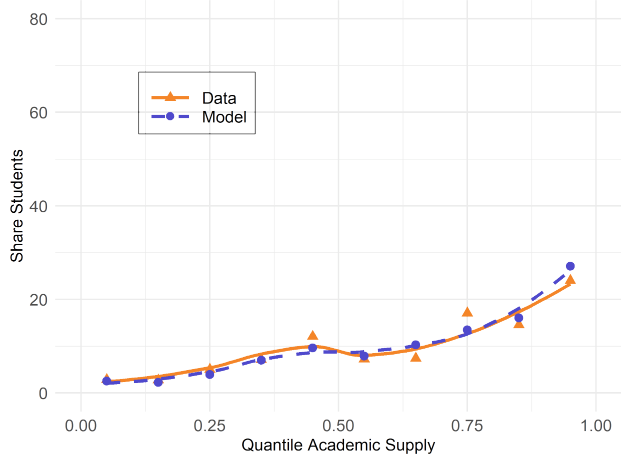

a) Enrollment Flows - Students from Deserts

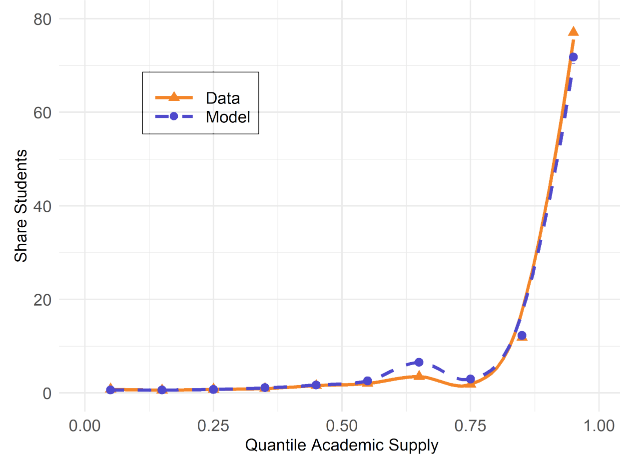

b) Enrollment Flows - Students from Hubs

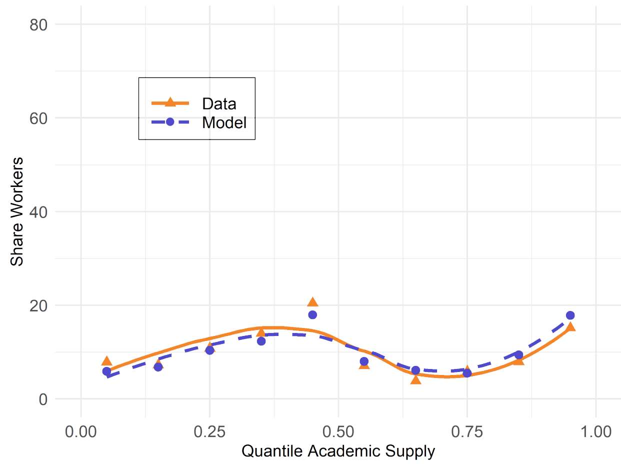

c) Work Flows - Workers from Deserts

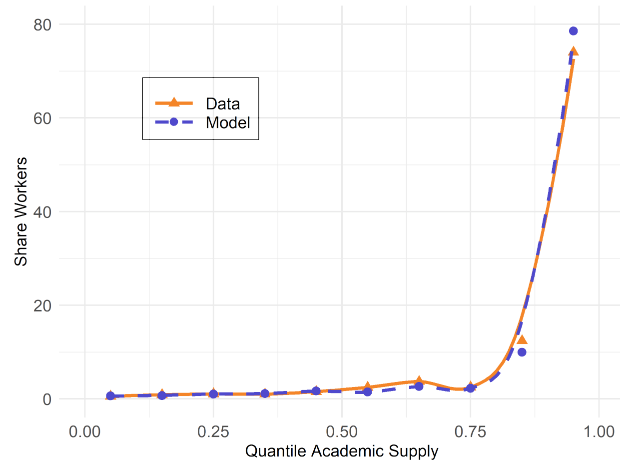

d) Work Flows - Workers from Hubs

NOTES. I rank provinces according to their supply of academic seats and group them into ten bins of equal size. Panel a) plots the share of students from higher education deserts enrolling in each of these bins. For example, about 25% of these students enroll in a program from a province in the top 10% of the academic seats supply distribution. Panel b) plots the same statistics, focusing on students from higher education hubs. Panel c) shows the share of workers from higher education deserts entering the labor market in each bin. For example, about 15% of them enter the labor market in a province belonging to the top 10% of the distribution of academic seats. Panel d) displays the same statistics for workers from higher education hubs. Each panel presents these flows in the data (orange triangles, plain line) and as predicted by the model (purple dots, dotted line).

I also evaluate the fit of the model using moments that are not directly targeted in the estimation. In particular, it is crucial for the model to be able to replicate migration flows both within the higher education market and across local labor markets. First, I simulate enrollment flows between provinces and contrast them with those observed in the data. I rank provinces according to their supply of academic seats and group them into ten bins. Panel a) of Figure 1.8 plots the share of students from higher education deserts enrolled in each of these bins. About 25% of students from higher education deserts enroll in an education hub, i.e. a province belonging to the top 10% of the academic seats supply distribution. This is in sharp contrast with enrollment flows of students from higher education hubs, as 75% stay in a hub to study (Panel b)). The model replicates these enrollment flows very closely, differently for students from hubs and deserts, while they were not targeted in the estimation.

I further use simulated enrollment flows at the province level to estimate the same gravity model as in Section 1.3, projecting them on origin and destination fixed effects, and on the distance between each location pair. Table 1.A.27 presents results from the model fitted on enrollment flows observed in the estimation sample in Column (1) and on simulated flows in Column (2). The latter yields an elasticity of enrollment flows to distance of -2.27, which is very close to that found in the data (-2.50).

I then turn to flows of workers across local labor markets, which will be key objects in counterfactual predictions. Panel c) of Figure 1.8 plots the share of individuals from higher education deserts entering the labor market in provinces according to their decile in the distribution of the supply of academic programs. Only about 15% of them enter the labor market in a higher education hub. Again, this contrasts sharply with the decision of individuals from higher education hubs, who are 60 percentage points more likely to enter the labor market in a hub. While these moments were not targeted in the estimation, the model is able to closely match them.

Finally, an alternative way to evaluate the fit of the model is to use preference estimates to predict the effect of a 500 euro moving subsidy, as the one evaluated in Section 1.3.2. This scholarship corresponds to a monthly payment of 55.5 euros.[26] Preference estimates allow us to express the utility cost of enrolling one kilometer farther away from home, which is about 11 euros/month. According to these estimates, receiving 55.5 euros/month would thus cover the cost of traveling an additional five kilometers, a figure belonging to the 99% confidence

[26]The academic year usually lasts about nine months.

interval of the effect of the subsidy found in the data (Table 1.A.8).

1.7 Counterfactuals

In this section, I first use the model to quantify how the uneven within-country distribution of higher education options interacts with mobility frictions to generate spatial inequalities, both between individuals and local labor markets. To do so, I eliminate mobility frictions within the higher education market, holding constant the distribution of programs. I then further remove moving costs at entry on the labor market as well. Second, I design and evaluate contracts allowing low-opportunity areas to outsource investments in their local human capital to higher education hubs, while providing students incentives to return home upon graduation.

1.7.1 The Geography of Higher Education and Spatial Inequalities

I first focus on spatial inequalities between students. Panel A of Table 1.3 presents several measures of these gaps at baseline (Column (1)). First, I compare individuals from higher education hubs and deserts. At baseline, we observe a 5 percentage point gap in the share of individuals graduating with at least a bachelor's degree in hubs (44.4%) and deserts (39.4%). The estimated model also allows us to quantify mean welfare differences, where welfare is computed based on Equation (1.4), for students across these regions, amounting to 991 euros/month. Second, I consider the variance of these different measures, across home locations, which captures the overall magnitude of regional inequalities.[27]

The model captures three key potential drivers of these gaps: differences in students' test scores entering programs' admission decisions, differences in students' preferences, and mobility frictions. It thus allows us to isolate how mobility frictions interact with the geography of the higher education market to generate the above spatial inequalities. To investigate this, I first break the link between the spatial distribution of programs and students' choices by eliminating mobility frictions. I simulate individuals' choices in this scenario and compare these counterfactual paths to the status quo. Details are provided in Appendix 1.E.1.

Column (2) of Table 1.3 presents results from these simulations. I find that 76% of the gap in the share of individuals graduating with a bachelor's degree between educational hubs and

[27] Details on the computation of all measures considered in this section are presented in Appendix 1.E.2.

Table 1.3: Mobility Frictions and Spatial Inequalities: Quantification

	Baseline (1)	Counterfactual (2)
Panel A: Spatial inequalities between individuals		
Share High-Skilled:		
Gap Hubs-Deserts	4.97	1.16 **(-76%)**
Cross-Province Variance	26.98	18.19 **(-33%)**
Mean Welfare (€/month):		
Gap Hubs-Deserts	991	614 **(-38%)**
Cross-Province Variance	230,050	93,490 **(-59%)**
Panel B: Inequalities across local labor markets		
Share High-Skilled:		
Gap Hubs-Deserts	7.67	3.72 **(-51%)**
Cross-Province Variance	26.94	20.88 **(-22%)**
Sum Wages (M€):		
Gap Hubs-Deserts	2.1	2.3 (+**10%**)
Cross-Province Variance	97.6	115.2 (+**18%**)

Notes. This table presents results from the main quantification exercise: removing mobility frictions within the higher education market. Panel A focuses on spatial inequalities between individuals, with respect to two variables: the share of individuals from a given province graduating with at least a bachelor's degree (share high-skilled, in percent) and welfare, based on Equation (1.4), expressed in monthly euros. Panel B focuses on spatial inequalities across local labor markets, measured by the share of their labor force with at least a bachelor's degree (in percent) and by the sum of entry-level workers' wages. Gaps in these variables are computed by comparing hubs and deserts — defined as the provinces respectively at the top and bottom of the academic seat distribution, both gathering 25% of the high school student population — and by measuring their variance across all provinces. Column (1) presents spatial inequalities in the status quo, while Column (2) shows their level and percentage change in the counterfactual. Appendix 1.E.1 details how counterfactual simulations are performed. The computation of each outcome is formalized in Appendix 1.E.2.

deserts can be attributed to mobility frictions within the higher education market. Overall, they account for one-third of the variance in educational attainment across home locations. Panel A of Table 1.3 also compares the average welfare of students from educational hubs and deserts in the status quo and after removing mobility frictions within the higher education market. While all students would benefit from removing such frictions, those from educational deserts would gain twice as much. As a result, the welfare gap between these two groups would decrease by 38%. The uneven distribution of higher education programs coupled with the existence of mobility frictions thus accounts for a substantial share of spatial inequalities in educational attainment and welfare across students.

Taking the point of view of local labor markets, however, leads to a very different conclusion (Panel B of Table 1.3). I find that inequalities across local labor markets, as measured by the variance of the sum of wages of entry-level workers in their labor force, increase by 18% in the counterfactual scenario (Column (2)). Such gaps between higher education hubs and deserts similarly rise by 10% with respect to the status quo. This is driven by the fact that hubs are able to retain in their labor force a significant portion of the new students going there to study, at the expense of deserts. We indeed observe a 14% increase in the number of individuals from deserts entering the labor market in hubs (Table 1.A.28). In particular, we observe an 11% increase in the share of entry-level workers in Paris (Column (2) of Table 1.A.29). This results from the importance of mobility costs, as well as from location-specific returns to education. As some individuals still return home, educational deserts benefit from the fact that they received a better education. In particular, the gap in the share of workers entering the labor force with at least a bachelor's degree decreases by 51%, from a baseline of 7.5 points. Yet, the former effect dominates. Overall, this quantification exercise allows us to highlight the key trade-off generated by alleviating mobility frictions within the higher education market. It would decrease the welfare gap between individuals from high- and low-opportunity areas, at the cost of accelerating regional divergence.

The importance of this trade-off would be magnified if we were to omit programs' admission decisions. Column (1) of Table 1.4 shows results from a counterfactual where we eliminate mobility frictions and let each student enroll in their favorite program, regardless of capacity constraints and programs' preferences. As students from hubs would no longer be crowded out from their favorite programs, the decrease in spatial gaps in educational attainment is slightly reduced (-58%). However, inequalities across local labor markets increase much more than when the higher education market is at equilibrium. As any student can now be admitted to a program located in a hub, the outflow of workers from deserts to hubs is more pronounced (Table 1.A.28). In particular, the share of workers choosing to enter the labor market in Paris would increase by 44% compared to the status quo (Column (3) of Table 1.A.29). It is thus important to take into account the distributional effects of programs' admission decisions and capacity constraints.

Capturing the dynamic link between the higher education and labor markets is also crucial in documenting the above trade-off. Column (2) of Table 1.4 illustrates this by presenting results from simulations where mobility frictions are eliminated, as in the main counterfactual, but individuals' labor market location choice is fixed to be the same as in the status quo. It thus ignores the fact that changes within the higher education market may alter individuals' labor market location decision. This omission would lead to conclude that eliminating mobility frictions within the higher education market would decrease not only

Table 1.4: Mobility Frictions and Spatial Inequalities: Quantification

	No Mobility Frictions within Higher Education Market			
	Only	& No Admission Restrictions	& Fixed Location Choice	& No Moving Costs
	(1)	(2)	(3)	(4)
Panel A: Spatial inequalities between individuals				
Share High-Skilled:				
Gap Hubs-Deserts	-76%	-58%	-76%	-33%
Mean Welfare (€/month):				
Gap Hubs-Deserts	-38%	-44%	-38%	-49%
Panel B: Inequalities across local labor markets				
Share High-Skilled:				
Gap Hubs-Deserts	-51%	-20%	-64%	-23%
Sum Wages (M€):				
Gap Hubs-Deserts	+10%	+32%	-1%	-11%

NOTES. This table presents results from a decomposition exercise. In all columns, mobility frictions within the higher education system are eliminated. Column (1) reports results from this change, corresponding to Column (2) of Table 1.3. Column (2) presents results from this counterfactual, when additionally ignoring programs' admission decisions: all students are admitted in their preferred program in stage 1. Column (3) presents results from a counterfactual where mobility frictions within the higher education market are removed, ignoring any effects from workforce re-allocation: location choices are fixed to be the same as under the status quo. Column (4) presents results from a counterfactual where both mobility frictions within the higher education as well as moving costs at the location choice stage are removed. Panel A focuses on spatial inequalities between individuals, with respect to two variables: the share of individuals from a given province graduating with at least a bachelor's degree (share high-skilled) and welfare, based on Equation (1.4), expressed in monthly euros. Panel B focuses on spatial inequalities across local labor markets, measured by the share of their labor force with at least a bachelor's degree and by the sum of entry-level workers' wages. Gaps in these variables are computed by comparing hubs and deserts — defined as the provinces respectively at the top and bottom of the academic seat distribution, both gathering 25% of the high school student population. Appendix 1.E.1 details how counterfactual simulations are performed. The computation of each outcome is formalized in Appendix 1.E.2.

spatial inequalities in students' welfare but also inequalities between local labor markets as it does not allow for workers' out-migration from deserts.

Finally, I explore the role of moving costs in generating the highlighted trade-off. The latter originates from the fact that alleviating mobility frictions within the higher education market leads to a brain drain of students toward higher education hubs. While human capital is portable, this initial effect harms low-opportunity areas, as a significant share of students who leave to study fail to come back. Such patterns could be generated, through the model, by location-specific returns to education and moving costs. I quantify the role of the latter by simulating individuals' study and location decisions when removing moving costs between their graduation and home location, in addition to eliminating mobility frictions within the higher education market. Column (3) of Table 1.4 shows that this would lead to further reducing spatial gaps in welfare, as students can now come back to their home location without incurring any moving costs. This would, however, reduce their incentives to obtain a bachelor's degree or more, given differences in returns to education across space. As a result, the decrease in educational attainment gaps across home locations would slow down compared to what is found in the main counterfactual. Yet, as more students would now come back, and with higher levels of education than in the status quo, inequalities across local labor markets would decrease. This hints at the fact that low-opportunity areas could benefit from outsourcing the education of their local labor force to higher education hubs, conditionally on incentivizing students to come back.

1.7.2 Outsourcing Education Through Contracts

Equalizing opportunities for students from different regions, and decreasing inequalities across local labor markets are two objectives often cited by policymakers. As one-third of the variation in educational attainment across provinces is explained by mobility frictions at the higher education stage, shutting down these barriers is key to reach the first goal. This could also allow low-opportunity areas to outsource the production of their local human capital to higher education hubs. The above simulations however reveal that many students would not come back, leading to a trade-off. Can policymakers circumvent this trade-off, and at what cost?

To go beyond this trade-off, higher education deserts need to decouple the educational investment of their local students from the local supply of programs, while incentivizing individuals to come back upon graduation. I explore the role of a policy instrument satisfying these two conditions. In particular, I design contracts offering a distance-based subsidy to students from higher education deserts. As a counterpart, they either have to enter the labor market in their home region or repay the scholarship. This is akin to scholarship programs already existing in Chile (Banco Central de Chile), Columbia (Banco de Republica), or Singapore, where students receive a scholarship to study abroad while agreeing to return to their home country upon graduation or repay.[28] The policy I evaluate thus corresponds to a within-country implementation of these programs.[29]

To evaluate such a policy, I simulate the counterfactual match by replacing the utility

[28]In Singapore, beneficiaries agree to pursue a career in public administration within the Singapore Public Service.

[29]Anecdotal evidence shows that some rural French communes implement such a scheme, for example to guarantee the presence of a general practitioner in their area.

Table 1.5: Policy - Change in Spatial Inequalities

	Change in Inequalities between:			
	Individuals (Mean Welfare)	Local Labor Markets (Sum Wages)	Share Repaying	Cost (% of Total Wages)
	(1)	(2)	(3)	(4)
$\tau = 0.25$	-10%	-17%	44%	6%
$\tau = 0.33$	-17%	-29%	39%	8%
$\tau = 0.50$	-33%	-52%	28%	15%

NOTES. This table evaluates a policy offering contracts to students from higher education deserts, conditional on entering their home labor market upon graduation, or having to repay. Such contracts offer them a distance-based subsidy covering τ% of the mobility frictions they face. Column (1) shows the effect of the policy on welfare inequalities across students from deserts and hubs. Column (2) presents its impact on gaps between hubs and deserts, as measured by the sum of the wages of entry-level workers in their labor force. Column (3) documents the share of individuals repaying the scholarship to avoid coming back home. This is also presented in Figure 1.A.15. Column (4) quantifies the cost of these subsidies, expressed as a share of the cohort of entry-level workers' total wages. Appendix 1.E.3 presents results from an additional policy counterfactual, where students commit to coming back if they participate in the program.

associated with a given program j, $\mathcal{U}_{ij}$, with the following object:

$$\begin{aligned} &u(x_i, z_j) - (1-\tau)c(x_i, d_{ij}) + \sigma\epsilon_{ij} \\ &+\frac{1}{\theta}\int \beta^{t(y_i)}(\gamma + \log(\exp(\sum_{\ell\in\mathcal{L}}\left[\theta(w_{i\ell} + \omega(x_i, y_i, m_\ell)) + \mathbb{1}\{\ell \neq \ell^H\}\tau c(x_i, d_{ij})))\right] f(y_i|x_i, h_i, z_j)dy_i \end{aligned} \tag{1.17}$$

In particular, the subsidy compensates for a fraction τ of mobility frictions associated with traveling distance d_{ij}, $c(x_i, d_{ij})$. The continuation value is also updated as students take into account that they will have to repay the subsidy if entering the labor market in a location that is not their home location, ℓ^H. I simulate individuals' paths when their utility is replaced by Equation (1.17).

Table 1.5 shows the effects of this policy on inequalities, both across individuals and local labor markets. I consider different degrees of generosity of the subsidy, with τ varying from 25% to 50%. When the subsidy covers 25% of the mobility frictions associated with distance, welfare inequalities across students from hubs and deserts decrease by 10%. As the contract incentivizes students to enter their home labor market upon graduation, by making them repay the policy otherwise, it leads to a 17% decrease in inequalities across places. As 44% of individuals prefer to repay the subsidy, its cost stays reasonable: it only represents 6% of the cohort's total wages. Overall, policies tying mobility scholarships to incentives to return would allow low-opportunity areas to leverage regional trade in education services to invest in their local human capital.

1.8 Conclusion

To what extent does the geography of the higher education market contribute to spatial inequalities, both between students and across local labor markets? This paper documents that, in France, higher education options are unevenly distributed across space, such that students with different home locations face very different college opportunities. This impacts their educational choices, as I show that their demand is sensitive to proximity. The local supply of colleges also correlates with the level of skills in each region's labor force. Breaking the dependence between educational choices and the local supply of higher education programs seems key to mitigating spatial inequalities. I, however, document that removing barriers to mobility within the higher education market has substantial distributional effects, and could lead to a long-lasting brain drain of students toward higher education hubs.

To quantify the role of the geography of the higher education market in generating spatial inequalities, I build and estimate a dynamic model, linking the higher education market at equilibrium to a location decision model at entry on the labor market. I find that students' demand for programs is highly sensitive to proximity and that there exist significant moving costs, generating stickiness in location choices. Quantification exercises highlight that the geography of the higher market, given the existence of mobility frictions, accounts for one-third of the gaps in educational attainment across space. However, eliminating such frictions would lead to an increase in the number of individuals from higher education deserts entering the labor market in hubs, increasing inequalities across regions. Policies tying mobility scholarships and incentives to return would allow low-opportunity areas to outsource the education of their local labor force, hence circumventing this trade-off.

The framework presented in this paper could be adapted to an international context to investigate the role of students' global migration in cross-country brain drain. Importantly, the administrative data used in this study record the college applications and educational paths of foreign students applying to French colleges. Future work should also investigate how local governments compete with each over through admission policies to attract and retain students in their labor force. Finally, endogenizing the production of human capital, to incorporate potential increasing returns and peer effects, would be key to studying the efficient distribution of higher education programs across space.

Appendices

1.A Additional Figures and Tables

Figure 1.A.1: Catchment Areas

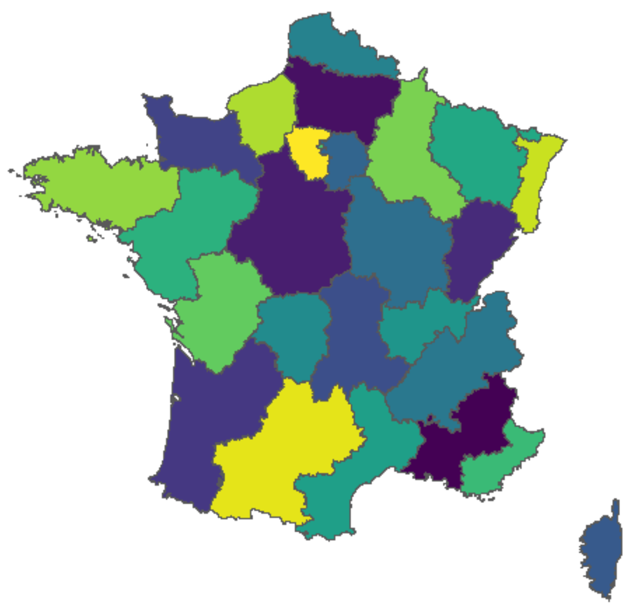

NOTES. This maps shows the geographic boundaries of the catchment areas in metropolitan France. Catchment areas outside metropolitan France are not shown.

Figure 1.A.2: Search Engine

NOTES. Screenshot from the Parcoursup website. It depicts the search engine allowing students to look for programs available.

Figure 1.A.3: Search Engine

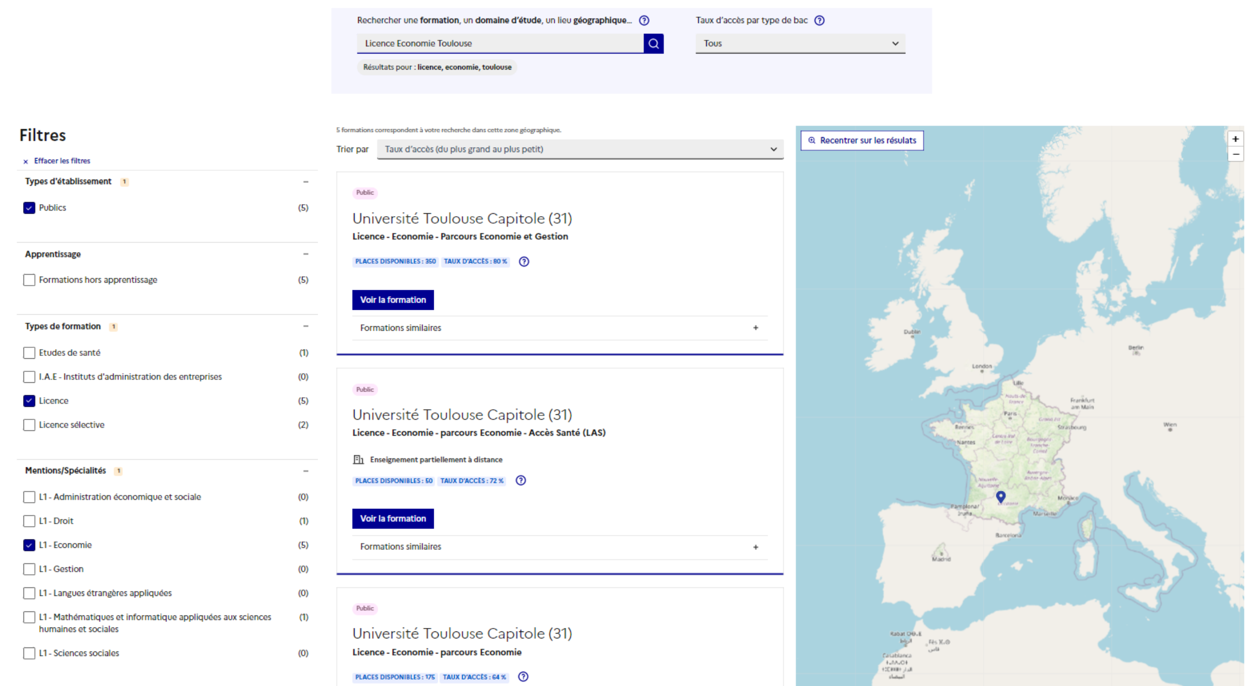

NOTES. Screenshot from the Parcoursup website. It depicts the result from a search for economics programs in universities located in Toulouse.

Figure 1.A.4: Search Engine

NOTES. Screenshot from the Parcoursup website. It depicts the type of information available to students when clicking on a specific program.

Figure 1.A.5: Search Engine

Vos résultats scolaires ou obtenus dans l'enseignement supérieur comptent pour 60%

Notes en mathématiques	Éléments évalués : Notes obtenues en première et terminale en mathématiques selon le programme et la nature de l'enseignement optionnel ou non, le cas échéant les notes du baccalauréat (selon la situation du candidat)	ESSENTIEL
Notes en anglais	Éléments évalués : Notes obtenues en anglais en première et terminale et le cas échéant les notes du baccalauréat (selon la situation du candidat)	TRES IMPORTANT
Notes en français	Éléments évalués : Notes des épreuves anticipées de français, ou en leur absence, les notes obtenues en français au cours de la scolarité de première	TRES IMPORTANT
Notes en philosophie	Éléments évalués : Notes obtenues en terminale en philosophie et le cas échéant les notes du baccalauréat (selon la situation du candidat)	IMPORTANT
Notes en histoire et géographie	Éléments évalués : Notes obtenues en histoire et géographie en première et terminale et le cas échéant les notes du baccalauréat (selon la situation du candidat)	IMPORTANT

Notes. Screenshot from the Parcoursup website. It depicts the type of information regarding admission criteria available to students when clicking on a specific program.

Figure 1.A.6: Search Engine

Vos compétences, méthodes de travail et savoir-faire comptent pour 35%

Compétence en analyse mathématique (continuité, dérivation, intégration, convergence), en algèbre (vecteurs, distance) et probabilité (indépendance, conditionnement, échantillons)	Éléments évalués : Notes obtenues en première et terminale en mathématiques et le cas échéant les notes du baccalauréat (selon la situation du candidat)	ESSENTIEL
Capacité à communiquer ses idées et de partager à l'écrit un raisonnement	Éléments évalués : Notes obtenues en philosophie en terminale et en français en première et le cas échéant les notes de baccalauréat (selon la situation du candidat)	TRES IMPORTANT
Culture économique	Éléments évalués : Notes obtenues en SES en première et terminale et le cas échéant les notes du baccalauréat (selon la situation du candidat)	IMPORTANT

Votre savoir-être compte pour 5%

Autonomie	Éléments évalués : Appréciations de bulletins de première et terminale. Fiche avenir	COMPLEMENTAIRE

Notes. Screenshot from the Parcoursup website. It depicts the type of information regarding admission criteria available to students when clicking on a specific program.

Figure 1.A.7: Search Engine

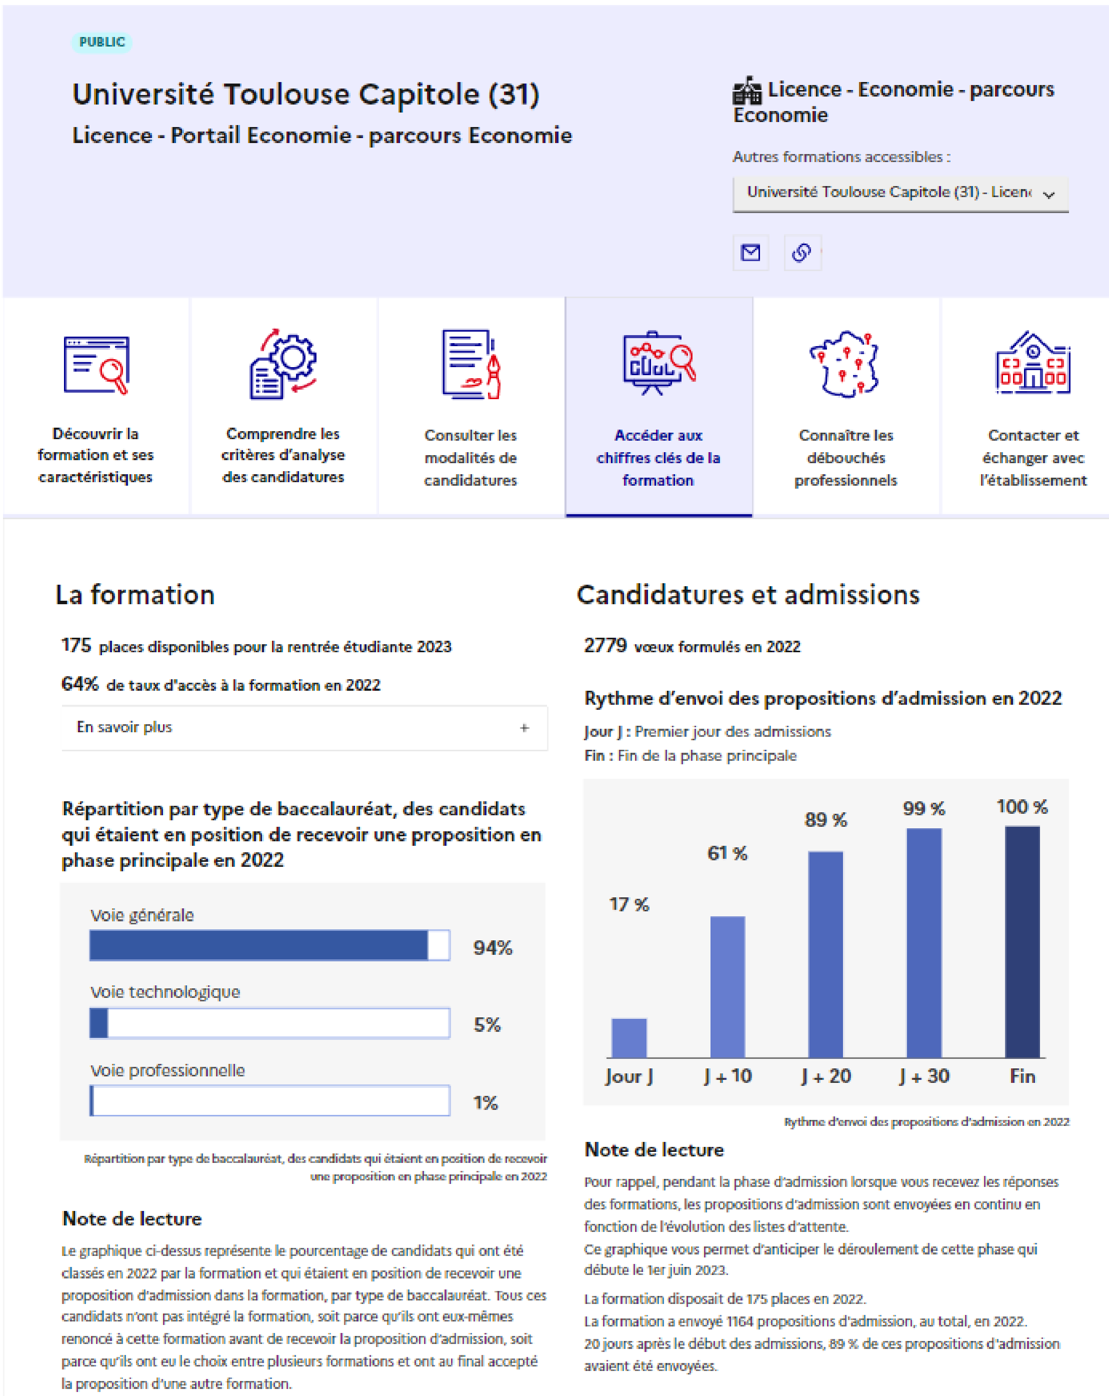

NOTES. Screenshot from the Parcoursup website. It depicts the type of information regarding admission criteria available to students when clicking on a specific program.

Table 1.A.1: Summary Statistics: Students

	Mean	St. Dev.
	(1)	(2)
Female	0.5396	0.4984
Socio-Economic Characteristics		
With Scholarship	0.1939	0.3954
High SES	0.3139	0.4641
Medium High SES	0.1566	0.3634
Medium Low SES	0.3119	0.4633
Low SES	0.2177	0.4127
End-of-High School Exam		
With Highest Honors	0.0946	0.2927
With High Honors	0.1744	0.3794
With Honors	0.2963	0.4566
Pass	0.4347	0.4957
High School Track		
General HS Track	0.6257	0.4840
Technological HS Track	0.2222	0.4157
Vocational HS Track	0.1521	0.3591
Observations	505,069	

Notes. This table presents summary statistics of college applicants in 2019. I restrict attention to high school students, living in metropolitan France.

Table 1.A.2: Summary Statistics: Programs

	Mean	St. Dev.
	(1)	(2)
Panel A: Program Type and Major		
Three-Year University	0.2235	0.4166
Humanities	0.4152	0.4929
Economics - Law	0.1639	0.3703
Social Sciences	0.1939	0.3954
STEM	0.2270	0.4190
Two-Year Elite	0.0743	0.2623
STEM	0.5253	0.4997
Economics - Law	0.3144	0.4646
Humanities	0.1603	0.3671
Engineering/Business	0.0716	0.2579
Engineering	0.5166	0.5000
Business	0.0601	0.2378
Arts & Architecture	0.4233	0.4944
Two-Year Technical	0.0712	0.2571
Production	0.5997	0.4903
Services	0.4003	0.4903
Two-Year Vocational	0.4405	0.4965
Production	0.3498	0.4769
Services	0.5492	0.4976
Agriculture	0.1011	0.3014
Nursing & Social Work	0.0520	0.2221
Nursing	0.6391	0.4807
Social Work	0.3609	0.4807
Other	0.0669	0.2498
Panel B: Other Characteristics		
Nber Seats	61	113
With Dorm	0.0572	0.2322
Observations	10,917	

NOTES. This table presents summary statistics of programs available on the platform in 2019, located in metropolitan France.

Table 1.A.3: Summary Statistics: Applications, Offers and Enrollment

	Mean	St. Dev.	25^{th} Pctile	50^{th} Pctile	75^{th} Pctile
	(1)	(2)	(3)	(4)	(5)
Nber Applications	10.587	9.291	5.000	9.000	13.000
Nber Offers	4.203	4.261	2.000	3.000	5.000
Distance of Enrollment	65.368	109.780	8.809	25.979	72.981

NOTES. This table presents summary statistics regarding applications, offers and enrollment of students in high school in 2019, living in metropolitan France. Characteristics of this sample are summarized in Table 1.A.1.

Table 1.A.4: Summary Statistics: Program Type Enrollment

	Mean (1)	St. Dev. (2)
Three-Year University	0.3920	0.4882
Two-Year Elite	0.0700	0.2552
Engineering/Business	0.0461	0.2098
Two-Year Technical	0.0896	0.2856
Two-Year Vocational	0.1685	0.3743
Nursing	0.0262	0.1597
Other	0.0178	0.1320
Unmatched	0.1898	0.3921
Observations	505,069	

NOTES. This table presents summary statistics regarding the program of enrollment of students in high school in 2019, living in metropolitan France. Characteristics of this sample are summarized in Table 1.A.1.

Table 1.A.5: Summary Statistics: Entry-Level Workers

	Mean	St. Dev.	Min	Max
	(1)	(2)	(3)	(4)
Panel A: Work and Study Locations				
Study Home	0.51			
Work Home	0.64			
Work and Study Home	0.40			
Work Home if Study Home	0.82			
If Study Out of Home Location:				
Work Home	0.46			
Work in Study Location	0.17			
Work Elsewhere	0.36			
Panel B: Earnings				
Monthly Earnings	1,477	804	0	14,000
Nber Individuals	26,347			

Notes. This table provides summary statistics of the study and work flows of individuals who exited the education system, as well as of their earnings. Locations are defined at the province level. Study (work, respectively) home refers to the share of individuals who studied (who are working, respectively) in their home province. 'Work and Study Home' refers to the share of individuals who both studied and work in their home province. The table then restricts the sample to students who studied out of their home province, and describes the share who moved back home to work, who stayed in their study location, or who entered the labor market outside their study or home province (work elsewhere).

Figure 1.A.8: Number of Academic Programs Within Commuting Distance

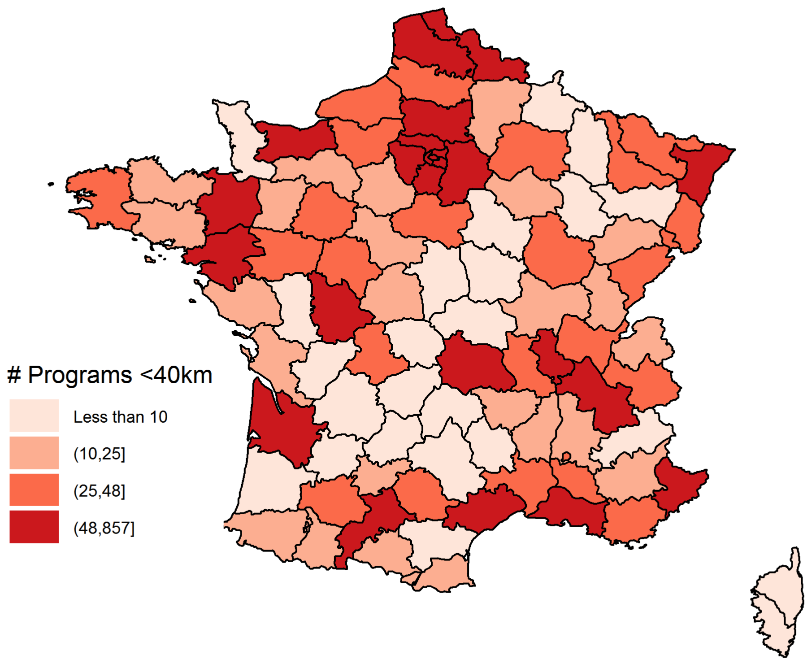

Notes. This map plots the number of academic programs the mean student of a given province has access to within a 40-kilometer radius from her home.

Figure 1.A.9: Spatial Distribution of Two-Year Technological and Vocational Programs' Seats

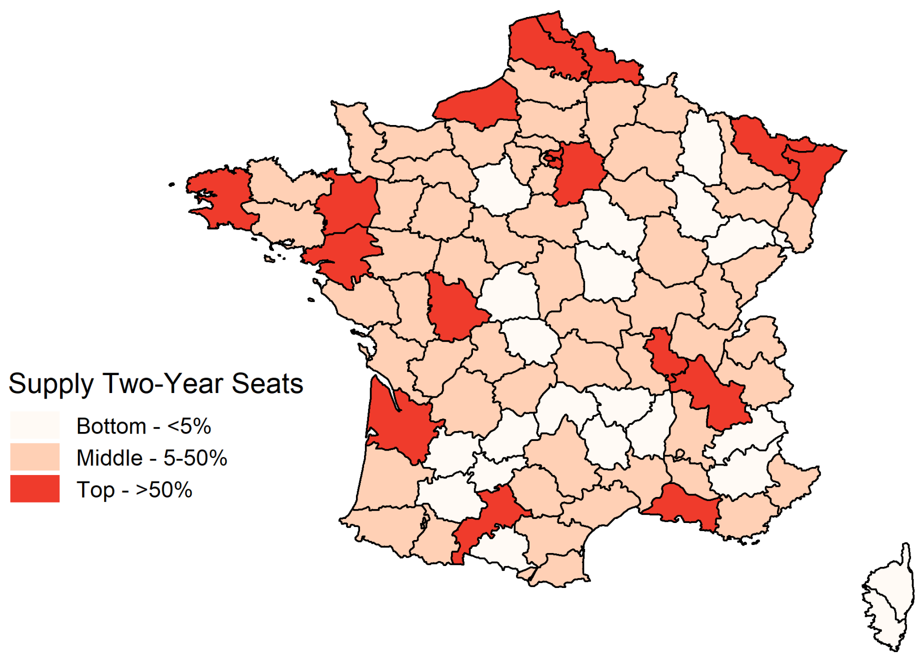

Notes. This figure maps the distribution of two-year technological and vocational programs' seats, at the province level. Provinces are ranked according to their position in the distribution of these programs' seats. Provinces in white are at the bottom of this distribution: they jointly account for less than 5% of the total supply. In contrast, provinces in red are at the top of this distribution: half of the total supply of two-year technological and vocational seats is concentrated there.

Figure 1.A.10: Cumulative Distribution of Two-Year Technological and Vocational Programs' Seats & Student Population

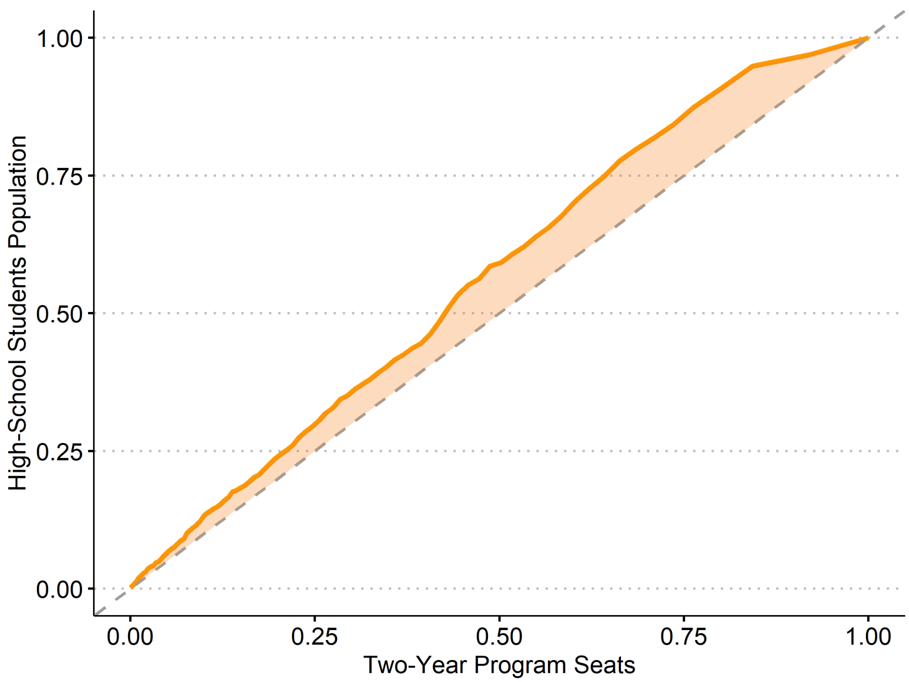

NOTES. This figure plots the relationship between the spatial distributions of two-year technological and vocational programs' seats and of high school students. It depicts the cumulative distribution of the supply of these seats across provinces on the x-axis, and the corresponding cumulative distribution of the high school student population on the y-axis. Provinces are ranked according to the share of vocational and technological seats they offer.

Table 1.A.6: Minimum Distance to Programs

	Quantile			
	5^{th}	10^{th}	25^{th}	50^{th}
	(1)	(2)	(3)	(4)
Three-Year University Programs				
Law	62	50	32	14
Economics	73	61	41	17
Medicine	108	92	62	30
Sports	77	63	43	18
Mathematics	90	73	46	18
Physics & Chemistry	97	78	48	20
Two-Year Elite Programs				
Sciences	57	46	26	11
Humanities	85	69	43	17
Economics & Business	67	56	37	15
Two-Year Technical Programs				
Legal Careers	256	210	150	73
Marketing	74	63	41	18
Chemistry	190	144	94	52
IT	98	81	56	27
Two-Year Vocational Programs				
Real Estate	110	93	60	28
Accounting	39	31	19	8
Construction	89	79	55	27
Engineering Schools	69	56	35	14
Management Schools	116	102	72	32
Nursing Schools	32	26	17	7

Notes. Minimum distance to travel, by type of programs. For example, 10% of students have to travel at least 92 kilometers to reach a medicine program.

Table 1.A.7: Gravity Model

	Enrollment Flows	
	OLS	PPML
	(1)	(2)
Distance$_{o,d}$	-1.912	-2.473
	(0.022)	(0.044)
Origin FE	Yes	Yes
Destination FE	Yes	Yes
Observations	5,878	8,930
Adjusted/Pseudo R-Square	0.7503	0.8734

NOTES. This table presents estimates of a gravity equation. It regresses students' enrollment flows between origin province o and destination province d on the distance (in kilometers) between the centroids of o and d, controlling for origin and destination fixed effects. Column (1) shows results from an estimation using OLS, while Column (2) presents results using PPML, as recommended by Silva and Tenreyro (2006). Robust standard errors are shown in parenthesis.

Figure 1.A.11: Enrollment in Medicine Program and their Local Availability

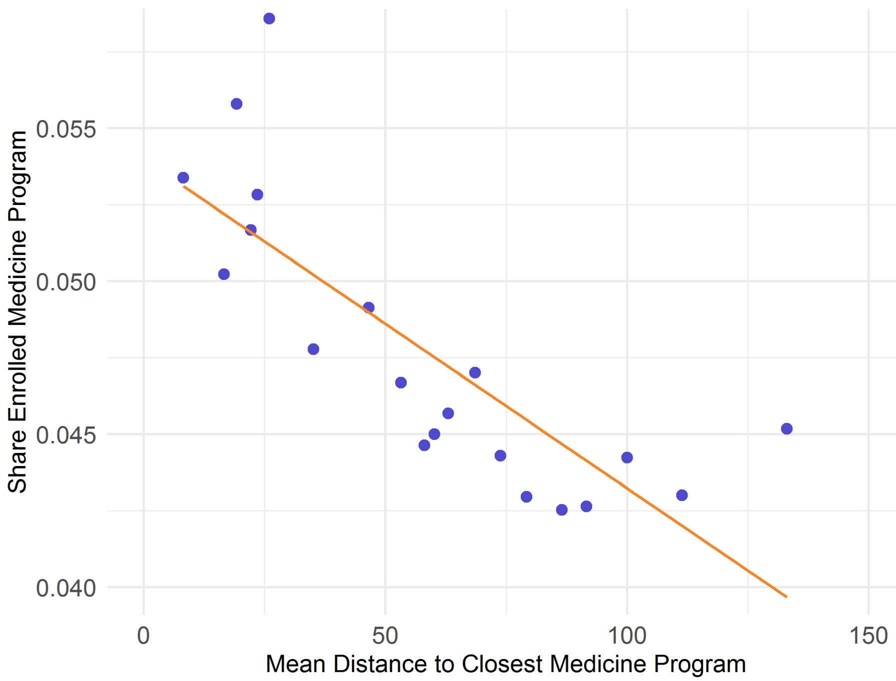

Notes. This figure is a binned scatter plot, at the province level, of the relationship between the share of students from a province enrolled in an medicine program and the mean distance to the closest medicine program for the average student of the province.

Figure 1.A.12: Enrollment in Two-Year Vocational Program and Local Availability of Academic Programs

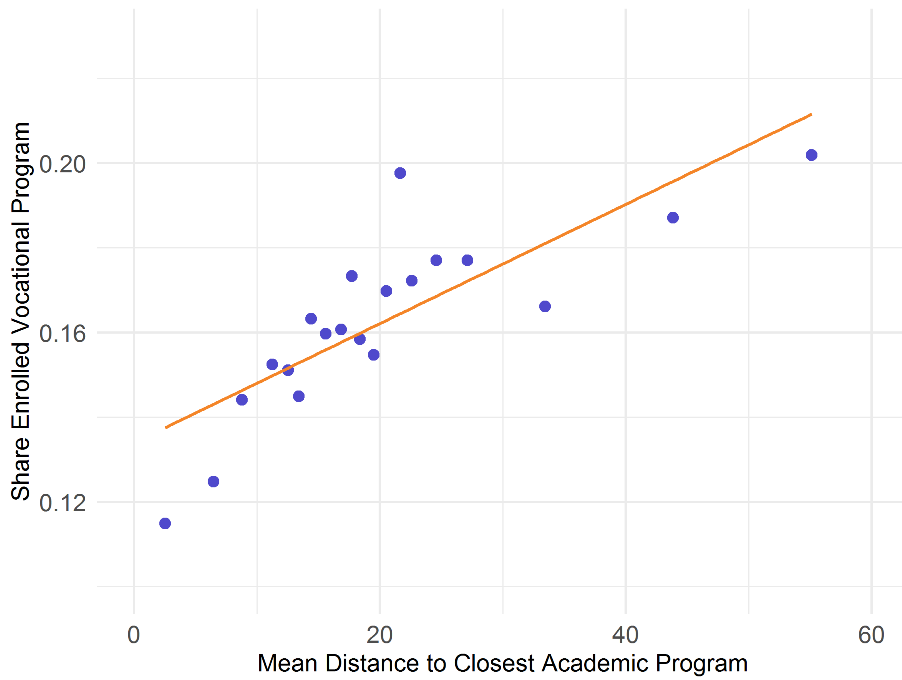

NOTES. This figure is a binned scatter plot, at the province level, of the relationship between the share of students from a province enrolled in a two-year vocational program and the mean distance to the closest academic program for the average student of the province.

Table 1.A.8: Effect of the Mobility Scholarship

		Out-of-Catchment Area		
	Distance	Applications	Offers	Accepted Offers
	(1)	(2)	(3)	(4)
β_{2019}	11.3616	0.0306	0.0073	-0.0248
	(2.8061)	(0.0089)	(0.0128)	(0.0246)
β_{2017}	0.8575	-0.0084	-0.0138	0.0039
	(2.8914)	(0.0096)	(0.0135)	(0.0333)
β_{2016}	-3.8352	-0.0250	-0.0443	-0.0882
	(3.0376)	(0.0100)	(0.0197)	(0.0312)
β_{2015}	-2.9885	-0.0045	0.0348	-0.0043
	(3.2062)	(0.0104)	(0.0205)	(0.0332)
Observations	3,675,961	3,675,961	907,240	226,926

NOTES. This table presents parameter estimates of β_τ from Equation (1.1). It shows the effect of the mobility subsidy, implemented in 2019, on the mean application distance (Column (1), see also Figure 1.4), on the probability to submit an of out-of-catchment area application (Column (2)), to receive an out-of-catchment area offer (Column (3)) and to accept an out-of-catchment area offer (Column (4)). Results shown correspond to defining low-supply areas as the five locations at the bottom of the supply distribution. Effects are robust to extending this definition to the ten locations with the least supply, while the magnitude of the estimates is slightly smaller. High-supply areas are defined as the ten locations at the top of this distribution. Results are robust to extending this set. Robust standard errors are shown in parenthesis.

Figure 1.A.13: Map of the Catchment Areas - Greater Paris Area

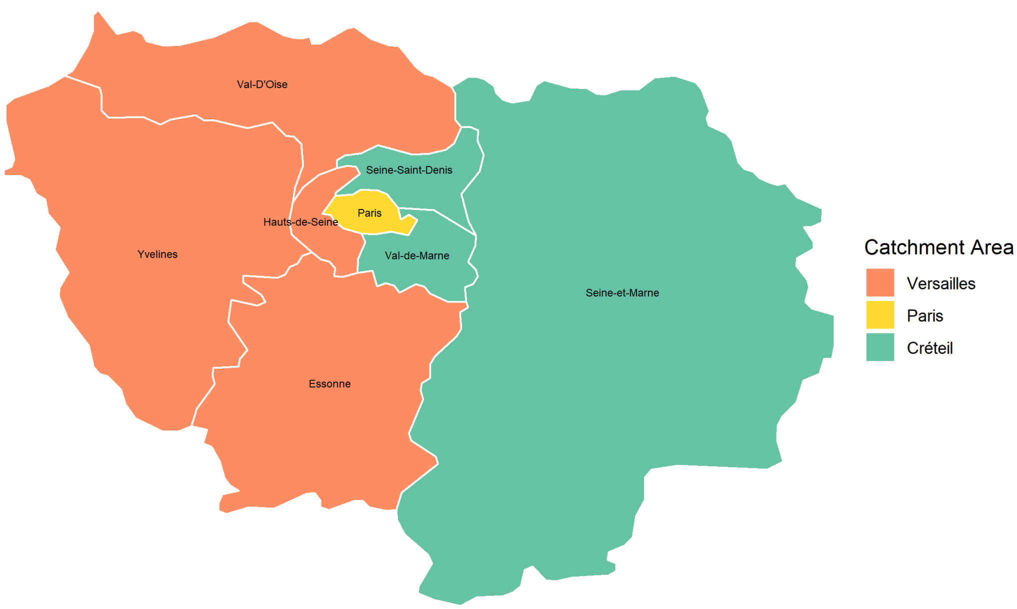

NOTES. The Greater Paris Area is composed of three catchment areas: Versailles, Paris and Créteil.

Table 1.A.9: Descriptive Statistics: Greater Paris Area Programs

	Catchment Area			
	Paris	Versailles	Créteil	Greater Paris Area
University				
Number University Programs	167	92	81	340
Number Universities/Branches	10	8	7	25
Number Ranked Applicants	3,073	2,884	2,635	2,923
Rank Last Admitted	1,000	1,224	1,033	1,064
% High-Ability Overall	27.17	14.38	8.32	19.22
% High-Ability from Paris	5.81	0.92	0.48	3.22
% High-Ability from Versailles	4.92	11.40	0.94	5.72
% High-Ability from Créteil	4.82	1.31	5.35	3.99
Non-University				
Number Programs	437	499	446	1,382
Number Institutions	113	186	158	457
Number Ranked Applicants	1,121	784	748	879
Rank Last Admitted	156	94	90	111
% High-Ability Overall	48.32	29.24	21.41	32.75
% High-Ability from Paris	13.62	2.21	1.20	5.49
% High-Ability from Versailles	10.94	19.85	3.38	11.72
% High-Ability from Créteil	6.85	2.28	14.01	7.53

Notes. Descriptive statistics of the programs offered in the Greater Paris area in 2018. The shares correspond to the mean percentage of high-ability students overall and from Paris/Versailles/Créteil admitted by each type of programs.

Table 1.A.10: Institutional Barriers to Mobility: Greater Paris Area

	Program's Catchment Area		
	Paris (1)	Versailles (2)	Créteil (3)
All Students			
From Versailles	-483	197	-180
From Créteil	-157	-373	171
From Paris	943	-400	-111
High-Ability Students			
From Versailles	-807	73	-200
From Créteil	-470	-472	91
From Paris	407	-522	-93

NOTES. Average difference between the rank of an applicant before and after the implementation of geographic priorities in 2018. On average, the application of a student from Versailles to a university program located in Paris looses 483 ranks after the implementation of geographic priorities.

Table 1.A.11: Removing Barriers to Mobility: Application Effects

	Share High-Ability Applicants		Share Low-Ability
	from Versailles	from Créteil	from Paris
	(1)	(2)	(3)
β_{Paris}	0.007***	0.002	-0.007**
	(0.002)	(0.002)	(0.004)
$\beta_{\text{Versailles}}$	-0.016***	0.004***	0.008***
	(0.003)	(0.001)	(0.001)
$\beta_{\text{Créteil}}$	0.003***	-0.009***	0.007***
	(0.001)	(0.002)	(0.002)
Baseline Paris	0.064	0.043	0.168
Baseline Versailles	0.143	0.019	0.048
Baseline Créteil	0.024	0.073	0.064
Program FE	Yes	Yes	Yes
Observations	680	680	680

NOTES. This table presents estimates from Equation (1.2). I study the effect of removing barriers to mobility on the share of high-ability students from Versailles (Column (1)), the share of high-ability students from Créteil (Column (2)) and the share of low-ability students from Paris (Column (3)) applying to university programs located in Paris, Versailles and Créteil. Standard errors clustered at the program level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table 1.A.12: Removing Barriers to Mobility: Sorting Effects

	Share High-Ability		Share Low-Ability
	from Versailles	from Créteil	from Paris
	(1)	(2)	(3)
$\beta^{DD}_{\text{Paris}}$	0.0360***	0.0061	-0.0787***
	(0.0082)	(0.0085)	(0.0173)
$\beta^{DD}_{\text{Versailles}}$	-0.0136***	0.0044	0.0195***
	(0.0128)	(0.0043)	(0.0063)
$\beta^{DD}_{\text{Créteil}}$	0.0099	0.0020	0.0232***
	(0.0042)	(0.0066)	(0.0069)
Baseline Paris	0.0492	0.0482	0.2205
Baseline Versailles	0.1140	0.0131	0.0303
Baseline Créteil	0.0094	0.0535	0.0472
Program FE	Yes	Yes	Yes
Observations	2,764	2,764	2,764

NOTES. This table presents estimates from Equation (1.19). I study the effect of removing barriers to mobility on the share of high-ability students from Versailles (Column (1)), the share of high-ability students from Créteil (Column (2)) and the share of low-ability students from Paris (Column (3)) enrolling in programs located in Paris, Versailles and Créteil. Standard errors clustered at the program level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table 1.A.13: Removing Barriers to Mobility: Application Effects

	Share High-Ability Applicants		Share Low-Ability
	from Versailles	from Créteil	from Paris
	(1)	(2)	(3)
$\beta^{DD}_{\text{Paris}}$	0.0132***	0.0005	-0.0068
	(0.0026)	(0.0020)	(0.0043)
$\beta^{DD}_{\text{Versailles}}$	-0.0091***	0.0032	0.0043**
	(0.0128)	(0.0011)	(0.0017)
$\beta^{DD}_{\text{Créteil}}$	0.0039	-0.0076	0.0056***
	(0.0013)	(0.0021)	(0.0021)
Baseline Paris	0.0636	0.0432	0.1685
Baseline Versailles	0.1433	0.0186	0.0484
Baseline Créteil	0.0245	0.0735	0.0643
Program FE	Yes	Yes	Yes
Observations	2,764	2,764	2,764

Notes. This table presents estimates from Equation (1.19). I study the effect of removing barriers to mobility on the share of high-ability students from Versailles (Column (1)), the share of high-ability students from Créteil (Column (2)) and the share of low-ability students from Paris (Column (3)) applying to programs located in Paris, Versailles and Créteil. Standard errors clustered at the program level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table 1.A.14: Students' Preferences - Value

		Heterogeneity								
		Female	Medium High SES	Medium Low SES	Low SES	Schol-arship	High Honors	Honors	Pass	Voc/Tech
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
University	-1.681	0.212	0.745	0.744	0.635	0.550	0.609	0.829	0.756	-2.354
	(0.167)	(0.083)	(0.126)	(0.110)	(0.122)	(0.119)	(0.175)	(0.162)	(0.153)	(0.096)
Two-Year Elite	0.868	-0.975	0.193	0.313	-0.225	0.191	0.079	-0.569	-1.536	-2.164
	(0.183)	(0.141)	(0.211)	(0.183)	(0.256)	(0.250)	(0.194)	(0.208)	(0.274)	(0.249)
Two-Year Technical	-2.196	-0.720	0.642	0.513	0.315	0.304	1.536	1.807	1.567	-0.573
	(0.284)	(0.120)	(0.172)	(0.154)	(0.176)	(0.167)	(0.291)	(0.279)	(0.278)	(0.128)
Nursing Schools	-6.085	2.108	0.586	0.861	0.342	-0.128	1.817	1.911	0.941	1.429
	(0.562)	(0.318)	(0.313)	(0.253)	(0.302)	(0.273)	(0.497)	(0.488)	(0.517)	(0.217)
Two-Year Vocational	-4.403	-0.416	0.736	0.870	0.854	0.220	0.996	1.402	1.314	1.412
	(0.275)	(0.098)	(0.164)	(0.138)	(0.146)	(0.127)	(0.285)	(0.265)	(0.260)	(0.117)
Engineering/Business	0.114	-0.881	0.312	0.196	-0.029	-0.089	0.825	0.432	-0.661	-2.001
	(0.233)	(0.156)	(0.224)	(0.202)	(0.250)	(0.280)	(0.239)	(0.239)	(0.272)	(0.215)
Other	-3.161	-0.645	0.078	1.257	1.076	0.197	1.276	0.787	0.866	-0.416
	(0.593)	(0.245)	(0.491)	(0.320)	(0.338)	(0.302)	(0.562)	(0.564)	(0.539)	(0.245)
Number Students	10,234									

NOTES. Standard errors in parentheses. Parameters correspond to $\frac{1}{\sigma}u(x_i, z_j)$. The estimated parameter for σ can be found in Table 1.A.15.

Table 1.A.15: Students' Preferences - Mobility Frictions

		Heterogeneity				
		Female	Medium High SES	Medium Low SES	Low SES	Schol-arship
	(1)	(2)	(3)			
Distance (100km)	2.154	0.075	0.341	0.441	0.896	0.007
	(0.061)	(0.059)	(0.000)	(0.077)	(0.091)	(0.000)
Distance^2	-0.166	-0.004	-0.029	-0.025	-0.070	-0.003
	(0.008)	(0.008)	(0.013)	(0.010)	(0.012)	(0.010)
σ^{-1}	2.170					
	(0.135)					
Number Students	10,234					

NOTES. Standard errors in parentheses. Note that the function capturing mobility frictions enters $\mathcal{U}_{ij}$ with a negative sign. Positive parameters associated with distance thus indicate a distaste for this variable. The negative parameters associated with the square of distance imply that these costs are concave. Parameters associated with Distance and Distance^2 correspond to $\frac{1}{\sigma}c(x_i, d_{ij})$.

Table 1.A.16: Students' Preferences - Compensation to Enroll 100km Away from Home

	Average	Female	High SES	Low SES	Scholar-ship
Monthly Subsidy (€)	1,102	1,125	933	1,317	1,192

NOTES. This table presents the monthly compensation, in euros, required by students to be willing to enroll in a program located 100 kilometers away from their home location. It is computed using parameters estimates in Table 1.A.15. Note that the continuation value is expressed in monthly euros, while its associated parameter is normalized to one in $\mathcal{U}_{ij}$. Parameters of $u(.)$ and $c(.)$ can thus be interpreted in monthly euros.

Table 1.A.17: Programs' Preferences

	University	Two-Year Elite	Two-Year Technical	Nursing	Two-Year Vocational	Engineer-ing/Business	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant	1.733	-1.543	0.846	1.743	2.018	1.461	2.486
	(0.062)	(0.080)	(0.063)	(0.081)	(0.015)	(0.053)	(0.059)
Grade	0.506	0.628	0.472	0.286	0.331	0.296	0.257
	(0.001)	(0.002)	(0.001)	(0.003)	(0.001)	(0.002)	(0.003)
Female	0.218	0.110	0.258	0.059	0.238	0.183	0.193
	(0.004)	(0.006)	(0.005)	(0.018)	(0.003)	(0.008)	(0.010)
Technological Track	-1.051	-1.093	-1.311	-0.424	-0.481	0.037	-0.501
	(0.008)	(0.051)	(0.074)	(0.013)	(0.007)	(0.010)	(0.012)
Vocational Track	-1.164	-1.990	-1.717	-0.856	-1.028	-0.506	-0.501
	(0.0146)	(0.187)	(0.018)	(0.018)	(0.011)	(0.015)	(0.012)
Medium-High SES	0.032	-0.000	0.054	-0.070	0.033	0.061	0.019
	(0.005)	(0.008)	(0.006)	(0.018)	(0.004)	(0.010)	(0.015)
Medium Low SES	0.009	0.031	0.028	-0.098	0.035	0.031	-0.004
	(0.005)	(0.007)	(0.005)	(0.016)	(0.003)	(0.009)	(0.012)
Low SES	0.034	0.080	0.037	-0.128	0.035	-0.007	-0.016
	(0.006)	(0.010)	(0.006)	(0.017)	(0.003)	(0.012)	(0.014)
With Scholarship	0.417	0.400	0.247	0.024	0.068	0.054	0.081
	(0.005)	(0.009)	(0.006)	(0.015)	(0.003)	(0.011)	(0.012)
Same Catchment Area	0.293						
	(0.004)						
High School with HE Program		0.133			-0.003		
		(0.006)			(0.003)		
Same High School		0.448			0.361		
		(0.012)			(0.004)		
Number Programs	10,917						

Notes. Standard errors in parentheses.

Table 1.A.18: Preferences over Local Labor Markets

Home Bias		
High School Diploma × Home Location	5.689	(0.046)
Two-Year Degree × Home Location	5.312	(0.051)
Bachelor's Degree and More × Home Location	4.214	(0.027)
Moving Costs		
High School Diploma × Graduation Location	3.877	(0.060)
Two-Year Degree × Graduation Location	3.463	(0.070)
Bachelor's Degree and More × Graduation Location	2.953	(0.034)
High School Diploma × Home Location × Graduation Location	-3.234	(0.099)
Two-Year Degree × Home Location × Graduation Location	-2.844	(0.119)
Bachelor's Degree and More × Home Location × Graduation	-2.195	(0.063)
Scale		
θ	1.425	(0.145)
Location FE	Yes	
Observations	2,502,965	
Number Individuals	26,347	

NOTES. Standard errors in parentheses.

Table 1.A.19: Transition Probabilities

	Two-Year Degree	Bachelor and More
	(1)	(2)
Entry Location	5.578 (0.004)	
Program Choice		
University	-0.301 (0.025)	2.387 (0.020)
Two-Year Elite	1.421 (0.034)	3.593 (0.030)
Two-Year Technical	2.300 (0.029)	3.295 (0.026)
Nursing	2.988 (0.039)	
Two-Year Vocational	2.719 (0.027)	2.547 (0.025)
Engineering/Business	1.350 (0.039)	3.823 (0.034)
Other	1.144 (0.040)	2.271 (0.039)
Outside Option	-0.749 (0.025)	0.678 (0.021)
Baseline Characteristics		
Female	0.061 (0.010)	0.148 (0.009)
With Scholarship	-0.089 (0.012)	-0.143 (0.011)
Medium-High SES	-0.051 (0.015)	-0.103 (0.013)
Medium-Low SES	-0.048 (0.013)	-0.165 (0.011)
Low SES	-0.062 (0.014)	-0.221 (0.013)
Medium-High GPA	-0.208 (0.025)	-0.404 (0.021)
Medium-Low GPA	-0.492 (0.024)	-0.964 (0.020)
Low GPA	-0.990 (0.023)	-1.888 (0.019)
Technical Track	-0.409 (0.012)	-1.335 (0.012)
Vocational Track	-1.195 (0.016)	-2.835 (0.020)
Province FE	Yes	
Observations	136,489,065	
Number Students	478,909	

NOTES. Standard errors in parentheses.

Figure 1.A.14: Location Choices - Moving Costs

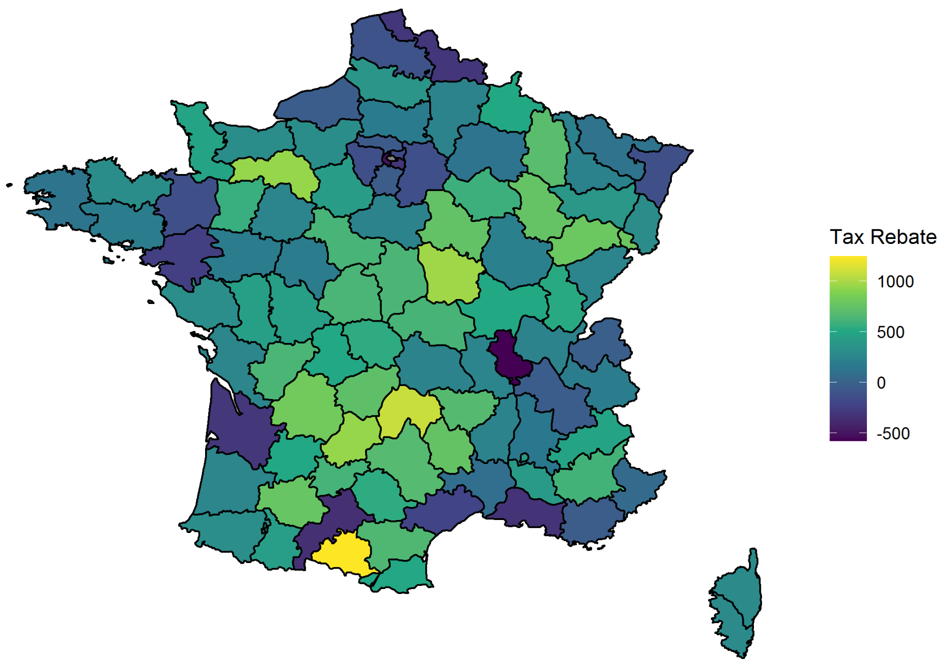

Notes. This map uses estimated preferences over local labor markets to compute the tax rebate, in monthly euros, an individual from a given province would need to receive to be willing to move back home rather than to stay in Paris, conditional on having studied in Paris and having reached at least a bachelor's degree. It corresponds to Column (2) of Tables 1.A.20-1.A.21. It abstracts from wage differences across locations.

Table 1.A.20: Compensation to Move to Paris - Three-Year Degree

	Studied Home	Studied Paris	Wage Premium Paris
	(1)	(2)	(3)
Ain	2,369.6129	-233.8147	326.3984
Aisne	2,368.6639	-234.7637	355.2361
Allier	1,985.5966	-617.8309	471.2733
Alpes de Hte Provence	2,001.6312	-601.7962	548.2294
Hautes Alpes	2,125.5078	-477.9197	665.6326
Alpes Maritimes	2,544.4763	-58.9512	330.3785
Ardèche	2,371.7680	-231.6596	607.5127
Ardennes	2,086.5887	-516.8389	401.0415
Ariège	1,358.9583	-1,244.4690	312.4867
Aube	2,022.0723	-581.3550	430.4088
Aude	1,950.3682	-653.0594	380.2688
Aveyron	1,924.8845	-678.5429	459.1324
Bouches du Rhône	2,928.8074	325.3799	417.2416
Calvados	2,289.4302	-313.9972	399.6028
Cantal	1,519.6078	-1,083.8197	362.2968
Charente	1,973.1458	-630.2817	365.7993
Charente Maritime	2,346.5526	-256.8749	513.4595
Cher	1,968.0151	-635.4125	417.8051
Corrèze	1,879.9699	-723.4577	418.7350
Corse	2,315.2206	-288.2070	622.1647
Côte d'Or	2,391.8662	-211.5612	425.7309
Côtes d"Armor	2,294.1312	-309.2964	519.4637
Creuse	2,060.5421	-542.8852	842.7283
Dordogne	1,782.8343	-820.5934	335.3272
Doubs	2,352.7405	-250.6870	339.9949
Drôme	2,455.5481	-147.8794	511.8137
Eure	2,294.8684	-308.5591	428.7651
Eure et Loir	2,175.4474	-427.9801	330.5469
Finistère	2,483.8316	-119.5958	489.8322
Gard	2,521.8089	-81.6187	494.3485
Haute Garonne	2,953.8333	350.4058	323.7789
Gers	1,832.4528	-770.9747	511.5211
Gironde	2,908.4976	305.0702	467.9077
Hérault	2,819.9723	216.5449	501.8954
Ille et Vilaine	2,734.8472	-131.4197	515.6054
Indre	1,973.4576	-629.9698	502.4423
Indre et Loire	2,419.2800	-184.1474	381.1048
Isère	2,644.9495	41.5221	440.0371
Jura	2,075.2228	-528.2046	478.8349
Landes	2,343.6603	-259.7672	490.7403
Loir et Cher	1,983.7581	-619.6693	466.3047
Loire	2,351.3026	-252.1251	418.0523
Haute Loire	1,926.1790	-677.2485	434.1151
Loire Atlantique	2,849.8087	246.3813	474.7519
Loiret	2,372.1628	-231.2647	324.5332
Lot	1,656.4006	-947.0268	372.4734
Total	2,537.3571	-66.0703	372.4073

Notes. This table presents monthly compensation payments, by home location, required for individuals to accept to enter the labor market in Paris rather than in their home location, if they studied home (Column (1)) or in Paris (Column (2)). Column (3) shows the wage premium between Paris and the other provinces.

Table 1.A.21: Compensation to Move to Paris - Three-Year Degree (Continued)

	Studied Home (1)	Studied Paris (2)	Wage Premium Paris (3)
Lot et Garonne	2,091.8518	-511.5757	471.4930
Lozère	1,851.6222	-751.8053	690.9455
Maine et Loire	2,424.8839	-178.5436	485.0795
Manche	2,126.3413	-477.0860	300.9559
Marne	2,482.1343	-121.2931	407.1520
Haute Marne	1,849.2807	-754.1468	330.7110
Mayenne	2,035.1841	-568.2434	419.8071
Meurthe et Moselle	2,383.3641	-220.0634	291.3181
Meuse	1,907.2369	-696.1907	457.0153
Morbihan	2,418.4473	-184.9802	532.3172
Moselle	2,490.5786	-112.8488	336.0319
Nièvre	1,630.8758	-972.5516	422.4279
Nord	2,912.8714	309.4440	384.7555
Oise	2,432.8463	-170.5809	195.1116
Orne	1,656.6796	-946.7479	320.8997
Pas de Calais	2,716.4795	113.0522	426.5190
Puy de Dôme	2,370.8833	-232.5442	371.9368
Pyrénées Atlantiques	2,282.8366	-320.5909	507.9615
Hautes Pyrénées	2,169.2061	-434.2213	549.2835
Pyrénées Orientales	2,096.8928	-506.5347	484.0320
Bas Rhin	2,740.4122	136.9845	340.6232
Haut Rhin	2,300.4759	-302.9515	246.9669
Rhône	3,184.0545	580.6269	390.7863
Haute Saône	1,812.4154	-791.0120	427.8329
Saône et Loire	2,086.4920	-516.9353	483.6682
Sarthe	2,354.2532	-249.1742	610.4579
Savoie	2,412.0451	-191.3823	459.0726
Haute Savoie	2,633.2440	29.8163	223.3103
Seine Maritime	2,645.9545	42.5270	353.5218
Seine et Marne	2,744.5728	141.1452	214.3282
Yvelines	2,740.9037	137.4763	122.6240
Deux Sèvres	2,158.9117	-444.5156	545.1390
Somme	2,235.0758	-368.3516	168.7533
Tarn	2,062.4704	-540.9571	493.1672
Tarn et Garonne	1,991.0916	-612.3359	469.2621
Var	2,653.2386	49.8111	530.9038
Vaucluse	2,193.4653	-409.9623	507.2162
Vendée	2,285.2377	-318.1898	574.0239
Vienne	2,162.4033	-441.0241	473.1853
Haute Vienne	2,090.1821	-513.2454	470.5154
Vosges	2,214.1920	-389.2355	517.9851
Yonne	1,858.8254	-744.6021	409.8497
Territoire de Belfort	1,817.0413	-786.3862	467.1307
Essonne	2,654.3795	50.9520	180.7714
Hauts de Seine	3,111.4564	508.0290	-20.3760
Seine St Denis	2,783.2324	179.8049	153.5027
Val de Marne	2,915.1990	311.7717	128.1345
Val d'Oise	2,661.1420	57.7145	182.9899
Total	2,537.3571	-66.0703	372.4073

NOTES. This table presents monthly compensation payments, by home location, required for individuals to accept to enter the labor market in Paris rather than in their home location, if they studied home (Column (1)) or in Paris (Column (2)). Column (3) shows the wage premium between Paris and the other provinces.

Table 1.A.22: Compensation to Move to Paris - Two-Year Degree

	Studied Home	Studied Paris	Wage Premium Paris
	(1)	(2)	(3)
Ain	3,042.4593	178.2033	94.9097
Aisne	3,041.5105	177.2544	282.6605
Allier	2,658.4430	-205.8130	331.2574
Alpes de Hte Provence	2,674.4778	-189.7785	302.0740
Hautes Alpes	2,798.3544	-65.9017	419.4772
Alpes Maritimes	3,217.3228	353.0668	84.2231
Ardèche	3,044.6145	180.3585	376.0240
Ardennes	2,759.4352	-104.8209	266.3717
Ariège	2,031.8050	-832.4512	181.7255
Aube	2,694.9191	-169.3370	295.7390
Aude	2,623.2148	-241.0414	213.7589
Aveyron	2,597.7313	-266.5247	328.3712
Bouches du Rhône	3,601.6540	737.3981	171.0862
Calvados	2,962.2767	98.0208	226.7728
Cantal	2,192.4544	-671.8016	222.2809
Charente	2,645.9922	-218.2638	93.2114
Charente Maritime	3,019.3993	155.1430	240.8716
Cher	2,640.8617	-223.3945	208.6330
Corrèze	2,552.8164	-311.4397	254.1644
Corse	2,988.0671	123.8110	113.9858
Côte d'Or	3,064.7129	200.4569	122.6522
Côtes d"Armor	2,966.9777	102.7216	225.6494
Creuse	2,733.3889	-130.8674	678.1577
Dordogne	2,455.6808	-408.5753	51.8349
Doubs	3,025.5872	161.3310	128.0431
Drôme	3,128.3947	264.1385	280.3250
Eure	2,967.7151	103.4589	227.6143
Eure et Loir	2,848.2939	-15.9622	121.3749
Finistère	3,156.6785	292.4223	196.0179
Gard	3,194.6556	330.3993	327.8386
Haute Garonne	3,626.6800	762.4239	193.0176
Gers	2,505.2995	-358.9567	380.7599
Gironde	3,581.3444	717.0881	184.4153
Hérault	3,492.8190	628.5629	335.3855
Ille et Vilaine	3,407.6937	543.4377	221.7911
Indre	2,646.3043	-217.9519	293.2704
Indre et Loire	3,092.1267	227.8706	171.9328
Isère	3,317.7960	453.5401	208.5485
Jura	2,748.0695	-116.1867	266.8831
Landes	3,016.5067	152.2507	207.2480
Loir et Cher	2,656.6049	-207.6513	257.1327
Loire	3,024.1490	159.8929	186.5636
Haute Loire	2,599.0256	-265.2304	294.0991
Loire Atlantique	3,522.6555	658.3994	136.2557
Loiret	3,045.0092	180.7533	115.3612
Lot	2,329.2474	-535.0088	241.7122
Total	3,210.2038	345.9477	194.0919

Notes. This table presents monthly compensation payments, by home location, required for individuals to accept to enter the labor market in Paris rather than in their home location, if they studied home (Column (1)) or in Paris (Column (2)). Column (3) shows the wage premium between Paris and the other provinces.

Table 1.A.23: Compensation to Move to Paris - Two-Year Degree (Continued)

	Studied Home	Studied Paris	Wage Premium Paris
	(1)	(2)	(3)
Lot et Garonne	2,764.6984	-99.5578	188.0007
Lozère	2,524.4689	-339.7871	524.4356
Maine et Loire	3,097.7306	233.4745	146.5833
Manche	2,799.1880	-65.0682	128.1259
Marne	3,154.9809	290.7249	272.4822
Haute Marne	2,522.1273	-342.1289	196.0412
Mayenne	2,708.0308	-156.2255	81.3109
Meurthe et Moselle	3,056.2107	191.9547	102.0846
Meuse	2,580.0833	-284.1727	267.7819
Morbihan	3,091.2937	227.0378	238.5029
Moselle	3,163.4253	299.1691	146.7984
Nièvre	2,303.7225	-560.5337	119.3491
Nord	3,585.7182	721.4620	253.9108
Oise	3,105.6931	241.4370	122.5359
Orne	2,329.5260	-534.7300	148.0696
Pas de Calais	3,389.3263	525.0701	295.6743
Puy de Dôme	3,043.7298	179.4737	231.9208
Pyrénées Atlantiques	2,955.6833	91.4272	224.4692
Hautes Pyrénées	2,842.0529	-22.2032	418.5222
Pyrénées Orientales	2,769.7395	-94.5168	317.5221
Bas Rhin	3,413.2586	549.0024	192.4934
Haut Rhin	2,973.3226	109.0664	98.8371
Rhône	3,856.9011	992.6450	159.2976
Haute Saône	2,485.2621	-378.9940	215.8811
Saône et Loire	2,759.3388	-104.9174	180.5895
Sarthe	3,027.0998	162.8437	271.9618
Savoie	3,084.8918	220.6357	227.5839
Haute Savoie	3,306.0904	441.8345	-8.1784
Seine Maritime	3,318.8013	454.5452	152.3710
Seine et Marne	3,417.4193	553.1633	214.3282
Yvelines	3,413.7505	549.4943	122.6241
Deux Sèvres	2,831.7585	-32.4976	272.5511
Somme	2,907.9226	43.6664	96.1777
Tarn	2,735.3169	-128.9391	362.4060
Tarn et Garonne	2,663.9383	-200.3179	338.5008
Var	3,326.0853	461.8291	284.7484
Vaucluse	2,866.3118	2.0558	261.0608
Vendée	2,958.0843	93.8284	235.5277
Vienne	2,835.2502	-29.0061	200.5974
Haute Vienne	2,763.0287	-101.2275	305.9449
Vosges	2,887.0388	22.7826	328.7516
Yonne	2,531.6721	-332.5841	106.7710
Territoire de Belfort	2,489.8878	-374.3681	255.1788
Essonne	3,327.2261	462.9701	180.7715
Hauts de Seine	3,784.3031	920.0469	-20.3760
Seine St Denis	3,456.0793	591.8230	153.5027
Val de Marne	3,588.0456	723.7896	128.1345
Val d'Oise	3,333.9886	469.7326	182.9900
Total	3,210.2038	345.9477	194.0919

Notes. This table presents monthly compensation payments, by home location, required for individuals to accept to enter the labor market in Paris rather than in their home location, if they studied home (Column (1)) or in Paris (Column (2)). Column (3) shows the wage premium between Paris and the other provinces.

Table 1.A.24: Compensation to Move to Paris - High School Diploma

	Studied Home	Studied Paris	Wage Premium Paris
	(1)	(2)	(3)
Ain	-3,324.2799	152.7627	78.7502
Aisne	-3,323.3308	151.8136	329.0525
Allier	-2,940.2634	-231.2536	255.1376
Alpes de Hte Provence	-2,956.2981	-215.2191	328.9836
Hautes Alpes	-3,080.1748	-91.3423	446.3868
Alpes Maritimes	-3,499.1432	327.6260	111.1327
Ardèche	-3,326.4350	154.9179	359.8644
Ardennes	-3,041.2555	-130.2615	280.7475
Ariège	-2,313.6253	-857.8918	33.5004
Aube	-2,976.7397	-194.7776	310.1148
Aude	-2,905.0350	-266.4821	200.3853
Aveyron	-2,879.5517	-291.9655	180.1461
Bouches du Rhône	-3,883.4744	711.9572	197.9958
Calvados	-3,244.0972	72.5801	87.6995
Cantal	-2,474.2748	-697.2423	146.1611
Charente	-2,927.8127	-243.7043	39.3705
Charente Maritime	-3,301.2196	129.7024	187.0307
Cher	-2,922.6822	-248.8352	182.9081
Corrèze	-2,834.6369	-336.8804	109.4091
Corse	-3,269.8875	98.3703	-69.6015
Côte d'Or	-3,346.5333	175.0161	84.8087
Côtes d"Armor	-3,248.7981	77.2809	172.1978
Creuse	-3,015.2093	-156.3080	533.4024
Dordogne	-2,737.5013	-434.0159	3.4210
Doubs	-3,307.4076	135.8903	51.8389
Drôme	-3,410.2152	238.6980	264.1655
Eure	-3,249.5355	78.0182	264.2183
Eure et Loir	-3,130.1144	-41.4028	95.6499
Finistère	-3,438.4989	266.9817	142.5664
Gard	-3,476.4760	304.9587	314.4651
Haute Garonne	-3,908.5004	736.9834	44.7925
Gers	-2,787.1199	-384.3973	232.5348
Gironde	-3,863.1646	691.6475	136.0014
Hérault	-3,774.6394	603.1222	322.0120
Ille et Vilaine	-3,689.5142	517.9971	168.3395
Indre	-2,928.1247	-243.3925	267.5453
Indre et Loire	-3,373.9472	202.4299	146.2078
Isère	-3,599.6165	428.0994	192.3888
Jura	-3,029.8899	-141.6272	190.6789
Landes	-3,298.3272	126.8101	158.8341
Loir et Cher	-2,938.4252	-233.0920	231.4077
Loire	-3,305.9694	134.4523	170.4040
Haute Loire	-2,880.8460	-290.6711	217.9793
Loire Atlantique	-3,804.4759	632.9587	71.0329
Loiret	-3,326.8297	155.3127	89.6362
Lot	-2,611.0679	-560.4494	93.4871
Total	-3,492.0242	320.5070	167.3803

Notes. This table presents monthly compensation payments, by home location, required for individuals to accept to enter the labor market in Paris rather than in their home location, if they studied home (Column (1)) or in Paris (Column (2)). Column (3) shows the wage premium between Paris and the other provinces.

Table 1.A.25: Compensation to Move to Paris - High School Diploma (Continued)

	Studied Home	Studied Paris	Wage Premium Paris
	(1)	(2)	(3)
Lot et Garonne	3,046.5187	-124.9983	139.5868
Lozère	2,806.2893	-365.2278	511.0620
Maine et Loire	3,379.5511	208.0337	81.3605
Manche	3,081.0085	-90.5090	-10.9474
Marne	3,436.8015	265.2843	286.8581
Haute Marne	2,803.9478	-367.5695	210.4170
Mayenne	2,989.8513	-181.6661	16.0880
Meurthe et Moselle	3,338.0311	166.5140	99.6413
Meuse	2,861.9038	-309.6133	265.3386
Morbihan	3,373.1141	201.5971	185.0514
Moselle	3,445.2458	273.7285	144.3551
Nièvre	2,585.5428	-585.9744	81.5056
Nord	3,867.5385	696.0213	249.9462
Oise	3,387.5136	215.9964	168.9280
Orne	2,611.3466	-560.1707	8.9963
Pas de Calais	3,671.1468	499.6294	291.7096
Puy de Dôme	3,325.5504	154.0332	155.8010
Pyrénées Atlantiques	3,237.5036	65.9864	176.0553
Hautes Pyrénées	3,123.8734	-47.6440	270.2972
Pyrénées Orientales	3,051.5596	-119.9575	304.1485
Bas Rhin	3,695.0791	523.5619	175.7847
Haut Rhin	3,255.1429	83.6257	82.1284
Rhône	4,138.7217	967.2043	143.1380
Haute Saône	2,767.0827	-404.4348	139.6769
Saône et Loire	3,041.1593	-130.3581	142.7460
Sarthe	3,308.9203	137.4031	206.7389
Savoie	3,366.7123	195.1950	211.4243
Haute Savoie	3,587.9109	416.3938	-24.3379
Seine Maritime	3,600.6215	429.1044	188.9750
Seine et Marne	3,699.2398	527.7227	214.3282
Yvelines	3,695.5710	524.0537	122.6241
Deux Sèvres	3,113.5789	-57.9382	218.7102
Somme	3,189.7430	18.2258	142.5697
Tarn	3,017.1373	-154.3797	214.1810
Tarn et Garonne	2,945.7586	-225.7584	190.2758
Var	3,607.9054	436.3884	311.6579
Vaucluse	3,148.1322	-23.3850	287.9704
Vendée	3,239.9048	68.3876	170.3049
Vienne	3,117.0705	-54.4468	146.7565
Haute Vienne	3,044.8490	-126.6681	161.1896
Vosges	3,168.8590	-2.6581	326.3083
Yonne	2,813.4923	-358.0248	68.9275
Territoire de Belfort	2,771.7083	-399.8088	178.9746
Essonne	3,609.0467	437.5295	180.7714
Hauts de Seine	4,066.1235	894.6063	-20.3760
Seine St Denis	3,737.8996	566.3823	153.5027
Val de Marne	3,869.8662	698.3489	128.1345
Val d'Oise	3,615.8090	444.2919	182.9899
Total	3,492.0242	320.5070	167.3803

NOTES. This table presents monthly compensation payments, by home location, required for individuals to accept to enter the labor market in Paris rather than in their home location, if they studied home (Column (1)) or in Paris (Column (2)). Column (3) shows the wage premium between Paris and the other provinces.

Table 1.A.26: Model Fit

	Data (1)	Model (2)
Panel A: Higher Education Market		
University	49.17	49.23
Two-Year Elite	8.68	9.57
Two-Year Technical	11.03	10.89
Two-Year Vocational	19.69	19.42
Nursing School	3.60	3.04
Ingeneering/Business	5.63	5.36
Other	2.20	2.49
Distance	67.21	69.88
Panel B: Work Location		
Same as Home & Study Location	40.31	40.34
Same as Home Location Only	23.55	23.56
Same as Study Location Only	9.02	8.81
Neither	27.11	27.28

Notes. This table evaluates the fit of the model by comparing statistics related to the match on the higher education market (Panel A) and to location choices at entry on the labor market (Panel B). Column (1) and Column (2) present the statistics observed in the data and in the simulated status quo, respectively.

Table 1.A.27: Model Fit: Gravity Model

	Enrollment Flows	
	Data	Model
	(1)	(2)
Distance$_{o,d}$	-2.506	-2.271
	(0.056)	
Origin FE	Yes	Yes
Destination FE	Yes	Yes
Obs.	8,930	8,930

NOTES. This table presents estimates of a gravity equation. It regresses students' enrollment flows between origin province o and destination province d on the distance (in kilometers) between the centroids o and d, controlling for origin and destination fixed effects. Column (1) shows results from an estimation using the estimation sample, while Column (2) presents results where the same model is fitted using the same sample, but where enrollment flows are determined by the province of the matched program in the simulated match. Estimation is performed using PPML, as recommended by Silva and Tenreyro (2006). Robust standard errors in parenthesis.

Table 1.A.28: Quantification - Brain Drain

	Counterfactual	Change
Panel A: No Mobility Frictions		
Work Home	0.56	-12%
From Desert - Work Hub	0.47	$+14\%$
Panel B: No Mobility Frictions + No Admission Restrictions		
Work Home	0.53	-17%
From Desert - Work Hub	0.50	$+20\%$
Panel C: No Mobility Frictions + No Moving Costs		
Work Home	0.69	$+8\%$
From Desert - Work Hub	0.45	$+9\%$

Notes. This table presents changes in the share of workers entering the labor force in their home location, and in the share of workers from higher education deserts entering the labor force in a higher education hub. Panel A presents results from a counterfactual where mobility frictions within the higher education market are removed. Panel B presents results from a counterfactual additionally eliminating admission restrictions: all students are admitted in their preferred program in stage 1. Panel C presents results from a counterfactual where both mobility frictions within the higher education as well as moving costs at the location choice stage are removed.

Table 1.A.29: Quantification - Labor Force in Paris

	Status Quo	No Mobility Frictions	No Admission Restrictions	No Moving Costs
	(1)	(2)	(3)	(4)
Share Work Paris	7.70	+11%	+44%	−9%

NOTES. This table presents changes in the share of workers entering the labor force in Paris. Column (1) presents the share of workers entering the labor market in Paris in the status quo. Column (2) presents results from a counterfactual where mobility frictions within the higher education market are removed. Column (3) presents results from a counterfactual additionally eliminating admission restrictions: all students are admitted in their preferred program in stage 1. Column (4) presents results from a counterfactual where both mobility frictions within the higher education as well as moving costs at the location choice stage are removed.

Figure 1.A.15: Policy Counterfactual - Share Repaying

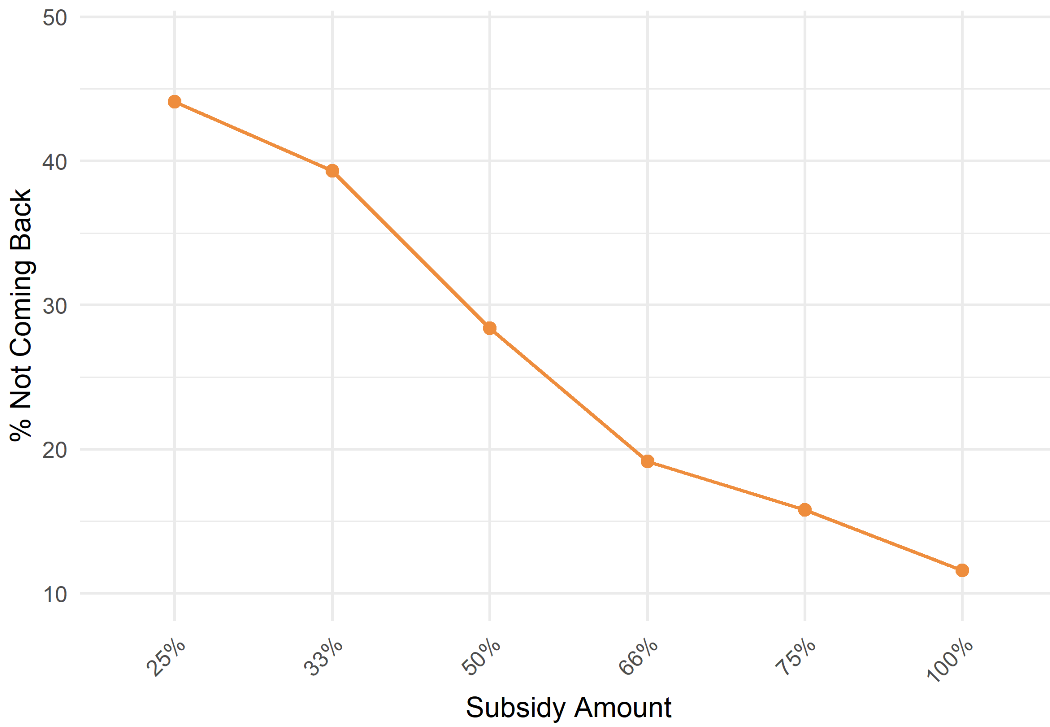

NOTES. This figure shows the results of the counterfactual in which students are offered a distance-based scholarship, which they have to reimburse if not returning home after graduation. It plots, for different values of the subsidy, the share of individuals from higher education deserts deciding to pay back the scholarship and entering the labor market outside of their home location.

1.B Context: Additional Details

1.B.1 High School Tracks and End-of-High School Exam

In the French high school system, students have three main educational tracks to choose from: general, technological, and vocational.

General track: In the general track, students can choose to specialize in one of three fields: sciences, social sciences, or humanities. While students in the general track share common classes, like English, they also have specialized courses tailored to their chosen field. For instance, a strong emphasis is made on mathematics, biology, and physics for students in the sciences field.

Technological track: This track focuses on applying scientific and technological knowledge to practical skills. Students can specialize in areas such as information technology, chemistry, or agronomy.

Vocational track: This track is designed to provide students with practical skills, with a particular focus on internships. They can specialize in various areas, such as automotive technology, hospitality, construction, agriculture, or healthcare.

The track and specialization of the student determine the centralized exam (*baccalauréat*) they take at the end of the three-year secondary education period. For example, students in the sciences general track will take tests particularly targeted at evaluating their knowledge in mathematics, biology and physics. The final score, between zero and 20, is associated to an Honors system (*Mentions)*. The passing score is 10. A score between 12 and 14 corresponds to graduating with honors (*Mention Assez Bien*). With a score between 14 and 16, the student graduates with high honors (*Mention Bien*). A student graduates with highest honors (*Mention Très Bien*) when reaching a score of 16 or above. Table 1.A.1 summarizes the share of students receiving each honors type, conditional on passing.

1.B.2 Program Types

There exist six main types of programs. Half of them are targeted at students willing to acquire a master's degree. First, students can apply to first-year university programs - *Licences.* Most universities are public, and tuition fees are very low. They amount to 260€/year and are waived if the student is eligible for a scholarship. The selectivity of such programs varies a lot, depending on its institution and on its major. These programs are designed such that students can first graduate with a bachelor's degree, after three year, and can continue to get a master's degree, after five years.

Second, students can apply to two-year elite programs, or *Classes Préparatoires aux Grandes Ecoles (CPGE).* These programs are highly selective, last two to three years and prepare students to take the competitive exams for admission to the most prestigious public and private higher education institutions (*Grandes Écoles*), such as the engineering schools *École Polytechnique* and *École des Ponts ParisTech.* They are offered by high school institutions. Tuition fees are the same as in university programs.

Third, there exist business and engineering schools, which offer admission after high school. They are often private, and their tuition fees and selectivity vary.

The three other main types of programs are targeted at providing a more technical and vocational training. First, there are technology-oriented programs, which last for two years and deliver a *Diplôme Universitaire de Technologie* (*DUT*). They are provided by public institutions (*Institut Universitaire de Technologie, IUT*). Their selectivity vary according to the institution and major, and tuition fees correspond to those in university programs.

There are also two-year vocational programs, offered by high school institutions. They are free, and are aimed at delivering a degree called *Brevet de Technicien Supérieur* (*BTS*).

Finally, there are social work and nursing schools. They can be public or private, with varying tuition fees.

1.B.3 Admission System: History

In France, the first online admission platform - 'Admission Post-Bac' (APB) - was introduced in 2002. It was specifically designed to match students to one type of programs: two-year elite programs. In 2009, the platform was extended to manage admission to other types of programs.

Within APB, students had to rank programs, before a Deferred Acceptance algorithm was run, taking into account students' ranking as well as programs' admission choices. Contrary to other programs, universities were legally forbidden from selecting students. The ranking of students by universities was determined by priorities, taking into account the rank at which the program was listed and the catchment area of the student. If over-subscribed, a lottery was run.

Since 2018, the APB platform has been replaced by Parcoursup. Two main changes were introduced. First, an admission algorithm is not used anymore to match students to programs. Students do not have to rank applications, and can receive several offers at the same time. Second, university programs are now allowed to choose applicants based on their own selection criteria.

1.C Proofs

1.C.1 Evaluation of the Moving Scholarship: Discussion of the Identifying Assumptions

Using a potential outcome framework, the assumption for the triple differences used in Section 1.3.2 to be valid writes:

$$\begin{aligned} &E\Big[\,(Y_{L,E,t}(0) - Y_{L,E,t-1}(0)) - (Y_{L,NE,t}(0) - Y_{L,NE,t-1}(0))\,\Big] \\ =&E\Big[(Y_{H,E,t}(0) - Y_{H,NE,t-1}(0)) - (Y_{H,NE,t}(0) - Y_{H,NE,t}(0))\Big] \end{aligned}$$

where $Y_{s,e,t}(0)$ denotes the average potential outcome of students with eligibility status $e \in \{E, NE\}$ living in provinces with supply level $s \in \{H, L\}$ in period t when not treated.

Under this parallel trends assumption, which can be tested, β_{2019} identifies the following object:

$$\begin{aligned} &E\Big[(Y_{L,E,2019}(1) - Y_{L,NE,2019}(0)) - (Y_{H,E,2019}(1) - Y_{H,NE,2019}(0))\Big] \\ &= ATT_{L,E,2019} - ATT_{H,E,2019} \end{aligned} \tag{1.18}$$

That is, it identifies the difference between the ATT of the eligible students living in low-supply areas, denoted $ATT_{L,E,2019}$, and the ATT of the eligible students living in high-supply areas, $ATT_{H,E,2019}$. The triple differences estimator thus identifies the differences in treatment effects of the policy in low-supply versus high-supply areas. While being eligible, it seems however safe to assume that the reform should not have impacted the application behavior of scholarship-eligible students living in high-supply areas. The scholarship was indeed very small, while these students are low-income students living in places where they can easily access higher education programs. Overall, the policy is unlikely to have an effect on students already living in a high-supply area, such that we can expect the ATT to be null in the latter locations. In that case, β_{2019} would identify the average treatment effect of the moving subsidy for students living in low-supply areas.

To limit concerns regarding a potential non-null effect of the subsidy on students part of the comparison group, I consider areas at the very top and bottom of the program supply distribution. I define low-supply areas as the five locations at the bottom of the supply distribution. High-supply areas are defined as the ten locations at the top of this distribution. The latter definition is chosen as a way to minimize the extent with which students in these areas may have benefited from the scholarship. Results are robust to restricting and extending this set. Regarding low-supply areas, the effects displayed in Table 1.A.8 are robust to extending the definition to the ten locations with the least supply, while the magnitude of the estimates is slightly smaller.

If we consider that the ATT in high-supply areas was non-zero but positive, β_τ would then identify the ATT in low-supply areas net of the one in high-supply areas. As long as the latter is positive, the fact that we observe a positive β_τ thus suggest that the treatment effect is positive for students living in low-supply areas, and larger than the one of those living in high-supply areas. If the treatment instead decreased the application distance of eligible students in high-supply areas, the difference in ATT would then be more difficult to interpret. However, there is no economic rationale for such a negative effect to arise.

1.C.2 Reverse Difference-in-Differences

Other factors than the reform discussed in Section 1.3.2 may have affected admission outcomes between 2018 and 2019. To control for them, I use a group of programs which admission criteria stayed the same over time: non-university programs, for which geographical priorities were never implemented. I thus augment the model as follows:

$$y_{j,\ell,t} = \alpha_j + \lambda_{\ell,t} + \sum_{\ell} \beta_\ell^{DD} \times D_{\ell,t} + \epsilon_{j,\ell,t} \tag{1.19}$$

where $\lambda_{\ell,t}$ captures location-specific trends in admission outcomes common to university and non-university programs. This is akin to implementing a difference-in-differences where the groups compared share the same treatment status at endline, but not at baseline (see Kim and Lee (2019) and Fabre (2023) for a discussion of reverse difference-in-differences). In particular, it compares the evolution of admission outcomes in university programs - for which institutional barriers to mobility were removed in 2019 - to those in non-university programs which never implemented such barriers.

I now discuss the identification assumptions required for such a comparison to identify the ATT of removing barriers to mobility for the group switching treatment status. Let us consider the following potential outcome framework. Program j belongs to group S if the program switched from implementing geographic barriers to mobility in period 2018 to not using such rules in 2019. This is the group switching treatment status. Program j belongs to group C if the program implemented non-geographic admission criteria in 2018 and 2019. $Y_{j,g,t}(d)$ is the potential outcome of program j belonging to group $g \in \{S, C\}$ in period t under treatment status $d \in \{0, 1\}$, where $d = 1$ under non-geographic admission criteria and zero otherwise. $Y_{g,t}(d)$ is the average potential outcome for group g in period t: $Y_{g,t}(d) = \frac{1}{n_g} \sum_{j \in g} Y_{j,g,t}(d)$, where n_g is the number of programs in group g.

In standard difference-in-differences, the treatment status of the two groups is the same at baseline and differs ex post. In contrast, here, the treatment status of the two groups differs at baseline, but is the same at endline. Kim and Lee (2019) and Fabre (2023) show that one can still identify the treatment effect of interest in this framework, under the following parallel trends assumption:

$$E\left[[Y_{S,t}(1) - Y_{S,t-1}(1)] - [Y_{C,t}(1) - Y_{C,t-1}(1)]\right] = 0 \tag{1.20}$$

That is, the potential outcomes of the two groups would have evolved similarly across time under non-geographic admission criteria. Note that, as the two groups do not have the same treatment status at baseline, one cannot test this assumption using periods prior to the implementation of the reform. However, one can test this assumption using outcomes posterior to the switch.[30]

Under this assumption, β^{DD} from Equation (1.19) identifies the ATT of non-geographic admission criteria for university programs:

$$\begin{aligned}
\beta^{DD} &= E\left[Y_{S,2019} - Y_{S,2018}\right] - E\left[Y_{C,2019} - Y_{C,2018}\right] \\
&= E\left[Y_{S,2018}(1) - Y_{S,2018}(0)\right] \\
&+ E\left[Y_{S,2019}(1) - Y_{S,2018}(1)\right] - E\left[Y_{C,2019}(1) - Y_{C,2018}(1)\right] \\
&= E\left[Y_{S,2018}(1) - Y_{S,2018}(0)\right]
\end{aligned} \tag{1.21}$$

β^{DD} thus identifies the ATT, for programs belonging to the switching group S, of removing institutional barriers to mobility.

Note that, given the nature of the market, one of the assumptions required for any

[30] I plan on testing this assumption with administrative data from the online admission platform from 2020 and onward.

difference-in-differences to be valid might not be satisfied in practice: the comparison group may be affected by the reform. Indeed, students' flows might respond to the removal of barriers to mobility by shifting away from other type of programs toward university programs located within Paris. To limit such concerns, I implement this difference-in-differences at the catchment area level. In that case, we would recover the effect of removing barriers to mobility in Paris, even if the reform would impact programs in Versailles.[31]

1.C.3 Identification of Preferences over Locations

Taking the logarithm on both sides of Equation (1.10), we have:

$$\log(\Pr(\ell^* = \ell | h_i, y_i, \boldsymbol{m}, \boldsymbol{w}_i)) = \theta(w_{i\ell} + \omega(h_i, y_i, m_\ell)) - \log(\sum_{\kappa \in \mathcal{L}} \exp(\theta(w_{i\kappa} + \omega(h_i, y_i, m_\kappa)))) \tag{1.22}$$

The same equation can be evaluated for $\ell = \ell_0$, the location choice for which we normalize $\omega(h_i, y_i, m_{\ell_0})$ to be equal to zero. Taking the difference of Equation (1.22) evaluated at location ℓ and ℓ_0, and considering $w_{i\ell_0} = 0$ we have:

$$\begin{aligned} &\log(\Pr(\ell^* = \ell | h_i, y_i, \boldsymbol{m}, \boldsymbol{w}_i)) - \log(\Pr(\ell^* = \ell_0 | h_i, y_i, \boldsymbol{m}, \boldsymbol{w}_i)) \\ &= \theta(w_{i\ell} + \omega(h_i, y_i, m_\ell)) \end{aligned}$$

where the left-hand side of the equation is directly observed in data from realized location choices. In particular, taking the derivative of this object with respect to $w_{i\ell}$ allows us to identify θ. Given θ, ω can then be identified.

1.C.4 Identification of Program-Specific Students' Score

Taking the logarithm on both sides of Equation (1.11), we have:

$$\log(\Pr(i^*_{j,r} = i | \boldsymbol{x}, \boldsymbol{h}, z_j)) = v(x_i, h_i, z_j) - \log\big(\sum_{a \in \mathcal{A}_j \setminus \{i^*_{j,k}\}_{k=1}^{r-1}} \exp(v(x_a, h_a, z_j))\big) \tag{1.23}$$

[31] For example, if an increase of students from Versailles admitted to programs in Paris is compensated by a decrease of students from Versailles admitted to non-university programs in Versailles, the comparison would still be valid. Note that it would, however, be a concern if this increase is rather compensated by a decrease of high-ability students from Versailles admitted to non-university programs in Paris.

The same equation can be evaluated for $i = 0_{\mathcal{I}}$, the outside option of the program, for which the flow utility is normalized to zero. Taking the difference of Equation (1.23) evaluated for student i and $0_{\mathcal{I}}$, we have:

$$\begin{aligned}
&\log(\Pr(i^*_{j,r} = i|\boldsymbol{x}, \boldsymbol{h}, z_j)) - \log(\Pr(i^*_{j,r} = 0_{\mathcal{I}}|\boldsymbol{x}, \boldsymbol{h}, z_j))\\
&= v(x_i, h_i, z_j) - v(x_{0_{\mathcal{I}}}, h_{0_{\mathcal{I}}}, z_j)\\
&= v(x_i, h_i, z_j))
\end{aligned}$$

where the left-hand side of the equation is directly observed in data from realized rankings of applicants, such that $v(.)$ can be recovered.

1.C.5 Identification of σ

This section details the proof of Proposition 1. Consider a student i and a student o, with $x_i = x_o$ and $h_i \neq h_o$. Consider a program j such that $u(x_i, z_j) - c(x_i, d_{ij}) = u(x_o, z_j) - c(x_o, d_{oj})$.

Let us start by considering student i. Let us take the ratio of the probabilities associated with enrolling in program j and the outside option:

$$\begin{aligned}
&\frac{\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z})}{\Pr(j \in \mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta})|x_i, h_i, \boldsymbol{z}) \Pr(\mu(i) = 0_{\mathcal{J}}|x_i, h_i, \boldsymbol{z})}\\
=&\frac{\int \frac{\exp(\frac{1}{\sigma}(u(x_i,z_j)-c(x_i,d_{ij})+\frac{1}{\sigma}\bar{U}(x_i,h_i,z_j)))}{\sum_{k\in\mathcal{F}(x_i,h_i,\eta)\bigcup 0_{\mathcal{J}}}\exp(\frac{1}{\sigma}(u(x_i,z_k)-c(x_i,d_{ik})+\bar{U}(x_i,h_i,z_k)))} g(\eta)d\eta}{\int \frac{\exp(\frac{1}{\sigma}(u(x_i,z_{0_{\mathcal{J}}})-c(x_i,d_{i0_{\mathcal{J}}})+\bar{U}(x_i,h_i,z_{0_{\mathcal{J}}})))}{\sum_{k\in\mathcal{F}(x_i,h_i,\eta)\bigcup 0_{\mathcal{J}}}\exp(\frac{1}{\sigma}(u(x_i,z_k)-c(x_i,d_{ik})+\bar{U}(x_i,h_i,z_k)))} g(\eta)d\eta}\\
=&\frac{\exp(\frac{1}{\sigma}(u(x_i, z_j) - c(x_i, d_{ij}) + \bar{U}(x_i, h_i, z_j)))}{\exp(\frac{1}{\sigma}(u(x_i, z_{0_{\mathcal{J}}}) - c(x_i, d_{i0_{\mathcal{J}}}) + \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}})))}
\end{aligned} \tag{1.24}$$

Taking the logarithm and subtracting from Equation (1.24) the same object evaluated at

(x_o, h_o), we find:

$$
\begin{aligned}
&\log\left(\frac{\Pr(\mu(i)=j|x_i,h_i,\boldsymbol{z})}{\Pr(j\in\mathcal{F}(x_i,h_i,\boldsymbol{z},\boldsymbol{\eta})|x_i,h_i,\boldsymbol{z})\Pr(\mu(i)=0_{\mathcal{J}}|x_i,h_i,\boldsymbol{z})}\right)\\
&-\log\left(\frac{\Pr(\mu(o)=j|x_o,h_o,\boldsymbol{z})}{\Pr(j\in\mathcal{F}(x_o,h_o,\boldsymbol{z},\boldsymbol{\eta})|x_o,h_o,\boldsymbol{z})\Pr(\mu(o)=0_{\mathcal{J}}|x_o,h_o,\boldsymbol{z})}\right)\\
=&\frac{1}{\sigma}(u(x_i,z_j)-c(x_i,d_{ij})-(u(x_o,z_j)-c(x_o,d_{oj})))\\
&+\frac{1}{\sigma}\Big[\big[\bar{U}(x_i,h_i,z_j)-\bar{U}(x_i,h_i,z_{0_{\mathcal{J}}}))\big]-\big[\bar{U}(x_o,h_o,z_j)-\bar{U}(x_o,h_o,z_{0_{\mathcal{J}}})\big]\Big]\\
=&\frac{1}{\sigma}\Big[\big[\bar{U}(x_i,h_i,z_j)-\bar{U}(x_i,h_i,z_{0_{\mathcal{J}}})\big]-\big[\bar{U}(x_o,h_o,z_j)-\bar{U}(x_o,h_o,z_{0_{\mathcal{J}}})\big]\Big]
\end{aligned}
$$

where the first equality stems from the fact that $u(x_i, z_{0_{\mathcal{J}}}) - c(x_i, di0_{\mathcal{J}}) = 0$, $\forall x_i$, and the second equality from the fact that $u(x_i, z_j) - c(x_i, d_{ij}) = u(x_o, z_j) - c(x_o, d_{oj})$. $\bar{U}(.)$ can be identified from the parameters identified from data over location choices and transition probabilities. We can thus express σ as a function of objects observed in the data and parameters previously identified, as long as $\bar{U}(x_i, h_i, z_j) - \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}}) \neq \bar{U}(x_o, h_o, z_j) - \bar{U}(x_o, h_o, z_{0_{\mathcal{J}}})$. This is made plausible by the fact that the home location enters in $\bar{U}(.)$, while students i and o have different home locations, $h_i \neq h_o$.

Note that, I include the distance between student i and program j in $c(.)$. The identification argument thus relies on considering a program j located at the same distance from both student i and student o's respective home location. $c(.)$ could include, on top, a home bias - i.e. a dummy capturing whether the program is located in the same province as the student's home. In this case, the identification argument would rely on finding a program j which is located out of students i and o's respective home province and equidistant to both of them. Note that the flow utility associated with a program can also depend freely on the location of the program, for example through location fixed effects.

The identification argument would also go through by using two different programs, j' and j'', as long as $u(x_i, z_{j'}) - c(x_i, d_{ij'}) = u(x_o, z_{j''}) - c(x_o, d_{oj''})$ and $\bar{U}(x_i, h_i, z_{j'}) - \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}}) \neq \bar{U}(x_o, h_o, z_{j''}) - \bar{U}(x_o, h_o, z_{0_{\mathcal{J}}})$. In particular, this requires the characteristics of programs j' and j'' to be the same, i.e. $z_{j'} = z_{j''}$. It also requires the distance between student i and program j' to be equal to the one between student o and program j''.

1.C.6 Identification of u(.) and c(.)

Taking the ratios $\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z})$ and $\Pr(\mu(i) = 0_{\mathcal{J}}|x_i, h_i, \boldsymbol{z})$, we have:

$$\begin{aligned}
&\frac{\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z})}{\Pr(j \in \mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta})|x_i, h_i, \boldsymbol{z}) \Pr(\mu(i) = 0_{\mathcal{J}}|x_i, h_i, \boldsymbol{z})} \\
=&\frac{\int \frac{\exp(\frac{1}{\sigma}(u(x_i,z_j)-c(x_i,d_{ij})+\bar{U}(x_i,h_i,z_j)))}{\sum_{k\in\mathcal{F}(x_i,h_i,\eta)\bigcup 0_{\mathcal{J}}} \exp(\frac{1}{\sigma}(u(x_i,z_k)-c(x_i,d_{ik})+\bar{U}(x_i,h_i,z_k)))} g(\eta)d\eta}{\int \frac{\exp(\frac{1}{\sigma}(u(x_i,z_{0_{\mathcal{J}}})-c(x_i,d_{i0_{\mathcal{J}}})+\bar{U}(x_i,h_i,z_{0_{\mathcal{J}}})))}{\sum_{k\in\mathcal{F}(x_i,h_i,\eta)\bigcup 0_{\mathcal{J}}} \exp(\frac{1}{\sigma}u(x_i,z_k)-c(x_i,d_{ik})+\bar{U}(x_i,h_i,z_k)))} g(\eta)d\eta} \\
=&\frac{\exp(\frac{1}{\sigma}(u(x_i, z_j) - c(x_i, d_{ij}) + \bar{U}(x_i, h_i, z_j)))}{\exp(\frac{1}{\sigma}(u(x_i, z_{0_{\mathcal{J}}}) - c(x_i, d_{i0_{\mathcal{J}}}) + \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}})))}
\end{aligned} \tag{1.25}$$

Taking logarithms, this writes:

$$\begin{aligned}
&\log(\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z})) - \log(\Pr(\mu(i) = 0_{\mathcal{J}}|x_i, h_i, \boldsymbol{z}) \Pr(j \in \mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta})|x_i, h_i, \boldsymbol{z})) \\
&= \frac{1}{\sigma}\left[u(x_i, z_j) - c(x_i, d_{ij}) + \bar{U}(x_i, h_i, z_j)\right] - \frac{1}{\sigma}\left[u(x_i, z_{0_{\mathcal{J}}}) - c(x_i, d_{i0_{\mathcal{J}}}) + \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}})\right] \\
&\iff \sigma\left[\log(\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z})) - \log(\Pr(\mu(i) = 0_{\mathcal{J}}|x_i, h_i, \boldsymbol{z}))\right] - \left[\bar{U}(x_i, h_i, z_j) - \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}})\right] \\
&= u(x_i, z_j) - c(x_i, d_{ij})
\end{aligned} \tag{1.26}$$

where $\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z})$ and $\Pr(\mu(i) = 0_{\mathcal{J}}|x_i, h_i, \boldsymbol{z})$ are directly identified from data on realized matched and $\bar{U}(x_i, h_i, z_j) - \bar{U}(x_i, h_i, z_{0_{\mathcal{J}}})$ is identified in a first step and σ is identified (see Appendix 1.C.5). $\Pr(j \in \mathcal{F}(x_i, h_i, \boldsymbol{z}, \boldsymbol{\eta})|x_i, h_i, \boldsymbol{z})$ is also identified as we identified programs' preferences, while we know the set of applicants to each program and the distribution of η. Thus, the left-hand side of this equation is observed. Normalizing $c(x_i, d_{ij}) = 0$ if $d_{ij} = 0$ ensures that we can separately identify $u(.)$ and $c(.)$.

1.D Estimation Details

1.D.1 Estimation Sample

For computational tractability, students' preferences are estimated using a random sub-sample of 10,234 students. Table 1.D.1-1.D.3 present the same summary statistics over students' characteristics and their enrollment outcomes as in Table 1.A.1 and Tables 1.A.3-1.A.4. It contrasts those statistics for the full sample and the random sub-sample, and show that all characteristics are very close in the two samples.

Table 1.D.1: Summary Statistics: Full vs. Random Sample

	Full Sample	Random Sub-Sample
	(1)	(2)
Female	0.5396	0.5479
Socio-Economic Characteristics		
With Scholarship	0.1939	0.1921
High SES	0.3139	0.3098
Medium High SES	0.1566	0.1588
Medium Low SES	0.3119	0.3155
Low SES	0.2177	0.2158
End-of-High School Exam:		
With Highest Honors	0.0946	0.0957
With High Honors	0.1744	0.1759
With Honors	0.2963	0.2958
Pass	0.4347	0.4326
High School Track:		
General HS Track	0.6257	0.6282
Technological HS Track	0.2222	0.2198
Vocational HS Track	0.1521	0.1518
Observations	505,069	10,235

NOTES. This table presents summary statistics of high school students, living in metropolitan France in 2019. Column (1) focuses on the full sample while Column (2) presents the statistics for a random sample.

Table 1.D.2: Summary Statistics: Applications, Offers & Enrollment

	Mean	St. Dev.	25^{th} Pctile	50^{th} Pctile	75^{th} Pctile
	(1)	(2)	(3)	(4)	(5)
Panel A: Full Sample					
Nber Applications	10.587	9.291	5.000	9.000	13.000
Nber Offers	4.203	4.261	2.000	3.000	5.000
Distance of Enrollment	65.368	109.780	8.809	25.979	72.981
Panel B: Random Sub-Sample					
Nber Applications	10.664	9.505	5.000	9.000	13.000
Nber Offers	4.250	4.332	2.000	3.000	5.000
Distance of Enrollment	66.701	113.148	9.145	26.145	74.1497

NOTES. This table presents summary statistics regarding applications, offers and enrollment of students in high school in 2019, living in metropolitan France. Characteristics of this sample are summarized in Table 1.A.1. Panel A and Panel B show these statistics for the full sample and random sub-sample, respectively.

Table 1.D.3: Summary Statistics: Program of Enrollment

	Full Sample	Random Sub-Sample
	(1)	(2)
Three-Year University	0.3920	0.3904
Two-Year Elite	0.0700	0.0706
Engineering/Business	0.0461	0.0481
Two-Year Technical	0.0896	0.0916
Two-Year Vocational	0.1685	0.1591
Nursing	0.0262	0.0275
Other	0.0178	0.0197
Unmatched	0.1898	0.1930
Observations	505,069	10,235

NOTES. This table presents summary statistics regarding the program of enrollment of students in high school in 2019, living in metropolitan France. Column (1) focuses on the full sample, while Column (2) presents the statistics for a random sample.

1.D.2 Estimation of Transition Probabilities

I use the universe of applicants to the online platform in 2019, regardless of their match status, and match them to the higher education census thanks to the student unique identifier. I am able to match 82% of students. 94% of students matched to a program on Parcoursup in 2019 are recovered in the higher education census dataset. This implies that 6% of the students who are matched to an on-platform program either do not register to their matched program, or do so but cannot be recovered in the census.

Note that the latter dataset allows me to track some unmatched students, as I observe whether i) they match to one of the few programs off-platform but present in the census or ii) whether they match to an on-platform program in the following years, implying they participated again to the admission process. Given that the administrative census gathers the large majority of higher education programs in France, I consider that a student who was unmatched on the platform in 2019 and who cannot be found in the censuses of the years 2019-2021 stayed out of the higher education system. I assign the home location of these students as their graduation location.

1.D.3 Extension: Adding Unobserved Heterogeneity

I augment the model in order to capture persistent unobserved preferences over locations over time. It may indeed be important to take into account such unobserved preferences when modeling location choices, as they often display persistence. Allowing for unobserved heterogeneity would help disentangle the stickiness in location choices coming from unobserved preferences from state dependence (Heckman, 1981). I thus augment the model in the second stage as follows:

$$\begin{aligned}\omega(h_i, \alpha_i, y_i, m_\ell) = \mathbf{Degree}_i' \boldsymbol{\phi}_{\mathbf{1}} \mathbb{1}\{\ell = \ell_i^G, \ell = \ell_i^H\} + \mathbf{Degree}_i' \boldsymbol{\phi}_{\mathbf{2}} \mathbb{1}\{\ell \neq \ell_i^G, \ell = \ell_i^H\} \\ + \mathbf{Degree}_i' \boldsymbol{\phi}_{\mathbf{3}} \mathbb{1}\{\ell = \ell_i^G, \ell \neq \ell_i^H\} + \alpha_i a_\ell\end{aligned} \tag{1.27}$$

where a_ℓ is an amenity index, and $\alpha_i \sim \mathcal{N}(0, \zeta)$. α_i is a persistent unobserved idiosyncratic taste for amenities.

Given the parameters of Equation (1.27) and a draw of α_i, $\forall i$, one can compute the expected discounted continuation value associated with choosing the optimal location in the second stage. Further assuming full persistence of the taste for the amenity index, the utility from enrolling in program j writes:

$$\mathcal{U}_{ij} = u(x_i, z_j) - c(x_i, d_{ij}) + \alpha_i a_j + \sigma\epsilon_{ij} + \bar{U}(x_i, h_i, \alpha_i, z_j)$$

where a_j is the amenity index associated with the location of program j. The probability that student i matches with program j simply rewrites:

$$\Pr(\mu(i) = j|x_i, h_i, \boldsymbol{z}) = \int \Pr(\mu(i) = j|x_i, h_i, \alpha_i, z) f(\alpha_i|\zeta) d\alpha_i$$

where:

$$\Pr(\mu(i) = j|x_i, h_i, \alpha_i, \boldsymbol{z}) = \int \Pr(\mu(i) = j|\mathcal{F}(x_i, h_i, \boldsymbol{\eta}), x_i, h_i, \alpha_i, \boldsymbol{z}, j \in \mathcal{F}(x_i, h_i, \boldsymbol{\eta}))g(\boldsymbol{\eta})d\boldsymbol{\eta}$$

where:

$$\begin{aligned}
&\Pr(\mu(i) = j|\mathcal{F}(x_i, h_i, \boldsymbol{\eta}), x_i, h_i, \alpha_i, \boldsymbol{z}, j \in \mathcal{F}(x_i, h_i, \boldsymbol{\eta})) \\
&= \frac{\exp(\frac{1}{\sigma}(u(x_i, z_j) - c(x_i, d_{ij}) + \alpha_i a_j + \bar{U}(x_i, \alpha_i, h_i, z_j)))}{\sum_{k \in \mathcal{F}(x_i, h_i, \boldsymbol{\eta}) \bigcup 0_{\mathcal{J}}} \exp(\frac{1}{\sigma}(u(x_i, z_k) - c(x_i, d_{ik}) + \alpha_i a_k + \frac{1}{\sigma}\bar{U}(x_i, \alpha_i, h_i, z_k)))}
\end{aligned}$$

Estimation of preferences in this framework would be more computationally costly, as it would require to jointly estimate preferences in stages one and two.

1.E Counterfactuals: Details

1.E.1 Counterfactuals: Simulation Details

This section describes how both the higher education market equilibrium and individuals' location choices are computed in counterfactual exercises.

Higher Education Market Equilibrium I use the estimated parameters and a draw of idiosyncratic shocks to simulate students' utility associated to each program, from Equation (1.4), as well as their program-specific score, from Equation (1.6).

In particular, when removing mobility frictions within the higher education system, the utility of student i from enrolling in program j becomes:

$$\mathcal{U}_{ij}^{\mathcal{C}^1} = u(x_i, z_j) + \sigma\epsilon_{ij} + \bar{U}(x_i, h_i, z_j) \tag{1.28}$$

This allows us to generate each student's ranking of programs, as well as each program's ranking of students. I feed these two rankings, together with programs' number of seats, into a student-proposing Deferred Acceptance algorithm, which matches each student to a unique program. This algorithm allows us to derive the student-proposing stable match. The properties of this market entails that it essentially has a unique stable equilibrium.

Higher Education Market Equilibrium - Change in Second Stage Payoffs Note that some counterfactual exercises directly modify payoffs in the second stage of the model — such as the counterfactual where mobility costs between the individual's graduation and home location are eliminated. In such cases, one needs to update the continuation value associated with each program before computing students' utility from Equation (1.4). I thus start by computing the new utility from entering each location ℓ in stage 2 from Equation (1.3), before using it to compute the log-sum formula, which is used to derive the continuation value. Taking the example of the counterfactual where mobility costs between the individual's graduation and home location are eliminated, denoted $\mathcal{C}^2$, the new utility from entering the labor market in location ℓ for individual i writes:

$$\mathcal{W}_{i\ell}^{\mathcal{C}^2} = w_{i\ell} + \omega(h_i, y_i, m_\ell) + \mathbf{Degree}_i^{'}\boldsymbol{\phi}_3\mathbb{1}\{\ell = \ell_i^H, \ell_i^H \neq \ell_i^G\} + \frac{1}{\theta}\nu_{i\ell} \tag{1.29}$$

where the term $\mathbf{Degree}_i^{'}\boldsymbol{\phi}_3\mathbb{1}\{\ell = \ell_i^H, \ell_i^H \neq \ell_i^G\}$ is added to remove moving costs between the graduation location of individual i, ℓ_i^G, and her home location, ℓ_i^H.

The continuation value of a program with characteristics z_j now becomes:

$$\bar{U}^{\mathcal{C}^2}(x_i,h_i,z_j) = \int \beta^{t(y_i)}\mathbb{E}\left[\max_{\ell\in\mathcal{L}} \mathcal{W}_{i\ell}^{\mathcal{C}_2}|y_i\right] f\left(y_i|x_i,h_i,z_j\right)dy_i \tag{1.30}$$

In this case, we can then derive each student's ranking of programs by computing the following object, for each i:

$$\mathcal{U}_{ij}^{\mathcal{C}^2} = u(x_i,z_j) + \sigma\epsilon_{ij} + \bar{U}^{\mathcal{C}^2}(x_i,h_i,z_j)$$

Higher Education Market Equilibrium - No Admission Restrictions One of the counterfactual exercises eliminates mobility frictions, while preventing programs from rejecting any students. In this case, each student can be admitted to their preferred program, regardless of capacity constraints and programs' preferences. The equilibrium is such that each student i matches to a program $\mu(i)$, such that:

$$\mu(i) = \arg\max_{\mathcal{J}\bigcup 0_{\mathcal{J}}} \mathcal{U}_{ij}^{\mathcal{C}^1}$$

Location Decisions at Entry on the Labor Market Once the new higher education market equilibrium is simulated, one can compute the counterfactual transition probabilities, and use Equation (1.3) to simulate the counterfactual location choice of individual i.

Note that all idiosyncratic shocks are kept fixed across the different simulation exercises, and are the same as the ones used to simulate the status quo paths.

1.E.2 Counterfactuals: Outcomes

Inequalities Between Students with Different Home Locations: Each counterfactual exercise is evaluated by comparing outcomes between the new equilibrium and the status quo simulation. In order to capture inequalities across individuals with different home location, I compute the share of individuals from a given province graduating with a bachelor's degree or more as follows:

$$s_\ell^{Home} = \frac{\sum_i \mathbb{1}\{\ell_i^H = \ell, \text{Degree}_i = \text{Bachelor}\}}{\sum_i \mathbb{1}\{\ell_i^H = \ell\}}$$

where $\mathbb{1}\{\ell_i^H = \ell, \text{Degree}_i = \text{Bachelor}\}$ is a dummy equal to one if ℓ is the home location of individual i and individual i graduates with at least a bachelor's degree. $\mathbb{1}\{\ell_i^H = \ell\}$ is a dummy equal to one if ℓ is the home location of individual i. I thus have 96 observations of the variable s_ℓ^{Home}, one for each province: I compute the variance of these values.

On top of looking at the variance, I also compare the share of individuals from higher education deserts and hubs graduating with at least a bachelor's degree as follows:

$$\frac{\sum_{i \in \{i:\ell_i^H \in Hubs\}} \mathbb{1}\{\text{Degree}_i = \text{Bachelor}\}}{n_{Hubs}} - \frac{\sum_{i \in \{i:\ell_i^H \in Deserts\}} \mathbb{1}\{\text{Degree}_i = \text{Bachelor}\}}{n_{Deserts}}$$

where $\mathbb{1}\{\text{Degree}_i = \text{Bachelor}\}$ is a dummy equal to one if individual i graduates with at least a bachelor's degree, *Hubs* is the set of education hubs, *Deserts* is the set of education deserts, n_{Hubs} is the number of individuals from higher education hubs and $n_{Deserts}$ is the number of individuals from higher education deserts.

Additionally, I compute the mean welfare of students in different provinces as follows:

$$\bar{W}_\ell = \frac{\sum_{i \in \{i:\ell_i^H = \ell\}} \mathcal{U}_{ij^*}}{\sum_i \mathbb{1}\{\ell_i^H = \ell\}}$$

where j^* correspond to the program where student i is matched in the scenario considered.

I compute the variance of this measure across provinces ℓ. I also contrast the mean welfare of students from hubs and deserts:

$$\frac{\sum_{i \in \{i:\ell_i^H \in Hubs\}} \mathcal{U}_{ij^*}}{n_{Hubs}} - \frac{\sum_{i \in \{i:\ell_i^H \in Deserts\}} \mathcal{U}_{ij^*}}{n_{Deserts}}$$

where *Hubs* is the set of education hubs, *Deserts* is the set of education deserts, n_{Hubs} is the number of individuals from higher education hubs and $n_{Deserts}$ is the number of individuals from higher education deserts.

Inequalities across Local Labor Markets: In order to capture inequalities across local labor markets, I first look at the composition of their labor force. To do that, I compute the share of individuals who graduated with at least a bachelor's degree and who entered the labor market in province ℓ, that is:

$$s_\ell^{Work} = \frac{\sum_i \mathbb{1}\{\ell_i^* = \ell, \text{Degree}_i = \text{Bachelor}\}}{\sum_i \mathbb{1}\{\ell_i^* = \ell\}}$$

where $\mathbb{1}\{\ell_i^* = \ell, \text{Degree}_i = \text{Bachelor}\}$ is a dummy equal to one if ℓ is the work location of individual i and individual i graduates with at least a bachelor's degree. $\mathbb{1}\{\ell_i^* = \ell\}$ is a dummy equal to one if ℓ is the work location of individual i. To measure inequalities across local labor markets, I compute the variance of s_ℓ^{Work} across the 96 provinces of metropolitan France. Additionally, I also compare the share of workers in hubs and deserts as follows:

$$\frac{\sum_{i \in \{i:\ell_i^* \in Hubs\}} \mathbb{1}\{\text{Degree}_i = \text{Bachelor}\}}{L_{Hubs}} - \frac{\sum_{i \in \{i:\ell_i^* \in Deserts\}} \mathbb{1}\{\text{Degree}_i = \text{Bachelor}\}}{L_{Deserts}}$$

where L_{Hubs} is the number of entry-level workers entering the labor market in a higher education hub, and $L_{Deserts}$ is the number of entry-level workers entering the labor market in a higher education desert.

The second measure I consider to capture inequalities across local labor markets is the sum of the wages of the entry-level workers in a given location:

$$Y_\ell = \sum_{i \in \{i:\ell_i^* = \ell\}} w_{i\ell}$$

I compute the variance of Y_ℓ across provinces, and additionally compare the sum of wages of entry-level workers in hubs and deserts:

$$\sum_{\ell \in Hubs} Y_\ell - \sum_{\ell \in Deserts} Y_\ell$$

1.E.3 Additional Policy Counterfactual

Distance-Based Subsidies and Commitment to Come Back As an alternative policy counterfactual, I consider contracts offering a distance-based scholarship conditional on coming back, with perfect enforcement. In particular, conditional on accepting the contract, students have to enter the labor market in their home location and cannot deviate.

To evaluate such a policy, I simulate the counterfactual match of a student by replacing the lifetime utility associated with a given program j, $\mathcal{U}_{ij}$, by the utility associated with this program under the contract. In particular, the subsidy compensates for (parts of) the disutility from traveling distance d_{ij}. Regarding the continuation value, it captures the fact that the student cannot choose where to enter the labor market anymore, as they commit to come back home. The utility from enrolling in program j under such a contract thus writes:

$$u(x_i, z_j) - (1-\tau)c(x_i, d_{ij}) + \sigma\epsilon_{ij} + \frac{1}{\theta}\int \beta^{t(y_i)}(\gamma + \theta(w_{i\ell^h} + \omega(h_i, y_i, m_{\ell^H})))f(y_i|x_i, h_i, z_j)dy_i \tag{1.31}$$

where $\tau \in [0,1]$ is the percentage of mobility frictions that the subsidy covers the cost for. Students take into account that, when taking this scholarship, they have to enter the labor market in their home location. I simulate students' match outcome in stage one when their utility is replaced by Equation (1.31), and denote the welfare associated to this choice $\mathcal{U}^{\mathcal{C}_3}(\tau)$, $\forall i$. The welfare of individual i associated to the choice made under the status quo is denoted $\mathcal{U}^{SQ}$.

To evaluate the impact of such a contract on the welfare of students living in low opportunities areas, I thus contrast the average utility under the contract with subsidy τ to the one under the status quo as follows:

$$\Omega_{\mathcal{D}}^{\mathcal{C}_3}(\tau) = \frac{\sum_{i\in\{i:\ell_i^H \in Deserts\}}\left[\mathcal{U}_i^{\mathcal{C}_3}(\tau) - \mathcal{U}_i^{SQ}\right]}{n_{Deserts}} \tag{1.32}$$

As long as $\mathcal{U}_i^{\mathcal{C}_3}(\tau) \geq \mathcal{U}_i^{SQ}$, student i is at least better off under the contract than in the status quo. The student would thus accept to tie her hands in the second stage to benefit from the subsidy. In contrast, if $\mathcal{U}_i^{\mathcal{C}_3}(\tau) < \mathcal{U}_i^{SQ}$, the student would not accept the contract, as she would be better off in the status quo. For the student to accept the contract, higher education deserts would need to compensate for this loss in welfare by offering tax rebates to students once they return home. I compute the tax rebates as follows:

$$\mathcal{T}_i^{\mathcal{C}}(\tau) = \begin{cases} \mathcal{U}_i^{SQ} - \mathcal{U}_i^{\mathcal{C}^3}(\tau) & \text{if } \mathcal{U}_i^{\mathcal{C}_3}(\tau) < \mathcal{U}_i^{SQ} \\ 0 & \text{if } \mathcal{U}_i^{\mathcal{C}_3}(\tau) \geq \mathcal{U}_i^{SQ} \end{cases} \tag{1.33}$$

Taking into account these tax rebates, I finally compute the average gain in terms of wages from making one student come back home thanks to the contract. To do that, I take into account that a student taking the contract may have decided to come back under the status quo. The average wage returns for low opportunity areas is thus defined as follows:

$$\begin{aligned} \mathcal{G}_{\mathcal{D}}^{\mathcal{C}_3}(\tau) = & \frac{1}{n_{Deserts}} \sum_{i \in \{i : \ell_i^H \in Deserts\}} \Bigg[\left(\int w_{i,\ell_i^H} f(y_i | x_i, h_i, z_{ij}^{\mathcal{C}^3}) dy_i \right) \\ & - \left(\int w_{i,\ell^{\mathcal{SQ}}(y_i)} \mathbb{1}\{\ell^{\mathcal{SQ}}(y_i) = \ell_i^H\} f(y_i | x_i, h_i, z_{ij}^{\mathcal{SQ}}) dy_i \right) - \mathcal{T}_i^{\mathcal{C}^3}(\tau) \Bigg] \end{aligned} \tag{1.34}$$

where $z_{ij}^{\mathcal{C}^3}$ and $z_{ij}^{\mathcal{SQ}}$ are the vectors of characteristics of the program chosen by student i under the contract and in the status quo, respectively. $\ell^{\mathcal{SQ}}(y_i)$ corresponds to the location at entry on the labor market that an individual with characteristics y_i, collecting the degree and location at graduation, would choose in the status quo. The first integral thus corresponds to the wage that an individual with characteristics y_i would earn when returning home. The second term of this equation takes into account that, with some probability, the student would have already chosen to return home upon graduation in the status quo. In this case, the contract would not benefit education deserts, as they would not manage to attract an additional worker through the contract. Finally, the last term takes into account that, in some cases, places will have to offer a tax rebate to students for them to accept the contract.

Figure 1.E.1 plots the average gain of this policy for students and places, as well as its cost for different values of τ, i.e. more or less generous subsidies. Offering the scheme covering 100% of the disutility associated with distance would allow the average student to benefit from an increase in welfare equivalent to 763 euros/month compared to the status quo. Such a student would not need any additional tax rebates to accept the contract. In this context, making one individual come back would allow education deserts to receive an additional wage of 822 euros/month on average. The cost associated with the distance-based subsidy is however very large, being equivalent to 2,816 euros/month.

Figure 1.E.1: Counterfactual: Distance-Based Scholarship Conditional on Coming Back Home

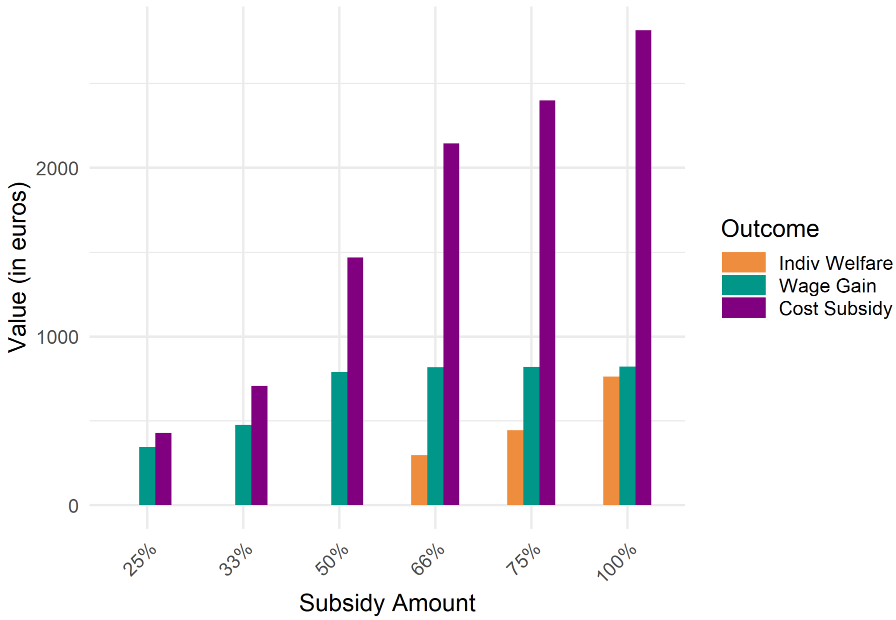

Notes. This figure shows the results of the counterfactual in which students are offered a contract under which they benefit from a distance-based scholarship, under the condition of coming back home upon graduation. It plots, for different values of the subsidy, the average welfare gain for a student taking the contract relatively to the status quo (in orange), the average wage gain for educational deserts from offering the scheme to one student (in green), and the total cost of the subsidy (in purple). All values are expressed in equivalent monthly euros.

As τ decreases, making the subsidy less generous, the cost associated with this policy decreases. The welfare gain for students born in educational deserts, however, also drops: when $\tau = \frac{2}{3}$, the average student accepting the contract would only get a welfare increase equivalent to 297 euros/month compared to the status quo. For subsidies covering 50% of the moving costs or less, tax rebates have to be offered for the average student to be willing to return home upon graduation. This implies that the net wage gain for low opportunity areas to make one student come back decreases. It reaches 344 euros/month when $\tau = 0.25$.

Chapter 2

College Admission Mechanisms and the Opportunity Cost of Time

with Olivier de Groote, Margaux Luflade and Arnaud Maurel

Abstract. College admission platforms aim at achieving a balance between avoiding congestion and allowing for ex-post flexibility in students' matches. The latter is crucial as the existence of off-platform options implies that some students will drop out of the platform in favor of their outside option, freeing up seats in on-platform programs. Sequential assignment procedures introduce such flexibility, by creating a dynamic trade-off for students: they can choose to delay their enrollment decision to receive a better offer later, at the cost of waiting before knowing their final admission outcome. We quantify this trade-off and its distributional consequences in a setting in which waiting costs can be heterogeneous across socio-economic groups. To do so, we use administrative data on rank-ordered lists and waiting decisions from the French college admission system to estimate a dynamic model of application and acceptance decisions. We find that waiting costs are a key determinant of the timing of students' acceptance decisions and of their final assignment. Nevertheless, we find substantial, but unequal, welfare gains from using a multi-round system.

2.1 Introduction

Centralized college and school admission platforms are increasingly used around the world. This shift away from decentralized mechanisms has been driven by theoretical work (Abdulkadiroğlu and Sönmez, 2003), which has put forward the desirable properties of centralized allocation systems, in particular in avoiding congestion problems (Roth and Xing, 1997). However, real-world implementations of such centralized solutions often fail to deliver all the benefits which could theoretically be achieved. In particular, solving congestion problems by sending unique offers to students comes at the cost of a lack of ex-post flexibility, stemming from the existence of off-platform options (Kapor et al., 2020b). Their presence implies that a number of students will end up rejecting the offer they received in favor of their outside option. This frees up seats in programs which might be preferred by some students to their current assignment. Welfare losses arise both for programs which end up with unfilled seats, and for students who are in a situation of justified-envy.

In this paper, we investigate the impact of allowing for ex-post flexibility in centralized matching mechanisms through the implementation of multiple single-offer rounds. In such systems, students first submit their rank-ordered list (ROL) to the platform, and then receive a unique offer, which they can choose to accept immediately, or hold on to, waiting for better opportunities in one of the following rounds. Despite being costless within the platform, the opportunity to wait for better offers comes at the cost of delaying the final admission decision. If students face some cost of waiting, sequential mechanisms thus introduce a dynamic trade-off. Students have to weigh the utility of receiving a potential better offer later against the disutility of waiting, which might encompass, for example, the cost of delaying the search for a student job or for an accommodation in the student's new location. The consequences of this dynamic trade-off on students' welfare are unclear, potentially heterogeneous, and call for an empirical investigation. On the one hand, the presence of a waiting cost might allow to take into account the strength of students' preferences even in a mechanism that allocates students to seats using a Deferred Acceptance algorithm. Indeed, only students who like a program very much would be willing to wait for the possibility to receive an offer. This might improve the quality of the resulting matches. On the other hand, sequential mechanisms might create inequalities if some students face a larger opportunity cost of time. This could in particular be the case of low socio-economic status (SES) students who may be credit constrained or face additional search frictions.

We quantify this equity-efficiency tradeoff using administrative data from the three-round French Admission Post-Bac (APB) system, which was used until 2018 to assign high-school graduates to public French universities. Data documenting applicants' decisions at each of round of this procedure, along with a dynamic model of applications and acceptance decisions, allow us to separately identify preferences for different programs and costs from waiting for better offers, following a two-step approach. In a first step, we use the ROL of students to identify the differences in utility between different programs for different students. In a second step, we estimate the probability for students to receive a particular offer, and recover the waiting costs and drop out utility by rationalizing their choices to accept, delay or drop out within a dynamic framework. Detailed information on ROLs and choices in the different rounds allows us to account for both observed and unobserved student heterogeneity. We propose a tractable Expectation-Maximization based estimation method by adapting the Conditional Choice Probability (CCP) approach to our context (Arcidiacono and Miller, 2011).

Estimated preferences for programs and waiting costs can be used in counterfactual exercises to explore the distributional consequences of waiting costs, and the welfare effects of sequential procedures. Notably, comparing total student welfare as well as the welfare gap between low-SES and high-SES students under the baseline procedure and a single-round assignment procedure makes it possible to quantify the equity-efficiency trade-off brought by the sequential admission procedure, between increased match quality and heterogeneous costs of waiting.

We obtain three main results. First, while students do prefer nearby college programs, being able to accept the offer early strongly reduces that effect. Second, low-SES students incur large costs of waiting. These costs are of a similar magnitude as the utility differences between colleges, thereby causing sub-optimal matches. These findings may reflect difficulties in finding affordable housing. Third, despite the aforementioned frictions, all groups benefit from a sequential matching mechanism, compared to a one-round mechanism. On average, the current sequential system provides a gain equivalent to enrolling 299 kilometers closer to home. However, the gains are not equally distributed. The two highest SES groups gain the equivalent of 299 to 333 kilometers, while the two lowest experience a somewhat smaller gain, equivalent to 259 to 275 kilometers.

While we do not directly simulate choices under a decentralized admission system, our study also speak more generally to the arbitrage between decentralized and centralized admission systems. Taking into account dynamic considerations and potential heterogeneity in waiting costs across students helps highlight other dimensions of the students' application decisions which are relevant to the choice of the admission mechanism. For example, heterogeneity in waiting cost would underline that decentralized admission systems might be particularly detrimental to disadvantaged students. This is especially true given the fact that academic ability and parental income are often positively correlated, meaning that, in decentralized markets, high-SES students will receive multiple offers early, and subsequent offers to low-SES students will only be made once high-SES students will have finalized their admission decision. Low-SES students might then be pushed to accept the lower quality early offers they received while they could have been admitted to a preferred option, or decide to wait but can only find low quality affordable accommodations.[1]

Related literature. A couple of other recent studies have investigated the properties of sequential assignment mechanisms, that is, mechanisms in which applicants are brought to interact multiple times with the clearinghouse (Bó and Hakimov, 2016; Luflade, 2018; Chen and Pereyra, 2019; Grenet et al., 2019). Focusing on the timing of offer acceptance, our paper is closest to Grenet et al. (2019), who find that students participating in the German university admission system are more likely to accept the earliest offer they receive. They rationalize this fact using a model of learning: at the time of application, students imperfectly know programs' qualities; at each period before the final-acceptance period, they decide whether to pay a cost to learn about one program's quality and, if so, optimally decide which program to learn about; in a sequential-offer setting, students are more likely to learn

[1] In fact, this is exactly what was feared when the French Ministry of Education decided to switch from APB, a centralized admission systems with three rounds of admission, to Parcoursup, a decentralized platform where students can receive multiple offers, in 2018. A report from the French Assemblée Nationale, in 2017, indeed warned about the fact that, in Parcoursup, disadvantaged students would be harmed as they will be more likely to receive offers later and would then face difficulties to find affordable accommodations. Some newspaper articles confirm that the platform change has triggered issues for students assigned very late in the procedure to find an accommodation before the academic year starts.(see `https://www.marianne.net/societe/parcoursup-la-recherche-d-un-logement-ou-l-autre-galere-des-etudiants-en-attente`, or `https://www.liberation.fr/france/2018/05/29/quand-parcoursup-complique-la-recherche-d-un-logement-social$_$1655070/`). This paper provides a framework allowing to partly capture such considerations.

first about (and then choose) the program they receive an offer from.
We explore an alternative explanation and investigate the role of waiting costs in decision-making in a sequential-offer mechanism. We estimate a dynamic model of application and offer acceptance to evaluate the welfare effects of the sequential procedure.

Our paper also examines the consequences of the co-existence of a centralized admission platform and off-platform options. As such, our analysis complements recent work by Kapor et al. (2020b), who study the welfare impacts of off-platform options in Chile, together with the presence of aftermarket frictions generated by the decentralized system of waitlists. They quantify the welfare cost of the presence of off-platform options through a counterfactual scenario where additional programs are added to the platform. While this counterfactual scenario can be understood as a first-best solution, it might be difficult in practice to incentivize programs to join the centralized platform.[2] In contrast, we focus in this paper on the evaluation of mechanisms with sequential rounds as an easy-to-implement solution aimed at mitigating the cost of off-platform options.

More broadly, and beyond the school choice setting, this paper adds to a small but growing set of papers modeling the dynamic considerations induced by centralized assignment mechanisms (Agarwal et al., 2021; Waldinger, 2021; Larroucau and Ríos, 2021). We contribute to this literature by quantifying and evaluating the welfare consequences of the dynamic trade-off that students face, in the presence of waiting costs, when deciding whether to accept an admission offer or instead delay acceptance by participating to the next round.

Our paper also contributes to the large and growing literature on the determinants of students' higher education choices, in particular joint choice of institution and field of study (see Altonji et al., 2016; and Patnaik et al., 2020 for reviews of this literature). We contribute to this literature by leveraging the sequential nature of the centralized APB college admission platform to quantify preferences for the different programs and the cost of waiting for the opportunity to receive an offer from a preferred program. This, in turn, allows us to explore the welfare implications of the dynamic trade-offs associated with this type of sequential college assignment mechanism. By doing so, our paper highlights the important role played by external factors and constraints, such as the ones generating these waiting costs, in students'

[2] In France, a variety of institutions, from art schools to engineering schools, decided to stay off-platform, for example to preserve the flexibility in their admission calendar and decisions.

higher education decisions.[3]

The remainder of the paper is organized as follows. Section 2.2 describes the French higher education system and its centralized college admission procedure, Admission Post-Bac (APB). Section 2.3 presents the data and some descriptive evidence of the dynamic trade-off faced by students. Section 2.4 delineates our structural model of students' dynamic application decisions, while Section 2.5 discusses our identification and estimation strategy. Section 2.6 presents the estimation results as well as the counterfactual exercises we perform. Finally, Section 2.7 concludes.

2.2 The French Higher Education System and the College Admission Procedure

Primary, secondary, and higher education in France are centralized and regulated by the Ministry of Education. The curriculum taught is the same across all schools, and there is no formal tracking of students until the beginning of high school (ninth to twelfth grades), when students make two decisions. First, they choose their *track* among three (*general*, *technological* or *vocational*), which will determine their post-secondary education choice set upon high school graduation. Second, within each track, students choose a *high-school major*, which will determine their exact high-school curriculum. Our empirical analysis focuses on students in the *general* track, who can choose one of three high-school majors: sciences, humanities, and social sciences. At the end of twelfth grade, students take a national exam (*Baccalauréat*). The content of the exam is determined by students' choice of track and high-school major, and obtaining a passing grade is necessary to graduate high-school and be eligible to enter higher education.

A semi-centralized system. For the purpose of understanding applications and admissions to post-secondary programs, the French system can be described as semi-centralized. On the one hand, a majority of programs participate in a nationwide centralized assignment procedure operated by the Ministry of Higher Education. This procedure involves students submitting rank-ordered lists to a clearinghouse via an online platform, programs assigning

[3] See also recent work by Humphries et al. (2023) who study and estimate a model of college investments in the presence of a centralized college application process.

priorities to their applicants, and the clearinghouse using a pre-specified algorithm to assign students to programs. On the other hand, other programs operate "off-platform", that is, collect applications and decide on admissions in a decentralized manner. Our analysis and model in Section 2.4 focuses on students' choices as they participate in the centralized procedure, and we simply take as given the existence of off-platform options. However, off-platform options matter for the outcome of the centralized procedure: receiving an offer from an off-platform program can induce students to decline the assignment received from the centralized procedure, thereby creating a vacancy that, in the case of a multi-round mechanism like the French APB, can be offered to another on-platform applicant. There is no unique or definitive criterion that determines whether or not a program participates in the centralized admission procedure, but off-platform programs generally pertain to one of three groups (Cour des Comptes, 2017): programs accredited by an agency other than the Ministry of Higher Education; programs jointly offered by multiple institutions (i.e. "double diplomas"); and programs wishing to retain full control over their admission calendar and decisions. Table 2.2.1 provides an overview of the different types of higher-education programs in France. The top and middle panels of Table 2.2.1 gather programs participating in the centralized admission procedure, which we discuss in detail below. The bottom panel summarizes characteristics of off-platform programs. An important take-away from the table is that the distinction between on- and off-platform programs does not perfectly reflect differences in other characteristics that typically determine preferences for programs —public (typically free) and private programs can be found both in- and off-platform, and so are different fields of study, such as STEM, Econ./Law/Business, or Humanities. As consequence, many high-school graduates wishing to enter higher education participate, in a given year, both in the centralized procedure and submit applications off platform.

We now proceed to present two unique features of the centralized assignment procedure that are important for the rest of this paper.

A centralized procedure with multiple rounds. The first one, motivating the research question, is that it is a multi-round assignment mechanism. After receiving an assignment from the platform, students can choose between three options. The student can *accept* the assignment immediately and proceed with enrollment; the student can refuse the assignment and *drop out* of the platform, thereby returning her assignment to the vacancy pool; and

Table 2.2.1: Types of post-secondary programs in France

On-platform programs excluding high-school performance from their admission criteria	
Bachelor programs	Three-year programs offered by public universities Graduates can enter the labor force or further pursue a Master's degree Deliver the degree of *Bachelor*
On-platform programs including high-school performance in their admission criteria	
Technical programs	Two-year programs offered by (public) institutes of technology Train mid-level technical workers Graduates can further their education by enrolling in Bachelor programs or engineering- and business-school programs Deliver the degree *Diplôme universitaire de technologie*
Vocational programs	Two-year programs offered by high schools Typically prepare students to enter the labor force, graduates can also further their education by enrolling in Bachelor programs Deliver the degree *Brevet de technicien supérieur*
Prep-school programs	Two-to-three-year programs offered by high schools Prepare students for the competitive exams to enter *Grandes écoles* (highly selective and public and private institutions)
Other programs	Typically offered by private engineering and business schools
Off-platform programs[1]	
Bachelors and other programs offered by public universities	(1) Programs accredited by other agencies than the Ministry of higher education —e.g., all programs in nursing and paramedical fields (accr. by the Ministry of Health and Social Affairs); most arts programs (accr. by the Ministry of Culture) (2) Programs jointly offered by multiple institutions (i.e. delivering "double diplomas") (3) competitive programs (e.g. programs in the selective Paris Dauphine and several political science institutes)
Programs offered by private schools	24 of the 36 private business schools recruiting after high-school 9 of the 76 private engineering schools recruiting after high-school

Notes. This table defines the main types of post-secondary programs in France. Each individual program is further defined by its *field of study* or *major*, and its host institution. [1]Source: Cour des Comptes (2017).

the student can decide to *delay* their decision and participate in the next round, that is, hold on to their assignment and maintain her applications to the programs ranked higher than her current offer in the her ROL.[4][5] The matching algorithm is then run again by the clearinghouse, excluding students who dropped out and using the same ROLs and priorities as in the first round, to assign the remaining vacancies to unmatched students and to those who chose to delay their decision. After receiving their second-round offer, students may again *accept*, *drop out*, or *delay*. In the third and final round, the clearinghouse produces a final set of offers, which students can only accept (or refuse and drop out). Figure 2.2.1 illustrates the timing of the process in 2015.

Figure 2.2.1: Multiple assignment rounds timeline in 2015

NOTES. This figure shows the timeline of the assignment procedure for 2015. In rounds one and two, students submit their decision to accept, drop out, or delay between the release of the offers and the end of the round.

ROLs, priorities, and the assignment algorithm. The other unique feature of the French APB procedure consists in the rules of the assignment mechanism itself. Regarding ROLs, rules are pretty straightforward: students can rank up to 36 programs, with a maximum of twelve per *type* of program. The four *types* of programs –Bachelor programs, technical programs, prep-school programs, and vocational programs— are described in Table 2.2.1. When it comes to priorities, rules are much more complex and obscure. From the perspective of students, no priority score is known at the time of application, nor at any point during or after the process. On the one hand, prep-school programs, technical programs, and vocational programs, typically do not disclose the criteria they use to determine priority, and are simply known to consider academic performance as a criterion. Bachelor programs, on the other hand, are not permitted by law to use academic performance as a determinant of applicants' priority. They rely on coarse priority groups defined based on the applicant's

[4]Note that students assigned to their top-ranked choice cannot *delay* and there is no program ranked higher than their current offer in their ROL.

[5]In practice there was also a fourth option where students reject the offer and still participate in the next round. As this is difficult to rationalize and only a small number of students used this option we exclude it from the analysis.

residential location (specifically, whether or not the student lives in the region the program is located in[6]) and the absolute rank of the program in the applicant's ROL. Then, a random lottery number is used to rank applicants within coarse priority groups. It is important to know that, in 2015, while the APB system was in use, the criteria used to determine Bachelor programs' priorities were unknown to the public. The use of lottery numbers and, more importantly, the role of their own ROL in the determination of their priority to programs was unknown to students, who were simply aware of the within-region advantage. Given the students' ROLs and the priorities set by each program, matches are determined by the clearinghouse using a college-proposing deferred acceptance (DA) algorithm. The limited amount of information provided to students at the time of application sets apart the French APB system from many other centralized school assignment mechanisms that have been studied in the literature. Not only priorities and their determinants but also the very algorithm used by the clearinghouse are unknown to students.[7] As a matter of fact, the only guideline provided to students was to report their preferences truthfully in their ROLs.

2.3 Data and Descriptive Evidence

The empirical analysis of this paper uses administrative data from French Ministry of Higher Education (Research and Innovation department, SIES). For every student participating in the college assignment procedure, the data show demographics, residential ZIP code, a high-school identifier, the track and major chosen in high school, performance at the national end-of-high-school exam, as well as the exact ROL submitted to the APB platform. The data also show the assignment offer received by the student in each round (if any) as well as her decision (accept, drop out, or delay) for that round. We supplement these data with publicly available information about the cost of housing in France's major cities. Details about the data sources and the definition of key variables are provided in Appendix 2.B.

Student sample. We focus our attention to the 2015 assignment procedure and further restrict the sample to students graduating from the *general* high-school track and partici-

[6] See Figure 2.B.1 for the map of the regions used for the construction of priorities (*académies*).

[7] Note also that "past admissions cutoffs" available to applicants in many other context are not disclosed in this setting.

pating in the college assignment procedure for the first time in 2015.[8] Descriptive statistics are shown in Table 2.A.1. Students submit on average slightly more than seven choices in their ROL. Of importance to motivate our research question, students' final assignment is on average ranked higher in their ROL than the first offer they receive (rank 1.95 vs. 2.19).

On the decision to *delay* acceptance. In this paragraph, we document the use by students of their option to delay in the different rounds of the assignment procedure. We show that there are gains from delaying, in the sense that a substantial fraction of the students who *delay* eventually receive a higher-ranked offer in a subsequent round. We also show that there is significant heterogeneity in the probability to choose *delay* across SES groups, potentially suggesting heterogeneity in the costs of waiting for one's final assignment.

First, Figure 2.A.1 illustrates the extent to which students choose to delay their decision. Focusing on students who did not receive an offer from their top-ranked choice (and are therefore eligible to use the delay option), it shows the share choosing to accept, delay, and drop out in each assignment round. In the first round, about 55% of these students decide to delay their decision, while slightly less than 40% decide not to wait for a better option and accept their offer. In the second round, almost 65% of participants choose to delay, while 20% accept their assignment. In the third round, when students cannot choose delay, 80% of participants accept their assignment while the rest drops out.

Table 2.3.1 illustrates the expected benefits of choosing the delay option. Twenty percent of the students choosing to delay in Round 1 receive a new (high-ranked) offer in Round 2, and 25% of the students delaying in Rounds 1 and 2 receive a better offer in Round 3.

Table 2.3.1: Share of students receiving new offers after choosing *delay*

	Share
Receiving better offer in Round 2, cond. on choosing *delay* in Round 1	0.20
Receiving better offer in Round 3, cond. on choosing *delay* in Rounds 1 and 2	0.25

NOTES. This table reports the share of student receiving a new (higher-ranked) offer conditional on choosing delay in the previous rounds.

Table 2.3.2 and Figure 2.3.1 illustrate the heterogeneity in the choice of the delay option across demographics. Table 2.3.2 reports the share of students choosing the delay option in each round (among those who did not receive an offer from their top-ranked choice) broken

[8] We also exclude students who completed high school outside of mainland France, whether it is abroad or in the French overseas regions.

Figure 2.3.1: Average marginal effects of demographics on the decision to *delay*, by round

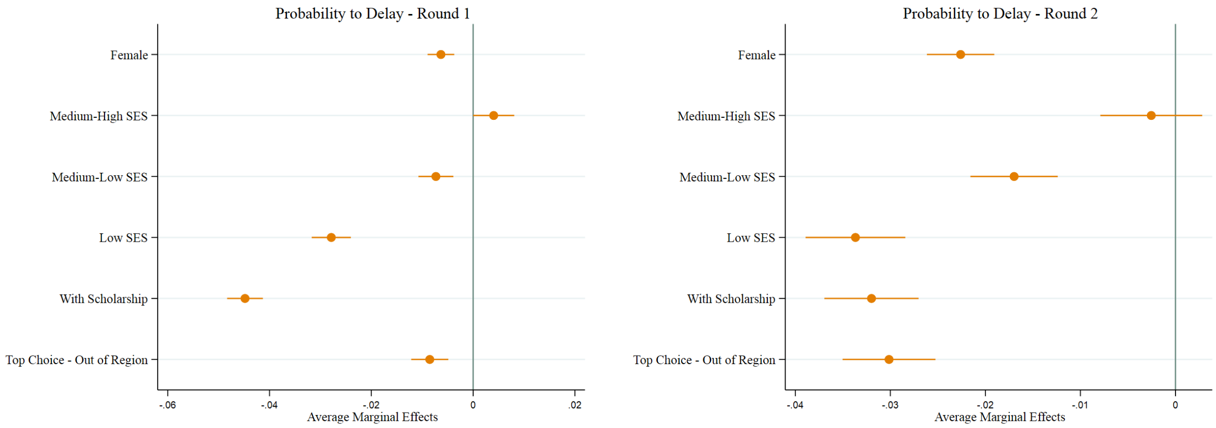

a) First Round **b) Second Round**

Notes. This figure shows the average marginal effects of demographics on the decision to *delay* estimated, for each round separately, using a logit regression.

down by SES. The table shows the existence of a gradient across SES in the use of the delay option. In particular, high-SES participants are eight (six) percentage points more likely to choose *delay* than low-SES participants in Round 1 (Round 2). To further explore this heterogeneity, we estimate a logit regression and report in Figure 2.3.1 the average marginal effects of several student characteristics on the probability to use *delay* in each round. While female students are less likely to choose to delay, indicators of lower income (namely, low SES and eligibility for means-based scholarship) remain the best predictors of the choice to delay. Applicants whose top choice is outside of their home region are also less likely to delay, consistent with housing costs being more salient in these cases.

Table 2.3.2: Share of eligible students choosing to *delay*, by SES

	Round 1	Round 2
High SES	0.57	0.67
Medium-high SES	0.56	0.66
Medium-low SES	0.54	0.64
Low SES	0.49	0.61

Notes. This figure shows the share of students who accept, delay, and drop out in each round (conditional on *not* having received an offer from their top-ranked choice), broken down by SES.

2.4 A Model of Students' Dynamic Application Decisions in a Sequential Admission Mechanism

We build a structural model of students' decisions in a sequential admission mechanism as the one implemented in France. It encompasses the different decision-making steps through which the student will go through on the platform. The model thus has two stages. In the first stage, students submit their ROL, based on their preferences between the different programs. The second stage of the model is dynamic: in a given round, students receive an offer and choose one of the three (or two in the last round) different answers available on the platform.

2.4.1 Stage 1: Rank-Ordered List Submission

When they connect to the platform, students are asked to submit a rank-ordered list of the programs they want to apply to. We assume students list truthfully their preferences, as recommended on the platform. We discuss this assumption at the end of this section.

Utility. A student i derives utility u_{ij} from enrolling in program $j \in \mathcal{J}$, where $\mathcal{J}$ is the set of programs available on the platform. We assume students rank programs according to the utility of being matched to them and a random shock:

$$u_{ij} + \eta_{ij} = u_j(S_i, \tau) + \eta_{ij}.$$

The utility is a j-specific function of a $L \times 1$ vector of observed characteristics S_i and a type τ. We assume a distribution of discrete types with finite support. The shock η_{ij} takes into account that they might have a 'trembling hand'.

Revealed ROL. Let $\bar{R}$ be the maximum length of submitted ROLs, and d_{ir}^{ROL} denote the program ranked in position $r \in \{1, ..., \bar{R}\}$ on student i's ROL. Assume an extreme value type 1 distribution on η_{ij} with scale parameter σ. Students revealing their preferences, the likelihood of observing one's ROL is given by the well-known exploded (or rank-ordered) logit probabilities:

$$\frac{\exp(u_{id^{ROL}_{i1}}/\sigma)}{\sum_{j\in\mathcal{J}}\exp(u_{ij}/\sigma)} \times \frac{\exp(u_{id^{ROL}_{i2}}/\sigma)}{\sum_{j\in\mathcal{J}\setminus\{d^{ROL}_{i1}\}}\exp(u_{ij}/\sigma)} \times ... \times \frac{\exp(u_{id_{iR_i}}/\sigma)}{\sum_{j\in\mathcal{J}\setminus\{d^{ROL}_{ik}\}^{R_i-1}_{k=1}}\exp(u_{ij}/\sigma)} \tag{2.1}$$

Note that many students do not submit a complete ranking. This could be the result of an effort cost in submitting a long list, against a limited gain if the list already contains safe choices. An alternative explanation is that the outside option (i.e. not attending a program in $\mathcal{J}$) is preferred over the unranked alternatives. We remain agnostic about the underlying reason by not including the outside option here and therefore only consider ranks between programs on the platform $\mathcal{J}$.

Remark. As described in Section 2.2, no information was disclosed to the public about either the algorithm used by the clearinghouse or the existence of strategic incentives, which stem from the fact that the applicant's priority for any Bachelor program is a function of the program's rank in her ROL. The only available guideline to applicants was to rank programs in order of preference. As a consequence, it is most reasonable to assume that students truthfully report their preferences in their ROL. This assumption in the setting of this paper is in line with the argument made in the literature that, even though the mechanism is not strategy-proof, truth-telling or a simple strategy close to truth-telling (truncation) may be the best action participants can take when they have limited information about the mechanism or others' preferences. Roth and Rothblum (1999) highlight that even in the simple case of the school-proposing DA, identifying profitable deviations from truthful reporting requires students to have more information than they often have in practice. Typically, a given manipulation is profitable under certain priorities and preference profiles for other applicants, but will yield a strictly worse outcome than truthful reporting given another state of the world . It is then crucial for a student to have good information about priorities and others' preferences —which students in our setting do not have— to be able to identify whether a particular action will be profitable. More recently, Troyan and Morrill (2020) formalize the concept of "obvious" profitable manipulations as profitable manipulations that can be identified as doing better than truthful reporting for agents who only know the range of outcomes that can result from each of their actions, that is, what would happen in the 'best-case scenario' and in the 'worst-case scenario' following their possible actions. They show that, in that sense, the school-proposing DA is not obviously manipulable.

2.4.2 Stage 2: Dynamic Model of Students' Enrollment Decisions

In the second stage of the model, students receive offers from the platform and can decide to match. Matching can occur in different rounds $t = \{1, 2, 3\}$. At the start of each round, each student receives a unique offer $j_t \in \mathcal{J} \cup \{0\}$, with $j_t = 0$ denoting no offer of a program. If $j_t \neq 0$ and $t < 3$, they can choose between options $k = \{1, 2, 3\}$ to maximize their expected lifetime utility. These denote respectively accepting the offer, delaying the decision, or dropping out from the platform. In the last round ($t = 3$), applicants can only decide to accept ($k = 1$) or drop-out from the platform ($k = 3$). If $j_t = 0$, only $k = 2$ and $k = 3$ are available. The expected lifetime utility is formalized by choosing the option d_{it}^{DDC} that yields the highest value of $v_{ikt} + \epsilon_{ikt}$, with v_{ikt} a conditional value function, and ϵ_{ikt} a random shock, revealed to students in round t. We now discuss the conditional value functions of each option.

Accept ($k = 1$). Accepting the offer yields the utility of being admitted to program j_t, potentially augmented by an early acceptance advantage ($e_{j_t,t}(S_i, \tau)$):

$$v_{i1t} = u_{j_t}(S_i, \tau) + e_{j_t}(S_i, \tau) \times I(t < 3) \tag{2.2}$$

with $I()$ the indicator function. The early acceptance advantage allows for deviations from the utility at enrollment to take into account that some students might prefer to know the characteristics of their program sufficiently in advance. More specifically, we expect housing opportunities to differ and will therefore estimate how distance and local rents affect this advantage.

Delay ($k = 2$). A student can also decide to wait for better options, while not losing the currently assigned alternative. In this case, they incur a waiting cost in the current period: $\omega_{it} = \omega_t(S_i, \tau)$. In contrast to the early acceptance advantage, this captures the impatience of students that is unrelated to the characteristics of the offer. This baseline cost of waiting can be financial. Low-SES households in particular could benefit from knowing that the student will enroll to prepare for the cost of studying. This could also reflect psychic costs, which are high if the student is impatient, low (and even negative) if they procrastinate (Akerlof, 1991). In addition to waiting costs, students also receive a continuation value that captures their (weakly) improved offer in the next round. The conditional value function is then given by:

$$v_{i2t} = -\omega_t(S_i, \tau) + \sum_{j' \in \mathcal{R}_i^{j_t} \cup \{j_t\}} \Pr(j_{t+1} = j' | \Omega_{it}) \bar{V}_{it+1}(\Omega_{it}, j_{t+1}) \tag{2.3}$$

where $\Pr(j_{t+1} = j'|\Omega_{it})$ denotes the probability to receive an offer from program j' in the next round, $t+1$, conditional on their information Ω_{it}. $\mathcal{R}_i^{j_t}$ is the set of options in $\mathcal{R}_i$ that are ranked above j_t. $\bar{V}_{it+1}(\Omega_{it}, j_{t+1})$ is the expected value of behaving optimally in the next round, conditionally on their current information and the offer received, j_{t+1}. We assume students keep track of their time-invariant individual characteristics S_i, type τ, ROL $\mathcal{R}_i$ and the time-varying round t and offer j_t:

$$\Omega_{it} = (S_i, \tau, \mathcal{R}_i, t, j_t). \tag{2.4}$$

Drop out ($k = 3$). Finally, the student can decide to drop out from the platform:

$$v_{i3t} = u_{0t}(S_i, \tau) \tag{2.5}$$

where $u_{0t}(S_i, \tau)$ captures the utility of their outside option.

Choice sets and solving the model. Assuming mean-zero extreme value type 1 taste shocks with scale normalized to 1, we can derive the choice probabilities in each round:

$$\Pr(d_{it}^{DDC} = k|\Omega_{it}) = \frac{\exp(v_{ikt})}{\exp(v_{i1t}) + \exp(v_{i2t}) + \exp(v_{i3t})} \tag{2.6}$$

The choice problem stops when students choose one of the terminal actions ($k = 1$ or $k = 3$). Note that in $t = 3$, only these two options are available (i.e. $v_{i23} = -\infty$) so the model becomes essentially static and can be solved for each realization of the time-varying state variable j_3. The rest of the model can be solved by backward induction, taking into account the expected value of behaving optimally in the future. This is facilitated by the closed-form solution resulting from the extreme value distribution:

$$\bar{V}_{it+1}(\Omega_{it}, j_{t+1}) = \ln(\exp(v_{i1t+1}) + \exp(v_{i2t+1}) + \exp(v_{i3t+1})) + \gamma \tag{2.7}$$

where γ denotes Euler's constant. Substituting this in (2.3), allows us to write the conditional value functions in $t = 2$, up to the data, utility parameters and state transitions. We can proceed in an analog way for $t = 1$.

Note that not every option is always available, i.e. some of the $v_{ikt} = -\infty$. Student i does not have the possibility to delay if they receive an offer to their top-ranked program. A student without an offer in the current round cannot choose to accept. As mentioned before, in the final period a terminal action needs to be chosen.

2.5 Identification and estimation

We first discuss the identification and parameterization of the model. We then discuss estimation for the case where type τ is known, how to estimate the model without solving it (CCP estimation), and how to allow τ to be unobserved by the econometrician.

2.5.1 Identification

We first discuss the case when types are observed and then generalize.

Program utility. The program utility $u_j(S_i, \tau)$ for $j \neq 0$ is identified up to scale using the ROLs, i.e. we identify $\tilde{u}_{j_3}(S_i, \tau) \equiv \frac{1}{\sigma} u_{j_3}(S_i, \tau)$. Without loss of generality, we normalize the utility of an arbitrary reference alternative to be 0. In practice, this corresponds to a hypothetical Bachelor-STEM program with other characteristics (such as distance) set to 0.

Scale and dropout in terminal period. Given the program utility, we can identify dropout utility in the last period ($t = 3$). To see this, note that in $t = 3$, there is no early acceptance advantage in v_{i1t} and there is no possibility to choose $k = 2$. Therefore, we obtain a static model with the utility of $k = 1$ known up to scale parameter σ, and dropout utility identified from the observed accept/dropout decision. We first identify the scale. Calculate choice probabilities in $t = 3$ and map them into utility functions:

$$\ln \Pr(d_{i3}^{DDC} = 1|\Omega_{i3}) - \ln \Pr(d_{i3}^{DDC} = 3|\Omega_{i3}) = v_{i13} - v_{i33}$$

$$
\begin{aligned}
&= u_{j_3}(S_i, \tau) - u_{03}(S_i, \tau) \qquad (2.8)\\
&= \sigma \tilde{u}_{j_3}(S_i, \tau) - u_{03}(S_i, \tau).
\end{aligned}
$$

with Ω_{it} given by (2.4) and the conditional value functions by (2.2) and (2.5). Let z be a program characteristic with $\frac{\partial \tilde{u}_{j_3}(S_i,\tau)}{\partial z} \neq 0$. As z is excluded from u_{0t}, taking derivatives and re-arranging yields the scale parameter:

$$\sigma = \frac{\partial(\ln \Pr(d_{i3}^{DDC} = 1|\Omega_{i3}) - \ln \Pr(d_{i3}^{DDC} = 3|\Omega_{i3}))/\partial z}{\partial \tilde{u}_{j_3}(S_i, \tau)/\partial z}.$$

With probabilities observed and $\frac{\partial \tilde{u}_{j_3}(S_i,\tau)}{\partial z}$ identified from the ROL, we have identified the scale, and therefore also $u_{j_3}(S_i, \tau)$. We can use (2.8) again to identify the utility of drop out $u_{03}(S_i, \tau)$.

Other primitives in the dynamic model. While state transitions are identified non-parametrically, utilities in dynamic discrete choice models are identified only after assuming a distribution on the taste shocks, specifying the discount factor, and normalizing the utility of an alternative in each state (Magnac and Thesmar (2002b)).

We assume a mean-zero extreme value type 1 distribution on the taste shocks and we normalize the discount factor to 1 (as different stages follow each other quickly). If the early acceptance advantage $e_{j_t}(S_i, \tau)$ was known, we would be in the standard case, with the utility of $k = 1$ known in every state. To identify the early acceptance advantage, we again exploit the program characteristics of the assigned alternative j_t. While student characteristics drive waiting costs and dropout utility, we set their baseline effect on the early acceptance advantage to be 0. Baseline effects of program characteristics of j_t, as well as their interactions with student characteristics, are identified as they do not enter waiting costs and dropout utility. Conceptually, waiting costs capture the disutility students derive from not knowing their final outcome yet, given that they are currently assigned the benchmark program, while the early acceptance advantage captures how program characteristics can change their willingness to wait.

Unobserved types. Both the ROL and the panel data of the dynamic model identify the unobserved types as only types are able to capture unobserved heterogeneity that is correlated within a ROL and between periods. For example, if students that rank Bachelor programs in rank 1 also usually put a Bachelor program in rank 2, it indicates the existence of an unobserved type that prefers Bachelor programs. Students that repeatedly delay their decision, despite a good j_t, suggest a low waiting cost type.

2.5.2 Parameterization

To obtain efficient estimates from our sample, we impose the following parametric structure.

Program utility. We parameterize the utility function of a student i when being matched to a program $j \in \mathcal{J}$ as follows:

$$u_j(S_i, \tau) = X'_{ij}(\Lambda S_i + \lambda_\tau) \text{ for } j \neq 0 \tag{2.9}$$

X_{ij} is a $K \times 1$ vector of observed program characteristics. Λ is a $K \times L$ matrix of parameters capturing the impact of observed student characteristics on program characteristics.[9] λ_τ is a $K \times 1$ vector of parameters capturing the impact of an unobserved student type.

A first set of program characteristics consists in indicator variables for each program type and major. Since only differences in utility are identified, we estimate the difference with a Bachelor program, in a STEM major, by including indicator variables for the other types (prep-school programs, technical programs, vocational programs, and others) and majors (Econ/Law, Humanities, Services, and Production[10]). A second set of characteristics flexibly captures the location characteristics. We include measures relative to the location of the students such as distance from home and indicators for the same catchment area and region.[11] We also include proxies for local housing costs by distinguishing between programs in and outside cities. For programs in cities, we further distinguish by measures of local rents. Finally, as some prep-school programs offer dorms, for this type of programs, we include indicators for whether dorms are available.

The vector of students' characteristics captures the gender of the student, her SES status, her grade at the centralized exam and indicators for each track in high school (Sciences, Humanities, Social Sciences). Type τ captures heterogeneity in preferences that is not explained

[9]We observe K program characteristics and $L-1$ student characteristics because we include a dummy in S_i to capture the baseline effect of program characteristics.

[10]The service and production majors are common among technical and vocational programs.

[11]Regions correspond to the geographic administrative units called *Département* in France. There are 96 such regions in metropolitan France.

by these characteristics.

Early acceptance advantage. The early acceptance advantage is parametrized in a similar way:

$$e_{j_t}(S_i, \tau) = X'_{ij}(\Phi S_i + \phi_\tau) \tag{2.10}$$

with Φ and ϕ_τ capturing the impact on the early acceptance advantage of respectively observed and unobserved student characteristics. To keep the model estimation tractable, we include only the program characteristics that capture mobility concerns. We include distance, a dummy for the program being located in a city, and (for cities) the local rent.

Waiting costs and utility of drop out. Finally, we specify heterogeneous waiting costs and outside options as follows:

$$\omega_t(S_i, \tau) = \psi_s S_i + \beta_\tau + \gamma_t \tag{2.11}$$

$$u_{0t}(S_i, \tau) = \alpha_s S_i + \kappa_\tau + \mu_t \text{ for } j = 0 \tag{2.12}$$

with the shifts of intercepts for the first type and period normalized to 0. Note also that $\omega_t(S_i, \tau)$ is only defined for $t < 3$.

State transitions. To solve the dynamic model, students take into account the probability to receive an offer $\Pr(j_{t+1} = j'|\Omega_{it})$ by using their information (2.4). We consider a parametric, but flexible relationship by estimating a logit model that predicts the probability to improve the offer $\Pr(j_{t+1} \neq j_t|\Omega_{it})$ and a conditional logit among the higher-ranked options $\Pr(j_{t+1} = j'|\Omega_{it}, j_{t+1} \neq j_t)$.

The binary logit includes the student's observed characteristics S_i and type τ and round t-specific intercepts. The program's type is also allowed to affect this, and we allow for heterogeneous effects by S_i and τ. Finally, we control for dorm availability and the rank of the current offer, and we add controls for the number of higher-ranked programs in the ROL of each type and each major. For the conditional logit, we first predict a program's selectivity.[12]

[12] To do this, we first estimate a logit on the full sample of applicants from 2015, where the dependent variable is equal to one if the student was admitted in the first round, and compute a program's selectivity index as the associated predicted latent variable. We restrict the estimation sample to the student-program

This index enters with a round-specific effect. Furthermore, we include a dummy for living in the catchment area of the program, the rank, major and type. We also allow the latter to depend on being inside or outside the catchment area.

2.5.3 Estimation with known types

With known types τ, we need to estimate the utility of being matched to a program $u_j(S_i, \tau)$, the early acceptance advantage $e_j(S_i, \tau)$, the outside option $u_{0t}(S_i, \tau)$ and the waiting costs $\omega_t(S_i, \tau)$. Moreover, we need to recover the state transitions $\Pr(j_{t+1} = j'|\Omega_{it})$ from the data.

Let θ capture all parameters to estimate. We can then write a likelihood function for the model where each individual's likelihood contribution is given by:

$$L_i(\theta) = L_{1i}^{ROL}(\theta_1) \prod_{t=1}^{3} (L_{it}^{TRANS}(\theta_2) L_{it}^{DDC}(\theta_1, \theta_2, \theta_3))$$

with $\theta = (\theta_1, \theta_2, \theta_3)$. $L_{1i}^{ROL}(\theta_1)$ is given by the probability of the observed ROL (2.1) and captures the utility parameters of each program up to scale, i.e. Λ/σ and λ_τ/σ. θ_2 are the parameters governing the state transitions. Finally, $L_{it}^{DDC}(\theta_1, \theta_2, \theta_3)$ is given by the choice probabilities in each round (2.6) with θ_3 the remaining parameters. These are the scale of the first stage trembling hand error (relative to the scale of utility) σ, the early acceptance parameters in Φ and ϕ, the waiting cost parameters in ψ and the dropout parameters α.

Note that the loglikelihood function is additively separable with loglikelihood contributions:

$$l_i(\theta) = l_{1i}^{ROL}(\theta_1) + \sum_{t=1}^{3} (l_{it}^{TRANS}(\theta_2) + l_{it}^{DDC}(\theta_1, \theta_2, \theta_3)).$$

Therefore, we can obtain consistent estimates by sequential estimation. We first obtain the estimates of θ_1 from a rank-ordered logit model on the ROLs.[13] We also obtain θ_2 from the (conditional) logit models that predict state transitions. We then use the estimated values of θ_1 and θ_2 to estimate the remaining parameters θ_3 in the dynamic choice model.

---

pairs where the student was an actual applicant, i.e. programs ranked weakly above the one from which the student received an offer in the first round. For the production and services majors, we allow for different effects in technical and vocational programs. The Econ./Law, STEM and Humanities majors are allowed to have different effects in prep-school, Bachelor and other programs.

[13]Since the choice set of students is very large, we exploit the properties of the logit probabilities and estimate the model by random sampling from the choice set, as explained by (Train, 2009a, page 65). In practice, we sample 450 alternatives out of the total set of 10,150.

CCP estimation

To estimate θ_3, we need to solve the model of stage 2 for each value of the state. This is cumbersome as the state space consists of every option that student i included in its ROL.

Dynamics enter through the choices allowing a student to stay on the platform in the following round ($k = 2$). We deal with the expected value of behaving optimally in the next round by rewriting the conditional value function of $k = 2$ as suggested by Hotz and Miller (1993) and Arcidiacono and Miller (2011). In particular, we can rewrite the ex ante value function (2.7) as a function of the conditional value function of dropout ($k = 3$), a terminal action that is always available, and the Conditional Choice Probability (CCP) of that action:

$$\bar{V}_{it+1}(\Omega_{it}, j_{t+1}) = u_{0t+1}(S_i, \tau) - \ln \Pr(d^{DDC}_{it+1} = 3|\Omega_{it}, j_{t+1})$$

No future value terms appear on the right-hand side because, with a terminal action, the finite dependence property trivially holds (Arcidiacono and Miller (2011)).

We can use this to re-write the conditional value function of $k = 2$:

$$\begin{aligned} v_{i2t} &= -\omega_t(S_i, \tau) + \sum_{j' \in \mathcal{R}_i^{j_t} \cup \{j_t\}} \Pr(j_{t+1} = j'|\Omega_{it}) \left(u_{0t+1}(S_i, \tau) - \ln \Pr(d^{DDC}_{it+1} = 3|\Omega_{it}, j_{t+1})\right) \\ &= -\omega_t(S_i, \tau) + u_{0t+1}(S_i, \tau) - \sum_{j' \in \mathcal{R}_i^{j_t} \cup \{j_t\}} \Pr(j_{t+1} = j'|\Omega_{it}) \ln \Pr(d^{DDC}_{it+1} = 3|\Omega_{it}, j_{t+1}) \end{aligned}$$

where the last line follows from $\sum_{j' \in \mathcal{R}_i^{j_t} \cup \{j_t\}} \Pr(j_{t+1} = j'|\Omega_{it}) = 1$.

Using these conditional value functions, we no longer need to solve the model to estimate the parameters if we estimate CCPs (i.e. $\Pr(d^{DDC}_{it+1} = 4|\Omega_{it}, j_{t+1})$ for $t + 1 = 2, 3$) in a first step and the estimator essentially becomes similar to the estimator of static model with a correction term.[14]

[14] We predict CCPs using a flexible binary logit. We include the observed and unobserved characteristics of the student, that we interact with characteristics of the current offer. We also include the numbers of programs ranked higher of different types and majors. For the characteristics of the current offer, we take into account the same variables that enter the utility function, but also the current rank and selectivity index.

2.5.4 Estimation with unknown types

Let there be M types τ with probability to occur π_τ. We then obtain the following log-likelihood contributions:

$$l_i(\theta,\pi) = \ln \sum_{\tau=1}^{M} \pi_\tau [l_{1i}^{ROL}(\theta_1) \prod_{t=1}^{3} (l_{it}^{TRANS}(\theta_2) l_{it}^{DDC}(\theta_1,\theta_2,\theta_3))]$$

with an additional vector of parameters to estimate: $\pi = (\pi_2, ..., \pi_M)$ and $\pi_1 = 1 - \sum_{\tau=2}^{M} \pi_\tau$. Note that the log-likelihood function is no longer additively separable, preventing us from estimating the model in stages. Moreover, our estimation approach requires CCPs to depend on types. We therefore follow the estimation approach of Arcidiacono and Miller (2011). This is an adaptation of the EM algorithm and can be summarized as follows. We first specify the number of types M and provide a set of starting values for θ. We use this to calculate the probability of each observation i to belong to each type τ and the type probabilities π. We then estimate the CCPs and the model as if types are known, using the individual probabilities as weights. The new set of estimates for θ are used to update the types. This is repeated until convergence of the likelihood function.

2.6 Estimation Results and Counterfactual Simulations

In this section, we report all utility estimates.[15] The estimates of state transitions and the CCP are available upon request.

2.6.1 Program utility

The results on the utility of each program are summarized in Table 2.A.2, Table 2.A.3 and Table 2.A.4. To interpret them well, it is important to note that we report the estimates Λ/σ and λ_τ/σ. To compare them to other utility estimates (waiting costs and drop out), they should be multiplied by the estimate of the scale parameter σ: 0.125 (with standard error of 0.005). This estimate implies that the standard deviation of the "trembling hand" error in the ROL is quite small as it amounts to about 12% of a standard deviation in the taste shocks that enter the dynamic choice model. The utility parametrization involves several indicator

[15]Standard errors do not correct for the uncertainty in estimating unobserved types and the use of predicted values for CCPs and state transitions.

variables; *Bachelor program* and *STEM* serve as benchmark categories for programs type and field of study. To interpret the role of individual characteristics, the benchmark is a male, high-SES, not eligible to a scholarship, who graduated high-school with a sciences major and highest honors, and with unobservable type 1.

The first part of Table 2.A.2 shows the utility of different program types and majors for the benchmark group, as well as the variables related to proximity and housing. From the latter, we can conclude that students value proximity. This is reflected in the negative coefficient on distance, as well as in the nonlinear effects coming from staying in the region and catchment area. The distance parameter also helps to make sense of the magnitude of other utility estimates. Note that distance is scaled in 100 km and can be used to quantify the utility impact of other variables. For example, prep-school programs are, on average, valued less than Bachelor programs by the benchmark group, equivalent to an increase in distance of 107 km ($= 100 \times (0.954/0.888) = 107$). To capture differences in local amenities housing conditions, we allow for different preferences to study in the main cities of France. Furthermore, we let the impact differ by local rents. This should not be interpreted as a causal effect of rents, but rather as a proxy for local amenities. Indeed, we find that at 0 rent, students prefer not to live in a city, but they prefer locations with better amenities.

The rest of Table 2.A.2, as well as Tables 2.A.3 and 2.A.4 describe the heterogeneity in program preferences, i.e. how utilities differ for the students that outside of the benchmark group. The main findings the following. As compared to the benchmark group, lower SES students and students with lower end-of-high-school performance are more likely to consider Vocational programs as an alternative to Bachelor programs, and less likely to consider Technical and Prep-school programs. All other things equal, female students have a relatively lower preference from STEM field that any for other field. All other things equal, students graduating with a humanities or social sciences high-school major have a relatively stronger preference for studying humanities, econ./law, or services in college as compared to STEM.

Table 2.A.4 concludes with the heterogeneity that is not explained by observable characteristics. Forty-seven percent of the sample is estimated to be of unobserved type 2. This type is capturing a much higher willingness to consider alternatives to a Bachelor program. For example, the difference in utility derived from attending a prep-school program instead of a Bachelor is equivalent to attending a program 562 km closer to home, all other things constant. This type is also much more likely to consider a program in Humanities or Econ./Law.

The magnitude of these estimates show the importance of allowing for unobserved heterogeneity in preferences for post-secondary programs.

2.6.2 Early acceptance

We now discuss the parameters entering the second stage of the model, starting with the parameters of the early acceptance advantage (Φ and ϕ_τ). Estimates are shown in Table 2.A.5.

Because of housing considerations, accepting a program early in the process can have substantial benefits for students. Table 2.A.5 reveals that students are indeed more likely to accept a program early if it is further away. For the benchmark student, accepting a program that is 100 km away in an early round, rather than the last one, increases their utility by 0.125 utils. To interpret this number, note that this is about equivalent to the disutility to travel 100 km for a program, which is estimated to be -0.111 ($= -0.888 \times 0.125$). This suggests that early knowledge on where to enroll can entirely offset the travel cost for a large group of students (abstracting from the nonlinear effects of living in the region where the program is located). As student demographic and academic background might affect the set of available (affordable) housing options, we allow for heterogeneity in early acceptance utility. However, we find that the differences along these dimensions are typically small and without a clear pattern. Another variable that captures housing considerations is the proxy for local amenities (the *Main City* indicator and the interaction with local rents). While the effects for rents are small and insignificant, students are more likely to accept early a program located in a city.

2.6.3 Waiting costs and drop out utility

Estimates for the waiting costs (ψ) and drop out utility (α) are shown in Table 2.A.6. Note that the tables report the negative of the waiting costs parameters (-ψ) so they should be interpreted as utility, rather than disutility.

While students on average experience a cost from using the drop out option, high SES, being female, and graduating high school with the humanities or social sciences major and a lower performance are individual characteristics associated with a relatively lower cost from dropping out. On the contrary, type-1 students on average derive a positively value from using the delay option (potentially capturing procrastination), while this value is much lower for type-2 students. The value of the delay option is also notably higher for highest-honors students.

To better interpret the magnitude of waiting costs, it is useful to compare waiting costs to the valuation of different program characteristics. To do this comparison, we convert waiting costs to express them in distance-equivalent terms (in kilometers, $= \psi/(0.00888 \times 0.125)$). We also express the utility gains from enrolling in programs of different types (relative to a benchmark Bachelor program) and in different fields of study (relative to the STEM benchmark) in distance-equivalent terms. Figure 2.A.2 shows the predicted distribution of waiting costs (in distance-equivalent terms) in comparison to the relative valuation of programs types (top panel) and fields of study (bottom panel). Focusing first on the distribution of waiting costs themselves, we find that waiting costs are positive (i.e. waiting utility is negative) for a majority of students. The median waiting cost corresponds to 141 kilometers, implying that many students would accept an offer today rather than waiting for an offer that is not substantially closer. Turning to the comparison of waiting costs with the relative valuation of program characteristics, we find that, overall, the magnitude of waiting costs offset the utility gains from being assigned to a program in more preferred type or field of study. In other words, for many students, differences in utility between majors or types are not large enough to overcome the cost of delaying acceptance.

Finally, Figure 2.6.1 further illustrates the heterogeneity in estimated waiting costs by showing their distribution for the highest and lowest SES groups separately. It shows that waiting costs tend to be higher for low SES students.

2.6.4 The benefit of a multi-round mechanism

Multi-round mechanisms have the flexibility to allow seats to open up after students decide to drop out. However, they can also affect the final match through the frictions it imposes. Substantial waiting costs can force students to accept a sub-optimal offer early. Certain characteristics (such as distance) can make this more likely because of housing considerations. At the same time, the trade-off between incurring waiting costs today, against a potentially better offer in the future incorporates the strength of agent preferences in the matching mechanism, potentially leading to increased efficiency. The overall welfare impact of running a sequential mechanism, as well as its distributional effects, is therefore an empirical question.

Figure 2.6.1: Waiting Costs by SES Group

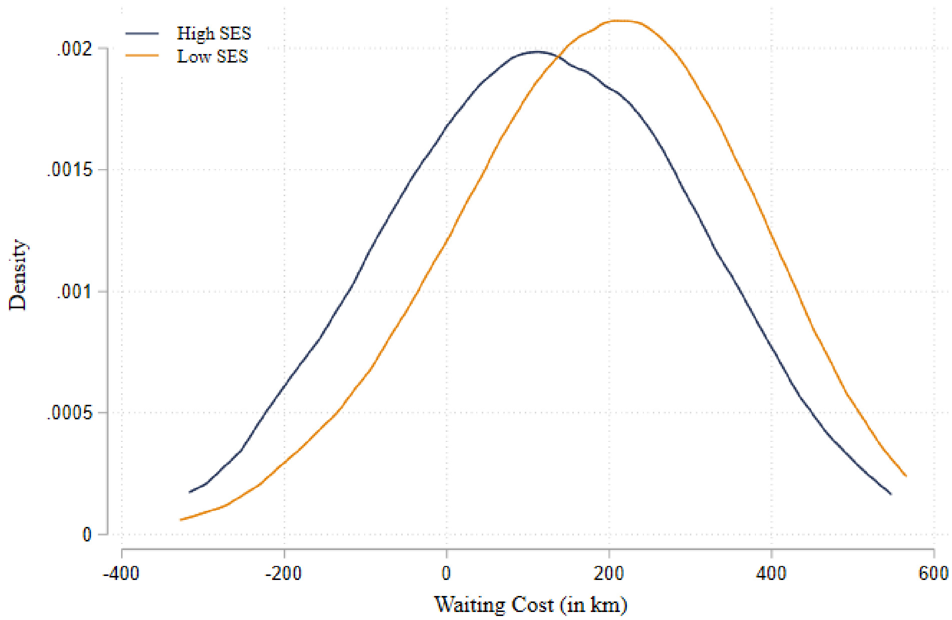

To answer this question, we simulate both the single-round and the multi-round mechanism and compare the effects on student welfare. Specifically, for the multi-round mechanism, we compute the value for a student of accepting immediately an offer, delaying, or dropping-out. In the counterfactual single-round scenario, we remove the option to delay at the end of the first round.[16] In that scenario, after receiving their first-round assignment, students can only choose to either accept their current offer or drop out of the platform. Students with no offer in the first round get the utility from dropping out from the platform; students who accept their offer benefit from the early acceptance advantage. Results are summarized in Table 2.6.1. On average, the multi-round procedure provides a gain equivalent to enrolling 299 kilometers closer to home relative to the single-round procedure. To better interpret this magnitude, it is useful to know that, on average, gaining admission to the top-ranked program instead of the second-ranked program corresponds to a gain of 79 kilometers. Gains from multiple rounds are heterogeneous across SES groups. Higher-SES groups welfare gains equivalent to a reduction in distance from home of 299 to 333 kilometers, against 259–275 for lower-SES groups. To better assess the effect of introducing multiple rounds on welfare

[16]Note that under the assumption that students report their preferences truthfully when submitting their ROLs, the ROL-submission stage is not affected by the presence or the absence of multiple assignment rounds and a future option to delay.

Table 2.6.1: Average welfare gains from a multi-round procedure relative to single-round

	Unconditional average		Conditional on not receiving a first-round offer from one's top-ranked program	
	Km-equiv. terms	Relative terms	Km-equiv. terms	Relative terms
	(1)	(2)	(3)	(4)
Overall	299	0.21	691	0.50
By SES				
High	333	0.25	698	0.53
Medium-high	299	0.20	713	0.50
Medium-low	275	0.19	673	0.50
Low	259	0.16	677	0.46

NOTES. Columns (1) and (3) measure average welfare gains in equivalent reduction is travel distance, all other things equal. Columns (2) and (4) present relative welfare gain, measured with respect to the welfare created by a single-round mechanism over the outside option — in distance-equivalent terms a reduction of 1,373 km (1,431 km) on average (and conditional on not receiving a first-round offer from one's top-ranked program), and for each of the four SES groups from highest to lowest: 1,345 km; 1,499 km; 1,428 km; 1,576 km (1,328 km; 1,435 km; 1,355 km; and 1,469 km, respectively).

inequalities, we renormalize utilities so that the outside option is equivalent to 0 utility and utility estimates capture the value of being assigned on the platform relative to the outside option. We find that the multi-round mechanism increases welfare by 56% more or highest SES group than for the lowest.

2.7 Conclusion

In this paper we investigate the consequences of using a sequential rather than a single-round centralized college admission mechanism. We focus on the context of the French APB admission mechanism, where students are allocated to higher education programs based on a sequential admission process with up to three offer rounds. We develop a parsimonious dynamic model of application and acceptance decisions, which captures a key dynamic trade-off between receiving a potentially better offer in the next round, and the disutility of waiting before enrolling in a program. We find that a multi-round system entails large gains for students of all socio-economic groups. Nevertheless, the gains are larger for the high SES students. While multi-round systems can improve outcomes, we also document the existence of large waiting costs, resulting in sub-optimal matches, particularly so for low-SES students. From a policy perspective, our findings point to the importance of reducing these cost by limiting the time needed to run the mechanism, and possibly also by facilitating access to local housing markets.

Appendices

2.A Additional Figures and Tables

Figure 2.A.1: Decision shares by round

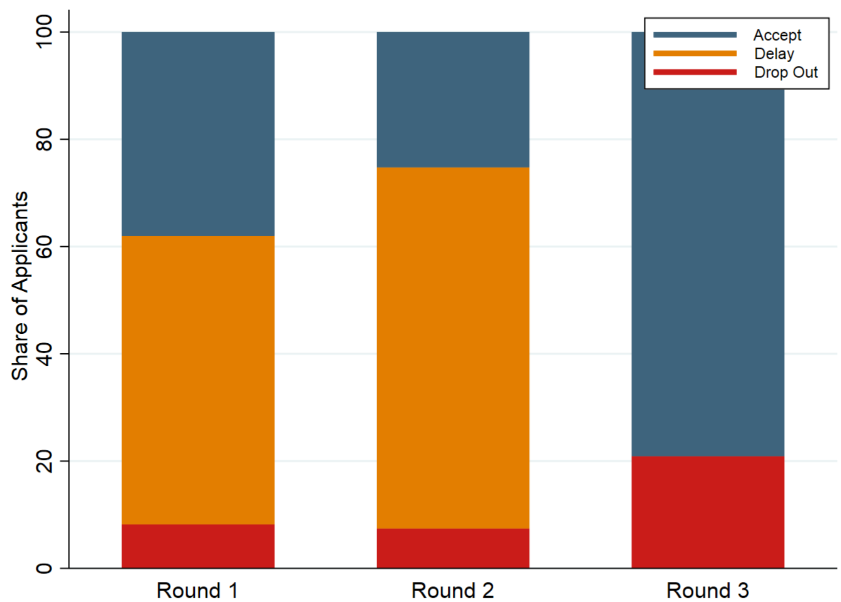

NOTES. This figure shows the share of students who accept, delay, and drop out in each round, conditional on *not* having received an offer from their top-ranked choice.

Table 2.A.1: Sample description

Demographics	
Female	0.57
High SES	0.40
Medium-high SES	0.16
Medium-low SES	0.27
Low SES	0.17
Means-based scholarship recipient	0.13
High-school major	
Sciences major	0.54
Social Sciences major	0.31
Humanities major	0.15
End-of-high-school-exam performance	
Highest honors (> 16/20)	0.12
High honors (14 to 16/20)	0.19
Honors (12-14/20)	0.28
Pass (10-12/20)	0.40
Applications and offers	
Average number of programs listed	7.09
Average rank of Round-1 offer	2.19
Average rank of accepted offer	1.95
Final match characteristics	
Bachelor program	0.52
Prep-school program	0.12
Technical program	0.10
Vocational program	0.05
STEM major	0.33
Econ./Law major	0.16
Humanities major	0.20
Other major	0.17
Urban location (Main city)	0.47
Avg. monthly rent (cond. on urban location, EUR)	448
Observations[1]	67,425

NOTES. This table shows descriptive statistics on the student sample and their final assignment. Unless otherwise specified, numbers provided are sample shares. See Appendix 2.B for details about the definition of SES categories, end-of-high-school-exam performance, and the construction of average rent. Statistics are shown for a random 25% sample of the population of interest (2015 graduates from the general high-school track participating in the college assignment procedure for the first time in 2015). This random sample is used for the structural estimation.

Table 2.A.2: Program utility (part 1 of 3)

Common component		
Program type (benchmark = Bachelor prog.)		
Prep-sch. prog.	-0.954	(0.057)
Technical prog.	1.717	(0.168)
Vocational prog.	-1.981	(0.203)
Other prog.	1.090	(0.054)
Program major (benchmark = STEM)		
Econ./Law	-1.014	(0.044)
Humanities	-3.419	(0.054)
Production	-2.942	(0.183)
Services	-4.279	(0.144)
Same catchment area	1.281	(0.014)
Same region	0.783	(0.013)
Main City	-0.489	(0.022)
Main City × Rent	0.205	(0.004)
Dorm Available	0.260	(0.031)
Dorm Available × Applied to Dorm	-0.397	(0.021)
Distance	-0.888	(0.007)
Heterogeneity		
SES Group (benchmark = High SES)		
Medium-High × Prep-sch. prog.	-0.260	(0.042)
Medium-High × Technical prog.	-0.233	(0.094)
Medium-High × Vocational prog.	0.044	(0.097)
Medium-High × Other prog.	-0.341	(0.048)
Medium-Low × Prep-sch. prog.	-0.418	(0.039)
Medium-Low × Technical prog.	-0.768	(0.083)
Medium-Low × Vocational prog.	-0.338	(0.085)
Medium-Low × Other prog.	-0.635	(0.046)
Low × Prep-sch. prog.	-0.668	(0.049)
Low × Technical prog.	-1.083	(0.101)
Low × Vocational prog.	-0.496	(0.101)
Low × Other prog.	-0.891	(0.059)
Medium-High × Econ./Law	-0.375	(0.042)
Medium-High × Humanities	0.036	(0.043)
Medium-High × Production	0.302	(0.102)
Medium-High × Services	0.185	(0.088)
Medium-Low × Econ./Law	-0.272	(0.037)
Medium-Low × Humanities	-0.032	(0.038)
Medium-Low × Production	0.958	(0.090)
Medium-Low × Services	0.692	(0.077)
Low × Econ./Law	-0.469	(0.045)
Low × Humanities	-0.258	(0.047)
Low × Production	1.045	(0.108)
Low × Services	0.689	(0.094)

Notes. Standard errors in parentheses.

Table 2.A.3: Program utility (part 2 of 3)

Heterogeneity (continued)		
Grade (benchmark = Highest honors)		
High honors × Prep-sch. prog.	-2.402	(0.054)
High honors × Technical prog.	-0.192	(0.182)
High honors × Vocational prog.	0.085	(0.216)
High honors × Other	-0.728	(0.056)
Honors × Prep-sch. prog.	-4.091	(0.057)
Honors × Technical prog.	-0.721	(0.175)
Honors × Vocational prog.	-0.272	(0.207)
Honors × Other prog.	-1.495	(0.060)
Pass × Prep-sch. prog.	-5.487	(0.061)
Pass × Technical prog.	-1.849	(0.173)
Pass × Vocational prog.	-0.637	(0.204)
Pass × Other prog.	-2.257	(0.063)
High honors × Econ./Law	0.187	(0.045)
High honors × Humanities	-0.380	(0.048)
High honors × Production	1.170	(0.201)
High honors × Services	1.709	(0.159)
Honors × Econ./Law	0.202	(0.045)
Honors × Humanities	-0.857	(0.048)
Honors × Production	3.250	(0.192)
Honors × Services	3.776	(0.153)
Pass × Econ./Law	0.030	(0.045)
Pass × Humanities	-0.936	(0.048)
Pass × Production	4.446	(0.189)
Pass × Services	5.020	(0.151)
High School Program (benchmark = Sciences)		
Social Sciences × Prep-sch. prog.	-0.941	(0.036)
Social Sciences × Technical prog.	0.189	(0.079)
Social Sciences × Vocational prog.	-0.013	(0.080)
Social Sciences × Other	-0.981	(0.055)
Humanities × Prep-sch. prog.	-0.107	(0.051)
Humanities × Technical prog.	0.226	(0.140)
Humanities × Vocational prog.	0.878	(0.134)
Humanities × Other prog.	-0.482	(0.059)
Social Sciences × Econ./Law	3.693	(0.039)
Social Sciences × Humanities	3.101	(0.041)
Social Sciences × Production	-0.021	(0.102)
Social Sciences × Services	3.809	(0.079)
Humanities × Econ./Law	3.582	(0.077)
Humanities × Humanities	5.535	(0.075)
Humanities × Production	-1.771	(0.246)
Humanities × Services	1.833	(0.147)

Notes. Standard errors in parentheses.

Table 2.A.4: Program utility (part 3 of 3)

Heterogeneity (continued)		
Scholarship Status (benchmark = Without Scholarship)		
Scholarship recipient × Prep-sch. prog.	0.188	(0.050)
Scholarship recipient × Technical prog.	0.033	(0.104)
Scholarship recipient × Vocational prog.	0.011	(0.103)
Scholarship recipient × Other prog.	0.026	(0.061)
Scholarship recipient × Econ./Law	0.072	(0.046)
Scholarship recipient × Humanities	-0.074	(0.048)
Scholarship recipient × Production	-0.245	(0.111)
Scholarship recipient × Services	-0.329	(0.098)
Sex (benchmark = Male)		
Female × Prep-sch. prog.	-0.419	(0.030)
Female × Technical prog.	-1.219	(0.066)
Female × Vocational prog.	-0.848	(0.067)
Female × Other prog.	-0.759	(0.035)
Female × Econ./Law	0.402	(0.029)
Female × Humanities	0.902	(0.030)
Female × Production	0.204	(0.072)
Female × Services	1.248	(0.061)
Unobserved Type (benchmark = Type 1)		
Type 2 × Prep-sch. prog.	4.994	(0.063)
Type 2 × Technical prog.	0.861	(0.078)
Type 2 × Vocational prog.	0.474	(0.081)
Type 2 × Other prog.	1.669	(0.045)
Type 2 × Econ./Law	-0.294	(0.032)
Type 2 × Humanities	1.786	(0.041)
Type 2 × Production	-3.558	(0.085)
Type 2 × Services	-4.217	(0.073)

NOTES. Standard errors in parentheses.

Table 2.A.5: Utility from Early Acceptance

Distance		
Distance	0.125	(0.022)
Distance × Female	0.013	(0.016)
Distance × Type 2	-0.027	(0.016)
Distance × Medium-High SES	-0.042	(0.022)
Distance × Medium-Low SES	-0.012	(0.019)
Distance × Low SES	-0.054	(0.026)
Distance × Scholarship recipient	-0.032	(0.025)
Distance × High honors	0.014	(0.022)
Distance × Honors	0.006	(0.022)
Distance × Pass	0.050	(0.022)
Distance × Social Sciences	-0.042	(0.018)
Distance × Humanities	-0.026	(0.022)
Main City		
Main City	0.125	(0.022)
Main City × Female	-0.083	(0.087)
Main City × Type 2	0.000	(0.087)
Main City × Medium-High SES	0.072	(0.129)
Main City × Medium-Low SES	-0.025	(0.109)
Main City × Low SES	-0.145	(0.134)
Main City × Scholarship recipient	0.054	(0.138)
Main City × High honors	-0.071	(0.148)
Main City × Honors	-0.325	(0.141)
Main City × Pass	-0.378	(0.136)
Main City × Social sciences	0.136	(0.095)
Main City × Humanities	-0.191	(0.124)
Main City × Rent		
Main City × Rent	-0.037	(0.026)
Main City × Rent × Female	0.027	(0.017)
Main City × Rent × Type 2	0.005	(0.017)
Main City × Rent × Medium-High SES	-0.002	(0.026)
Main City × Rent × Medium-Low SES	-0.002	(0.021)
Main City × Rent × Low SES	0.016	(0.027)
Main City × Rent × Scholarship recipient	0.014	(0.027)
Main City × Rent × High honors	0.000	(0.027)
Main City × Rent × Honors	0.048	(0.026)
Main City × Rent × Pass	0.040	(0.025)
Main City × Rent × Social sciences	-0.034	(0.019)
Main City × Rent × Humanities	0.030	(0.023)

NOTES. Standard errors in parentheses.

Table 2.A.6: Utility from Waiting & Drop Out

Common component		
Delay	0.190	(0.064)
Drop Out × Round 1	-2.546	(0.061)
Drop Out × Round 2	-2.229	(0.063)
Drop Out × Round 3	-1.709	(0.061)
Heterogeneity		
SES (benchmark = High SES)		
Medium-High × Delay	0.013	(0.048)
Medium-Low × Delay	-0.021	(0.041)
Low × Delay	0.026	(0.051)
Medium-High × Drop out	-0.312	(0.046)
Medium-Low × Drop out	-0.285	(0.039)
Low × Drop out	-0.548	(0.049)
Scholarship Status (benchmark = Without Scholarship)		
Scholarship recipient × Delay	-0.048	(0.054)
Scholarship recipient × Drop out	-0.325	(0.052)
Sex (benchmark = Male)		
Female × Delay	0.047	(0.034)
Female × Drop out	0.294	(0.033)
End-of-high-school-exam performance (benchmark = Highest honors)		
High honors × Delay	-0.153	(0.064)
Honors × Delay	-0.267	(0.060)
Pass × Delay	-0.462	(0.058)
High honors × Drop out	-0.081	(0.063)
Honors × Drop out	0.158	(0.059)
Pass × Drop out	0.397	(0.057)
Unobserved Heterogeneity (benchmark = Type 1)		
Type 2 × Delay	-0.158	(0.033)
Type 2 × Drop out	0.026	(0.031)
High-school Major (benchmark = Sciences)		
Social Sciences × Delay	0.102	(0.036)
Humanities × Delay	-0.130	(0.050)
Social Sciences × Drop out	0.775	(0.037)
Humanities × Drop out	0.768	(0.051)

Notes. Standard errors in parentheses.

Figure 2.A.2: Distribution of the Supply of Academic Seats and High School Student Population

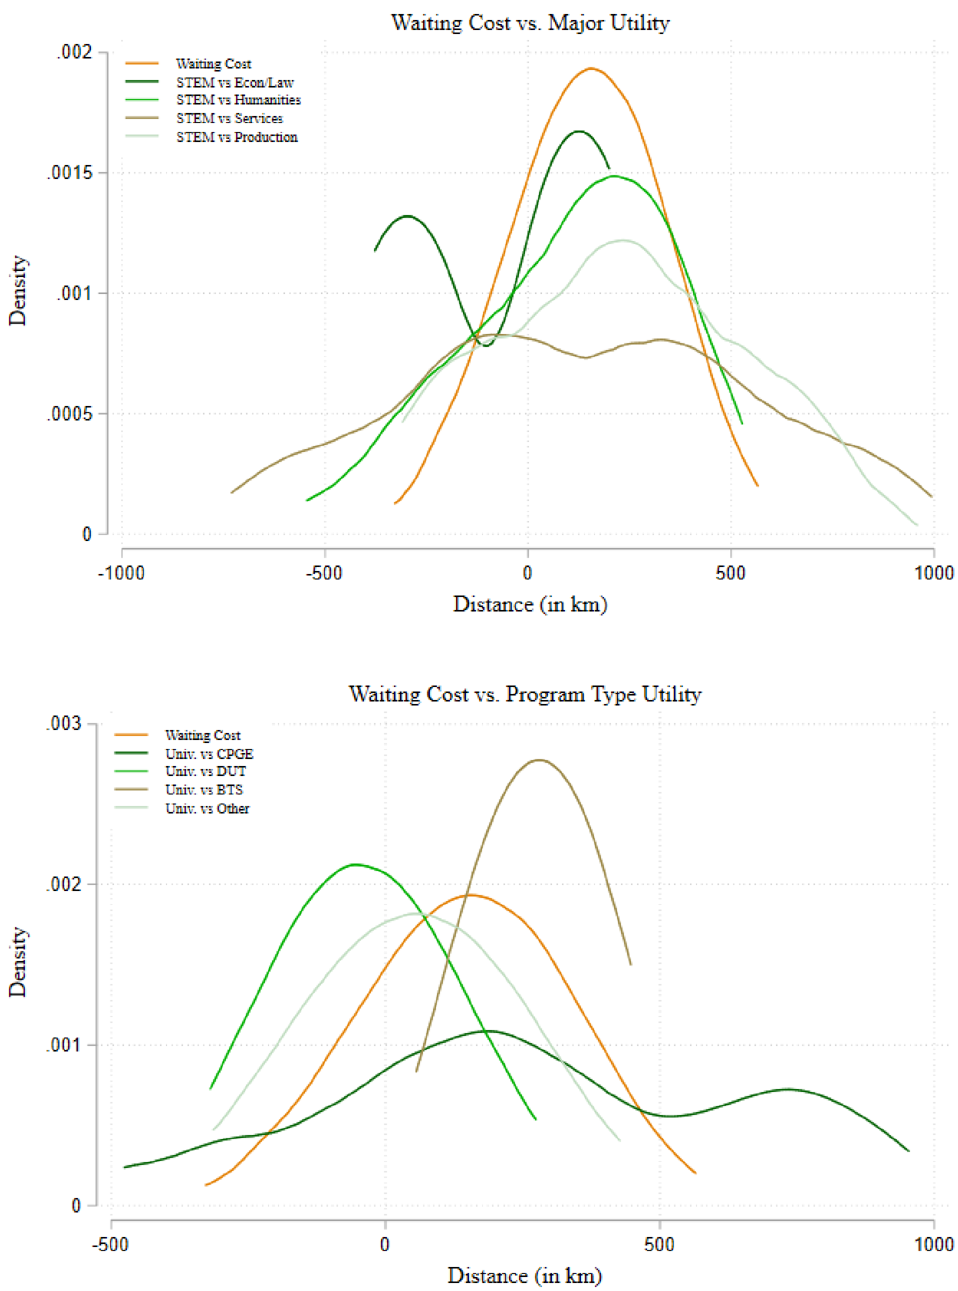

NOTES. Waiting costs and preferences for types of programs and fields of study

2.B Data appendix

Academic regions. Figure 2.B.1 shows the regions used by Bachelor programs to determine coarse priority groups.

Figure 2.B.1: Catchment Areas

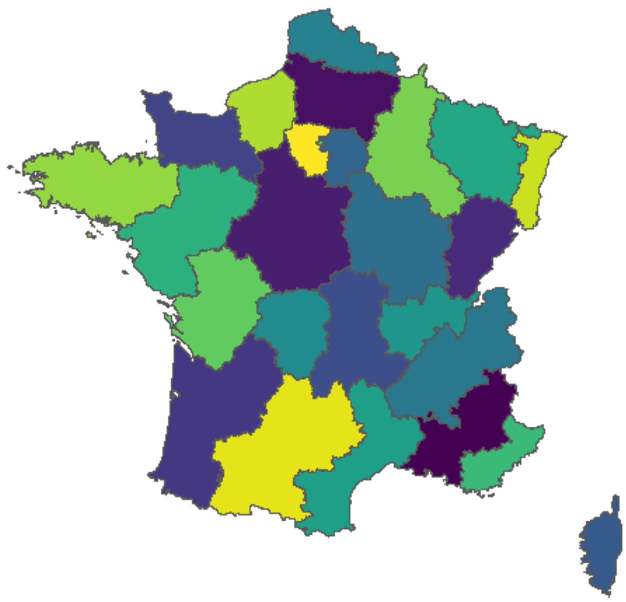

NOTES. This maps shows the geographic boundaries of the catchment areas in metropolitan France. Catchment areas outside metropolitan France are not shown.

SES. Following the classification used by the Ministry of Higher Education, we define SES based on the socio-professional category of the student's legal guardian, which consists in four categories in the data: High (company managers, executives, liberal professions, engineers, intellectual occupations, arts professions), Medium-High (technicians and associate professionals), Medium-Low (farmers, craft and trades workers, service and sales workers), Low (manual workers and unemployed individuals). Details on this classification can be found in Merle (2013).

End-of-high-school-exam performance. The end-of-high school exam is graded out of 20 points. A four-category discrete measure of this exam score is reported in our college-application data: highest honors (score above 16 out of 20); high honors (score between 14 and 16 out of 20)l honors (12–14), and pass (10–12). Note that 10 out of 20 is the minimum score required to pass the end-of-high school exam, graduate from high school, and access post-secondary education. All students in our sample are eligible to enter post-secondary education.

Cost of housing. Measures of local rents for 99 cities in France are taken from: https://www.century21.fr/pdf/logement_etudiant/2015/logement_etudiant_2015.pdf.

Chapter 3

Incomplete Information and the Complexity of Centralized College Application Processes

Abstract. Centralized application systems have been argued to have the potential to help close the socio-economic gap in access to higher education, by making information about the college application process transparent and easily accessible. Yet, exploiting data from the Chilean centralized admission mechanism, I show that 26% of the students submit an invalid application, suggesting that misinformation about the rules of the system is widespread. Taking this feature into account, I build a structural model of students' application decisions in order to study the impact of being misinformed on a student's admission outcome. I simulate the rank-ordered lists submitted by students in a counterfactual scenario where they would have been aware of the requirements for an application to be valid. I find that more than 80% of previously unassigned uninformed students would have been assigned under this scenario. More than half of these students belong to one of the first three deciles of the income distribution, indicating that the complexity of the application process disproportionately harms disadvantaged students.

3.1 Introduction

Despite evidence of large and positive returns to college, gaps in the access to higher education with respect to students' socio-economic background remain large (Chetty et al. (2018a)). Low income students' lack of information regarding the college-application process has been presented as being one of the key factors explaining this fact (Hoxby and Avery (2013), Hoxby and Turner (2015), Dillon and Smith (2017)), together with application fees. Centralized application systems have thus been argued to have the potential to remove these barriers as they are usually associated with the elimination of application costs, and greatly simplify the application process by offering a platform centralizing information.

Yet, centralized systems are difficult to design. In particular, they often come with a large set of rules and requirements determining whether a student's application is valid. This might make it difficult for students who are misinformed about the application process to navigate it, potentially undermining its functioning. It is thus important to understand whether there is room for improvements in the design of centralized application systems, in particular in terms of removing information barriers, which is key in allowing every student to submit the application which would maximize their welfare.

Through this paper, my goal is then to contribute to our understanding of the role of heterogeneity in the access to information in students' application behavior within centralized application systems, by answering two questions. First, do all students have the same knowledge of the rules of the college-application process? If not, what is the impact of such differences on students' enrollment outcomes?

These questions, that have not yet been answered in the literature, have a crucial policy relevance. Answering them would indeed help design fairer assignment mechanisms, for example by mitigating their complexity or by implementing policies assisting misinformed students in the application process.

Overall, this project brings three main contributions. I first contribute to the literature studying centralized college application systems by focusing on one particular determinant of students' application decisions that has not been studied thoroughly yet: the role of information about programs' requirements and the platform's rules. Centralized mechanisms are increasingly studied given the richness of the data they generate, and the fact that, often, universities' requirements are fixed ex ante, making it easier to estimate students' preferences. It has thus been generally considered that programs' priorities and requirements should be common knowledge for the applicants. Yet, the fact that information on the application requirements is easily accessible does not mean that all applicants will be aware of them, especially when they can vary both across programs and universities. In fact, whether students have a perfect knowledge of these rules remains an open question. In this paper, I contribute to closing this gap in the literature by documenting the extent of misinformation regarding the requirements embedded in a specific centralized college application system, the Chilean one.

Second, I contribute to understanding further the implications of such misinformation by quantifying its cost with respect to students' admission prospects. In order to do this, I build a structural model of students' application decisions which incorporates the fact that some students might not be aware of the requirements for their application to be valid. This model indeed allows to investigate students' admission outcomes in a counterfactual scenario where uninformed applicants would have been informed about the application requirements, by simulating the lists they would have submitted in this case.

The third contribution of this paper relates to the data-driven way by which incomplete information is captured and incorporated to the model. In particular, I directly exploit data from the centralized application system in Chile in order to learn about students' information set, through a 'revealed information' approach. More specifically, I use the requirements that programs and universities are allowed to impose for an application to be valid. Indeed, each higher-education program is allowed to set minimum test scores requirements and can impose to students to take a particular optional test. Each university can also decide not to consider any application which would be ranked below a certain rank. As in many other systems, if a student fails to meet these conditions, her application will be considered as invalid, implying that she has a probability zero to be admitted to this program. These rules, together with the fact that the number of applications which can be submitted is limited, imply that students should take into account their probabilities of admission when forming their application list (Fack et al. (2019)). Submitting an invalid application is thus at odds with the expected behavior a student should display if being aware of the programs' and universities' requirements.

I will thus consider that submitting an invalid application indicates that the student is not informed about the requirements implemented within the centralized system. Using this strategy, I am able both to quantify the prevalence of misinformation among applicants in Chile, as well as to construct a structural model of students' application decisions incorporating differences in information sets. In particular, I allow for some heterogeneity in students' beliefs: applicants categorized as misinformed fail to set their beliefs over their probabilities of admission to zero for programs for which they do not meet the requirements.

This approach allows to complement the strand of the school choice literature developing methods to estimate students' preferences when they are strategic. Considering that some students might fail to understand the strategic incentives embodied in the mechanisms they study, Agarwal and Somaini (2018) and Calsamiglia et al. (2020) have allowed for two 'types' of students: students can be either truthful, ranking programs according to their true preferences, or strategic, submitting the list which maximizes their expected utility while holding rational expectations over their probabilities of admission. However, there exist no micro-economic foundations to a model where students' equilibrium strategy would be to be truth-telling in the contexts they study (Fack et al. (2019)).

I thus depart from this assumption, and instead rationalize the observed 'obvious mistakes' from the data by allowing students to differ in their information regarding the college application system. Doing this greatly extends the scope of counterfactual policies one can study. Indeed, one of the limitations of using strategic types is that they are black boxes. Then, since it is unclear what makes a student strategic or truth-telling, it is difficult to design policies directly influencing applicants' strategic type. In contrast, linking the fact that some students make obvious mistakes to gaps in their information set makes it clearer how to prevent such behavior.

Agarwal and Somaini (2018) have also considered the fact that students might be strategic, but might not hold rational expectations. They have thus allowed for two additional forms of expectations: adaptive expectations and coarse expectations. The first one considers beliefs that are based on the application process of the previous year. The second allow beliefs to be based only on aggregate information about the number of applicants and capacities at each school. Yet, in contrast to what I suggest, Agarwal and Somaini (2018) do not link these alternative forms of beliefs to what is observed in the data, and rather consider models where all students hold either adaptive expectations, or coarse expectations.

Only two papers are more general with respect to the type of expectations held by students over their probabilities of admission. In particular, they allow for students to be strategic, while having subjective, and thus potentially biased, beliefs. This is the case of Kapor et al. (2020a), who directly collect data over individuals' beliefs in order to incorporate them to a structural model of students' preferences, allowing for rich heterogeneity in their beliefs. In the absence of subjective beliefs data, taking into account this heterogeneity remains extremely challenging. To the best of my knowledge, only one paper estimates separately students' beliefs over their admission probabilities and preferences without relying on subjective beliefs data. Luflade (2019) is able to do so by exploiting the sequential implementation of the Deferred Acceptance (DA) algorithm in the centralized college admission system in Tunisia, which, she argues, incentivizes a subsample of students to report their preferences truthfully, allowing to recover their preferences in a first stage.

One of the drawbacks of observing or estimating subjective beliefs is that it does not tell us how they are formed. If the beliefs recovered do not coincide with rational expectations, it is then difficult to argue that we can correct them, as we do not understand what made them biased in a first place. In contrast to the two above papers, I do not depart from the rational expectations assumption but instead allow for a single dimension of heterogeneity in beliefs - whether or not students have the right information about programs' requirements. This complements the two above papers by suggesting a way to introduce some heterogeneity in students' beliefs without requiring the collection of subjective beliefs data, nor assumptions about students' truth-telling behavior. One of the additional advantages of focusing on a single reason why beliefs are biased, although it allows for less heterogeneity, is that it can be easily linked to a policy intervention having the potential to correct them.

Overall, using the 'revealed information' approach I suggest, I first present descriptive evidence that a large proportion of applicants behave in a way indicating that they are not aware of the requirements implemented in the Chilean centralized application platform. In fact, 26% of the applicants submit a rank-ordered list with at least one invalid application. It also appears that low-income students are significantly more likely to be misinformed about the college-admission mechanism compared to high-income students. Indeed, I find that students whose family belongs to the bottom decile of the income distribution are 6 percentage points more likely to submit an invalid application compared to students from the top income decile. Such mistakes seem to directly affect students' welfare as making a mistake is, in turn, correlated with a lower overall probability of being admitted to any program available within the centralized platform. These descriptive patterns thus suggest that making a mistake is likely to be associated with worse admission outcomes, which will disproportionately affect low income students.

I then exploit the model of students' application decisions I construct, which takes into account the fact that some applicants are not aware of the rules of the system, in order to investigate how making such mistakes impacts students' admission outcomes. In particular, I use the estimated preference parameters in order to simulate the rank-ordered list that uninformed students would have submitted in a scenario where they would have been aware of the application rules. I then determine whether, keeping other students' rank-ordered list constant, the admission status of each uninformed student would have changed compared to the status quo scenario, where the student is uninformed.

Results from this exercise suggest that the complexity of the Chilean application system undermines students' admission chances. Indeed, the fraction of uninformed students that would have been admitted to a program through the centralized application platform in a scenario where they would have submitted a fully valid rank-ordered list is 13 percentage points higher than it is in the status quo scenario. More than 80% of the uninformed students unmatched in the latter scenario would have seen their assignment status changed had they been informed about the platform's application requirements. Since low-income students are the most likely to be misinformed, the complexity of the application system disproportionately harms them: more than 50% of the students who go from being assigned to being unassigned due to misinformation belong to one of the first three deciles of the income distribution in Chile.

The rest of the paper is organized as follows. Section 3.2 describes the college application system in Chile, and documents the existence of heterogeneity in students' information about the rules of the mechanism. Section 3.3 delineates a simple model of students' strategic consideration under biased beliefs, which is consistent with the descriptive evidence. Section 3.4 details the data used for the analysis. Section 3.5 outlines the model of students' preferences, which is estimated using the strategy specified in Section 3.6. Section 3.7 and 3.8 present the results from the estimation and counterfactual policy simulations, respectively. Finally, Section 3.9 lays the ground for future work, while Section 3.10 concludes.

3.2 The Chilean Application Process

3.2.1 Context

The university application system in Chile is semi-centralized: out of the 60 universities present in the country, 41 participate in a centralized application system. Only the most selective universities are included in this centralized mechanism. Among them, 29 are traditional public universities, while 12 are private universities. In order to apply to those universities' programs, students need to take a standardized test (PSU), including Math, Language, Science and Social Sciences. Note that only the Math and Language tests are mandatory. On top of these test scores, a score is also associated to students' high-school ranking and GPA.

Each program can freely choose the weights they assign to those different components in order to compute students' program-specific test scores. Some programs require students to take a particular optional test, while others allow them to use either their Science or Social Sciences test score. Programs can also add some specific requirements for a student's application to be valid, such as a minimum application score or minimum tests scores. Information about each programs' weights, number of seats, past admission score cutoffs and additional requirements is publicly available for students before they take the PSU.

Note that, in Chile, students apply directly to majors. A 'program' will thus designate a university/major pair. Once they received their scores, students can rank up to 10 programs to which they want to apply. The final assignment is determined by an algorithm that only takes into account students' program-specific test score and programs' capacities. Larroucau and Rios (2018) describe the algorithm as a variant of DA, in which tied students in the last seat of a program are not rejected if vacancies are exceeded. The allocation rule is thus as follows:

1. Each student proposes to her first choice according to her submitted rank-ordered list (ROL). Each program rejects any unacceptable student, and if the number of proposals exceeds its vacancies (c), rejects all students whose scores are strictly less than the c-th most preferred student.

2. Step $k \geq 2$. Any student rejected in step $k-1$ proposes to the next program in their submitted ROL. Each program rejects any unacceptable student, and if the number of

proposals exceeds its vacancies (c), rejects all students whose score is strictly less than the c-th most preferred student.

The final allocation is obtained by assigning each student to the most preferred program in her ROL that did not reject her. The score of the last student admitted to a program j can be interpreted as the cutoff score $\bar{t}_j$ and can be written as follows:

$$\text{student } i \text{ is assigned to program } j \Leftrightarrow t_{ij} \geq \bar{t}_j \text{ and } t_{ij} < \bar{t}_{j'}, \forall j' \in R_i \text{ s.t. } j' \succ_{R_i} j$$

where $j, j' \in \mathcal{J}$ with $\mathcal{J}$ the set of programs, R_i denotes student i's ROL and t_{ij} is the test score of student i for program j. $\succ_{R_i}$ is a total order induced by R_i over the set $\{j : j \in R_i\}$, such that $j' \succ_{R_i} j$ if and only if program j' is ranked above program j in R_i.

3.2.2 Access to Information and Socio-Economic Background

While being centralized and information being publicly available, the Chilean application system might still be difficult to navigate for some students. It is particularly likely to be the case because of the requirements that programs can impose to students. Such program-specific rules can take the following forms:

- Programs can impose a minimum math-language average score.
- They can also impose a minimum weighted score.
- They can also impose a particular optional test (either social sciences and history or sciences).

On top of these program-specific requirements, universities are allowed to restrict the type of applications they will consider as valid. In particular:

- They can limit the number of programs a student is allowed to apply to. For example, some universities allow students to apply to a maximum of four of their programs.
- They can restrict the ranks to which a student can list one of the programs they offer. For example, some universities consider only applications if they appear in the students' top four choices.

Finally, two other types of rules at the system-level are imposed. First, applications from students with a math-language score average below 450 are considered as invalid. Second, additional rules are implemented regarding students' applications to pedagogy programs.

Table 3.A.1 describes the programs' key characteristics, as well as their requirements. In particular, about 30% of the programs make it mandatory for students to take a particular optional test. Among them, more than 70% require students to take the test in sciences. More than 80% of the programs also impose a minimum weighted score. The average minimum score imposed is 503. The share of programs setting a minimum math-language average score is slightly lower (65%), which might be explained by the fact that there already exists a minimum math-language average score imposed throughout the platform. The average minimum math-language average score is about 490. 8% of the programs are also subject to a university-specific requirement. In particular, the two most selective universities in Chile require students to rank their programs within the first four ranks on their ROL.

Table 3.A.2 illustrates that these rules add complexity to the application process: more than 25% of the students make at least one mistake in their ROL.[1] Among them, 30% simply do not meet the average math-language average, so that their application is considered as invalid by the system altogether. Among students who have a math-language average high enough in order to be able to submit valid applications to the system, Columns (3) and (4) show that the program-specific and university-specific requirements lead students to fill up their ROL with programs to which they are ineligible. Column 3 illustrates that almost 90% of these applicants submit at least one application which does not meet the programs' requirements, while 16% fail to comply with one of the university-level requirements (Column 4).

Figure 3.A.1 suggests that making such a mistake is correlated with students' socio-economic background. In particular, conditional on test scores, students from higher income households are significantly less likely to make any mistake in their ROL compared to students from low-income families: there is a 6 percentage points difference in the probability to make at least one mistake between the top and bottom income decile (Figure 3.A.2).

Some patterns suggest that variations in the programs' requirements make mistakes more likely. In particular, Figures 3.A.1, 3.A.3 and 3.A.4 show that students applying to programs

[1] A student makes a mistake if she applies to a program while she does not meet at least one of the requirements imposed in order for an application to be valid.

which changed their requirements compared to the previous year are significantly more likely to make mistakes. The same result holds for students applying to programs not available in the year before.

Another informative result from these figures is that, for every type of mistakes, students who already took the PSU test the year before are significantly less likely to make mistakes. This suggests that students learn about the mechanism when navigating it a first time. Another possibility is that some students taking the test for the first time already know that they will apply again in the following year, so that they do not have to carefully check the rules of the programs.

Finally, as illustrated by Figure 3.A.5, making a mistake is not inconsequential when it comes to students' admission outcomes. Students who do not make mistakes in their ROL are significantly more likely to be admitted to one program within the centralized system, even after controlling for students' test scores and number of applications. For example, students making only one mistake are more than 10 percentage points less likely to be assigned to a program within the centralized admission system compared to students submitting fully valid ROL. Each additional mistake significantly reduces the student's probability of admission.

Overall, the above descriptive evidence points to the presence of widespread misinformation regarding the requirements implemented within the Chilean application system, which seems particularly prevalent for low-income students. This analysis also suggests that such misinformation is costly for applicants, as it is correlated with lower admission chances. Yet, measuring the actual cost of being misinformed for students in terms of their admission chances would require to know what would have been their ROL and admission outcome had they been informed about these requirements. In order to perform such a counterfactual exercise, I thus need to build a model of student's application decisions which takes into account that some students are not aware of these rules.

3.3 Illustration : The Consequence of Being Misinformed on the Probability to be Assigned

A toy model of students' strategic considerations under biased beliefs can help rationalize these behaviors. Suppose there are two programs A and B, each with two seats, and three students, $i = 1, 2, 3$. The programs' preferences are assumed to be the same, student 1 being preferred to student 2, who is preferred to student 3. Students 1 and 2 meet the requirements of both programs. In contrast, student 3 does not meet the requirements of program A: she can only be admitted to program B. Preferences for programs are given by u_{iA} and u_{iB}. A student who is not assigned to program A or B chooses the outside option, for which the utility is normalized to zero. Note that u_{iA} and u_{iB} are assumed to be large enough such that the probability that a student prefers the outside option is negligible. Students know the programs' capacities, as well as their priority ranking. Yet, preferences for programs are private information: only their distribution is common knowledge.

Suppose students can only rank one program in their ROL. Consider first the case where all students know the programs to which they can submit valid applications. Since students know their priority ranking, students 1 and 2 can simply truthfully rank programs according to their preferences, and be assigned to their preferred program. Student 3 knows that she does not meet the requirements to program A. She will thus apply to program B. She will remain unassigned only if program B has already reached full capacity.

In contrast, consider now the case where student 3 does not know that her application to program A is invalid. Since she knows that she is ranked third, there is uncertainty regarding her admission chances, given that she does not observe the ROL of the other students. She thus has to take into account her beliefs over her probabilities of admission, q_{3A} and q_{3B}, and will rank the school which maximizes her expected utility $q_{3j}u_{3j}$, for $j \in \{A, B\}$. Since she does not know that her application to program A is invalid, she does not know that her actual probability of admission to this program is 0, but thinks it is positive: $q_{3A} > 0$. She is thus not aware that her actual expected utility from applying to program A is 0. This implies that even in the case where B still has not reached full capacity, the probability that the student ends up unassigned is non null and equal to:

$$\Pr[q_{3A}u_{3A} > q_{3B}u_{3B}]$$

This simple example indicates that not being informed about the requirements imposed by a program might increase the probability that a student ends up unassigned, while she would have been better off assigned to a program which had still some slots available. My goal is thus to estimate a model of student's application decisions in order to investigate more thoroughly the impact of being misinformed on students' application decisions and admission outcomes.

3.4 Data

The empirical analysis I will perform relies on several data sets. The first one gathers students' college applications for the year 2019. It contains the universe of students who took the PSU test in 2019 (306,197 students). In particular, since the PSU test is required in order to apply to one of the 41 universities that participate in the centralized admission process, this sample is composed of all the students who applied to a program available through the centralized system (155,281 students). I restrict this sample to students who only applied through the regular application system (139,947 students).[2]

Observations are at the student-application level. For each application submitted by a student, I observe the admission outcome, her test scores as well as additional details such as whether the student's application to a program was valid. This data set also includes information about students' socio-demographic characteristics. Table 3.A.4 presents summary statistics of some of the key characteristics of applicants to the Chilean centralized system.

The second data set I use contains data at the program level, for the years 2004 to 2019. For each program available through the centralized admission process, I observe their location, the number of seats available, the test scores weights they use, as well as the minimum score requirements they impose in order to consider an application as valid. In 2019, 1,806 programs were offered by universities participating in the centralized mechanism. I complete this dataset with a dataset from MINEDUC, which I use in order to have additional information about the available programs, such as tuition fees, but also their length in terms of number of semesters. Table 3.A.1 summarizes some of the programs' key characteristics.

3.5 Model

In this paper, I build on the literature modeling students' application decisions when they are strategic. In particular, I will use the model developed by Rho (2019), herself building on the

[2]There are two additional ways to apply to the centralized application system, targeting disadvantaged students. 90% of the applicants apply through this regular system only. The remaining students apply to the regular system but also duplicate their applications in order to be considered by a scholarship scheme and/or by the PACE scheme (Programa de Acompanamiento y Acceso Efectivo a la Educacion Superior).

sequential interpretation of a ROL by Calsamiglia et al. (2020). Rho (2019) rewrites explicitly the student's decision as a dynamic discrete choice model, and shows that estimation can be easily performed using a CCP (Conditional Choice Probability) estimation strategy.

Here, I will first consider a case where students have rational expectations over their probabilities of admission and complete information about the programs', universities' and system's requirements, as the rest of the literature (Agarwal and Somaini (2018), Calsamiglia et al. (2020), Rho (2019)). Second, I will relax this later assumption and allow some students not to be aware about the rules of the mechanism, based on what can be observed in the data.

3.5.1 The utility function

Consider an economy with I students, indexed by $i \in \mathcal{I} = \{1, ..., I\}$, and J programs, indexed by $j \in \mathcal{J} = \{0, 1, ..., J\}$, where $j = 0$ denotes the outside option. The outside option includes programs not participating in the centralized admission process, as well as the choice of not pursuing college.

The utility of student i if she attends program j writes:

$$u_{ij} = x_{ij}\theta$$

where x_{ij} is a vector of student-program characteristics including the major of the program, tuition fees, the length of the program in terms of number of semesters, the quality of the program based on its cutoff in 2018 and the distance between the campus of program j and where student i lives. I also add interaction terms between students and programs' characteristics. In particular, I include an interaction term between the gender of the student and a dummy equal to one if the program belongs to one of the STEM majors. This allows students from different gender to have a different taste for STEM majors. I also include interaction terms between the student's income group (divided in quintiles) and the program's tuition fees, as well as between the student's income group and the program's quality.

Note that there is no unobserved heterogeneity here: two students with the same characteristics have the same valuation for a program j. This assumption is further discussed in the next subsection. The utility derived from the outside option is normalized to 0: $u_{i0} = 0$.

3.5.2 Students' Choice Problem

I follow Agarwal and Somaini (2018), Calsamiglia et al. (2020) and Rho (2019) who consider that students submit the ROL which maximizes their expected utility. Estimation of students' preferences is not computationally feasible in the framework developed by Agarwal and Somaini (2018) and Calsamiglia et al. (2020). I will thus use the one suggested by Rho (2019).

As Calsamiglia et al. (2020), Rho (2019) twists the model of Agarwal and Somaini (2018) in order to take into account the dynamic structure of the students' problem, who rank programs sequentially, consistently with the assignment mechanism. This implies that the ROL submitted by students can be viewed as resulting from a dynamic discrete choice problem, where each slot of the list can be considered as a 'time period'.

In particular, each student submits a ROL to the platform, while programs' selection rules determine a ranking among applicants. Given ROL, programs' selection rules and the number of seats available, the DA algorithm determines to which program each student will be matched. Students fill their ROL sequentially, consistently with the assignment mechanism: the DA algorithm sequentially considers programs ranked by a student and only moves down the ROL if the student is rejected from all programs ranked before. Yet, students do not know ex ante whether they will be admitted to a program or not. I thus assume that students will take into account their beliefs over their admission probabilities to the different programs when forming their ROL.

Following Rho (2019), I also assume that a student receives a shock ϵ_{ijr} when ranking a program j at rank r. Since an applicant does not observe future shocks, she has to form expectations over them. A student i then faces the following choice specific expected utility (or conditional value function):

$$v_{ijr} + \epsilon_{ijr} = q_{ijr}u_{ij} + (1 - q_{ijr})\bar{V}_{ir+1} + \epsilon_{ijr} \quad \text{if} \quad j \neq 0 \tag{3.1}$$

and

$$v_{ijr} + \epsilon_{ijr} = u_{ijr} + \epsilon_{ijr} = \epsilon_{ijr} \quad \text{if} \quad j = 0 \tag{3.2}$$

where q_{ijr} is the belief held by student i regarding her probability to be admitted to program j at rank r, and $\bar{V}_{r+1}$ is the expected value of behaving optimally from rank $r+1$ on, conditional on being rejected by j at rank r. If a student fills her ROL up to R then $u_{ijR+1} = u_{i0}$, where

u_{i0} is the value of the outside option, which is normalized to 0.

ϵ_{ijr} is assumed to be identically and independently distributed, following a Type I Extreme Value distribution. The fact that this shock is independent across individuals, programs and ranks and enters additively in the conditional value function is very convenient: it allows to use a CCP estimation strategy, following the insights from Hotz and Miller (1993). This is what is exploited by Rho (2019) in order to make the estimation of students' preferences tractable in a sample with a large choice set, contrasting with the frameworks developed in Agarwal and Somaini (2018) and Calsamiglia et al. (2020).

However, it also comes at the cost of several limitations which should be underlined here. First, this shock is independently distributed across individuals, programs and ranks. This implies that the shock of a same individual for a certain program will not be correlated across ranks. In other words, this model does not incorporate unobserved heterogeneity in preferences that is constant across ranks. This issue is akin to the standard problem implied by the Conditional Independence assumption in dynamic discrete choice models. In order to allow for some unobserved heterogeneity in such models while using a CCP estimation strategy, Arcidiacono and Miller (2011) offer an adaption of the expectation–maximization (EM) algorithm. This solution could thus be applied to the model presented in this section. Yet, it would only allow for a very limited number of unobserved preference types, such that unobserved preference heterogeneity would still be limited.

Second, the fact that this shock enters additively makes it difficult to interpret. Indeed, it implies that a student will choose her application strategy, at each rank, by taking into account her preferences, her beliefs over her probabilities of admission, but also a shock which is generated by the simple fact of ranking a program, whatever the probability of admission associated to it. In contrast, following a revealed preference approach, a model of students' application decisions should take into account that a student will receive a flow utility from a program only if she is admitted to it. Section 3.9 presents an alternative model, which would solve the two above issues. The estimation of this model is left for future work.

I abstract from these issues for now, and show how the assumptions imposed on ϵ_{ijr} can be used to identify students' preferences from data on ROL and estimate them easily. Since ϵ_{ijr} follows a Type I Extreme Value distribution, we have that:

$$\bar{V}_{ir+1} = \gamma + \ln[\sum_{j' \in \{0, \Omega_{ir+1}\}} \exp(v_{ij'r+1})] \tag{3.3}$$

where Ω_{ir+1} is the set of programs still available for student i at rank $r+1$. γ is the Euler constant.

Now, let's arbitrarily consider the choice of the outside option at rank $r+1$. The probability of choosing the outside option, conditional on being rejected from the programs ranked before writes:

$$Pr(d_{i0r+1} = 1 | x_{ij}, \Omega_{ir+1}) = \frac{1}{\sum_{j' \in \{0, \Omega_{ir+1}\}} \exp(v_{ij'r+1})} \tag{3.4}$$

Combining (3.3) and (3.4) gives:

$$\bar{V}_{ir+1} = \gamma - ln[Pr(d_{i0r+1} = 1 | x_{ij}, \Omega_{ir+1})] \tag{3.5}$$

Plugging (3.5) in the conditional value function (3.3) gives the following expression:

$$v_{ijr} + \epsilon_{ijr} = q_{ijr} u_{ij} - (1 - q_{ijr}) \ln[Pr(d_{i0r+1} = 1 | x_{ij}, \Omega_{ir+1})] + (1 - q_{ijr})\gamma + \epsilon_{ijr} \tag{3.6}$$

where both q_{ijr} and $Pr(d_{i0r} = 1 | x_{ij}, \Omega_{ir+1})$ can be estimated from the data, such that the parameters of students' preferences can be identified from Equation (3.6).

3.5.3 Identification of the Probabilities of Admission

Rational Expectations and Complete Information

In the baseline model, I assume that all students have correct beliefs over their probabilities of admission given their program-specific weighted scores, t_{ij}, and the population distribution of weighted scores and ranking strategies. Overall, I assume that each student believes that the lottery, when she reports $R_i \in \mathcal{R}_i$, is as follows:

$$q_i^I = E_{(R_{-i}, t_{-i}^I)}[\Phi((R_i, t_i^I), (R_{-i}, t_{-i}^I)) | R_i, t_i^I] \tag{3.7}$$

where Φ denotes the assignment mechanism function corresponding to the standard DA algorithm, which takes as given students' score vector and ROL in order to determine their admission status to each program. $\Phi_{ij}((R_i, t_i^I), (R_{-i}, t_{-i}^I))$ denotes the admission status of student i at program j and $E_{(R_{-i}, t_{-i}^I)}[.]$ denotes the expectation taken over the random variables $(R_{i'}, t_{i'}^I)$ for $i' \neq i$. Note that $q_i^I = (q_{i1}^I, ...q_{iJ}^I)$, where $q_{ij}^I = \{q_{ijr}^I\}_{r=\{1,...,R\}}$.

Here, I thus follow Agarwal and Somaini (2018) by assuming that students have correct beliefs over their probabilities of admission given their program-specific scores, t_{ij}^I, and the population distribution of scores and ranking strategies. This would for example be a good approximation if students form their beliefs by asking their friends where they apply, as well as their test scores. This assumption contrasts with a model of complete information with respect to others' strategies, where students would be assumed to perfectly observe others' ROL and scores.

On top of this, this model also incorporates the fact that students perfectly know the rules embodied in the mechanism implemented in Chile. In particular, they are aware of the requirements at the mechanism-, university- and program-level that each student must meet for an application to be valid. This assumption is introduced through the notation (t_i^I, t_{-i}^I), with $t_i^I = \{t_{ij}^I\}_{j=\{1,...,J\}}$ and $t_{ij}^I = \{t_{ijr}^I\}_{r=\{1,...,R\}}$ where:

$$t_{i'jr}^I = \mathbf{1}\{i' \text{ satisfies all the requirements imposed by } j \text{ at rank } r\} t_{i'j} \quad \forall i' \in \mathcal{I}, \forall j \in \mathcal{J}$$

where $t_{ij} = \sum_{c \in C} \omega_{jc} t_{ic}$ where C denotes the different score components taken into account by the programs, $C = \{\text{Math, Language, History, Sciences, GPA, High-School Ranking}\}$, ω_{jc} is the weight associated by program j to component c, and t_{ic} is student i's score associated to component c. This notation implies that any student would be perfectly able to tell whether their application or the application of another student is invalid, and would thus be discarded by the algorithm.

Rational Expectations and Incomplete Information

Yet, descriptive evidence presented in Section 3.2 indicates that some students seem to be unaware of these additional rules. My goal is then to capture the fact that some students are misinformed about these requirements by relaxing the assumption of complete information of students regarding the rules of the mechanism. Misinformed students' beliefs about their probabilities of admission will then write as follows:

$$q_i^{NI} = E_{(R_{-i},t_{-i}^{NI})}[\Phi((R_i,t_i^{NI}),(R_{-i},t_{-i}^{NI}))|R_i,t_i^{NI}] \tag{3.8}$$

where Φ denotes the assignment mechanism function corresponding to the standard DA algorithm. Similarly to the above case, students are thus assumed to be aware of the population distribution of ROL and test scores. Yet, here, they are not aware of the requirements put in place by the programs, universities and application system. In particular, they consider program-specific test scores to be as follows:

$$t_{i'j}^{NI} = t_{i'j} = \sum_{c\in C} \omega_{jc} t_{ic} \quad \forall i' \in \mathcal{I}, \forall j \in \mathcal{J}$$

This captures that, in contrast to the above case, students fail to take into account that a student submitting an invalid application cannot be admitted to the program in question.

3.6 Estimation

3.6.1 Probabilities of Admission

In a first stage, I estimate students' beliefs over their probabilities of admission. As described in Section 3.5, our objective is to capture the uncertainty created by the fact that a student does not know the realization of (R_{-i}, t_{-i}) at the time of submitting the report. This implies that the individual has to form expectations for this realization based on the population distribution of test scores and ROL in order to form beliefs over her admission probabilities.

Agarwal and Somaini (2018) propose a resampling procedure in order to approximate this source of uncertainty. Here, I follow their approach. In contrast to them, I have to estimate two types of beliefs, for both informed and uninformed students. I proceed as follows:

Informed Students: Estimation of q^I

1. Set students' test score to zero for programs where they do not meet the admission requirements: this takes into account that students know the requirements of the system.

2. Sample at random with replacement from the observed ROL and test scores.

3. Run the DA algorithm.

The score of the last admitted student of a draw b, denoted $\bar{t}^I_{j,b}$, corresponds to the cutoff score of program j. Each draw thus gives a vector of cutoff scores $\bar{t}^I_b = \{\bar{t}^I_{j,b}\}_{j \in \mathcal{J}}$. We thus have the following estimator for q^I_{ijr}:

$$q^I_{ijr} \approx \frac{1}{B}\sum_{b=1}^{B}[\Phi((R_i, t^I_i), (R_{-i}, t^I_{-i})_b)|R_i, t^I_i]$$

$$\approx \frac{1}{B}\sum_{b=1}^{B}\mathbf{1}\{t^I_{ijr} \geq \bar{t}^I_{j,b}\}$$

$$= \hat{q}^I_{ijr}$$

where $(R_{-i}, t^I_{-i})_b$ is the b^{th} sample of n ROL, B is the number of samples drawn and $\hat{q}^I_{ijr}$ denotes the estimate of q^I_{ijr} obtained through this procedure.

Uninformed Students: Estimation of q^{NI} The procedure in order to estimate beliefs of the uninformed students is the same as what is described above, with the difference that step 1 is not implemented: students' test scores are not set to zero for programs where they do not meet the admission requirements. This change takes into account that students with beliefs q^{NI} are not aware of the requirements implemented through the centralized system. They thus do not take into account that their test score is zero for their own invalid applications, but also fail to take into account that the test scores of other applicants would be set to zero if they try to submit applications to programs where they do not meet the requirements.

3.6.2 Preference Parameters

Case 1: Model under Rational Expectations and Complete Information

I first consider a model where all students have rational expectations about their probabilities of admission, as well as complete information about the rules of the system. In this case, one can simply maximize the following log-likelihood to recover students' preference parameters:

$$\mathcal{L}^I(\theta) = \sum_i \sum_r \ln L^I_{ir}$$

where, from Equation (3.6):

$$L^I_{ir} = \prod_{j\in\{0,\Omega_{ir}\}} [\frac{exp(\hat{q}^I_{ijr}u_{ij}(\theta) - (1-\hat{q}^I_{ijr})\ln[Pr(d_{i0r+1}=1|x_{ij},\Omega_{ir+1})] + (1-\hat{q}^I_{ijr})\gamma)}{1+\sum_{j'\in\{0,\Omega_{ir}\}} exp(\hat{q}^I_{ij'r}u_{ij'} - (1-\hat{q}^I_{ij'r})\ln[Pr(d_{i0r+1}=1|x_{ij},\Omega_{ir+1})] + (1-\hat{q}^I_{ij'r})\gamma)}]^{d_{ijr}} \tag{3.9}$$

where d_{ijr} is a dummy equal to one if student i applies to program j at rank r.

Case 2: Model under Rational Expectations and Incomplete Information

I now consider a second version of the model, where we use the information contained in the data in order to assign an information type to students. More specifically, students who do make a mistake are considered as misinformed, while students who do not make a mistake are considered as having perfect information about the rules of the mechanism. In this case, since types are observed, one can simply maximize the following log-likelihood to recover students' preference parameters:

$$\mathcal{L}^{obs}(\theta) = \sum_i \ln[\mathbb{1}\{i \text{ is informed}\}\Pi_r L^I_{ir} + \mathbb{1}\{i \text{ is not informed}\}\Pi_r L^{NI}_{ir}]$$

where L^I_{ir} writes as in Equation (3.9) and, from Equation (3.6):

$$L^{NI}_{ir} = \prod_{j\in\{0,\Omega_{ir}\}} [\frac{exp(\hat{q}^{NI}_{ijr}u_{ij} - (1-\hat{q}^{NI}_{ijr})\ln[Pr(d_{i0r+1}=1|x_{ij},\Omega_{ir+1})] + (1-\hat{q}^{NI}_{ijr})\gamma)}{1+\sum_{j'\in\{0,\Omega_{ir}\}} exp(\hat{q}^{NI}_{ij'r}u_{ij'} - (1-\hat{q}^{NI}_{ij'r})\ln[Pr(d_{i0r+1}=1|x_{ij},\Omega_{ir+1})] + (1-\hat{q}^{NI}_{ij'r})\gamma)}]^{d_{ijr}} \tag{3.10}$$

3.7 Results

3.7.1 Conditional Choice Probabilities of Choosing the Outside Option

Figure 3.A.6 presents the density of the CCP of the outside option at different ranks. As expected, the probability to choose the outside option increases with the rank of the list. These CCP are estimated using a flexible logit, including different programs' characteristics, interactions between students' and programs' characteristics, the distance between students and programs, as well as its square.

Note that it includes interactions between some programs' characteristics and a dummy indicating whether the student is informed or not. This allows CCP to vary according to students' information type, which will be used in order to update the CCP of the uninformed students in the counterfactual scenario where they receive information about the mechanism's rules.

3.7.2 Students' Preferences

Table 3.A.5 presents the parameter estimates for different models and specifications of students' preferences. Columns (1) and (3) present the estimates from a model where all students are assumed to have complete information about the rules of the mechanism. Columns (2) and (4) present the estimates from a model where students who are observed to make a mistake are assumed to have an imperfect knowledge of the rules of the mechanism. In particular, this implies that they are not aware that their probability of admission to a program is null when they fail to meet its requirements.

In both models, the most parsimonious specification indicates that students have a strong distaste for private schools. This might be explained by the fact that, historically, the most renowned universities in Chile are public. As expected, students do not like programs that are far away from their hometown. Interestingly, while applicants value positively STEM programs, female students have a strong distaste for the latter. In fact, female students are willing to travel about 84 kilometers more in order to enroll in a non-STEM program. The low magnitude of the coefficient associated to the length of the program indicates that this variable is not a key determinant of students' application choices. The positive coefficient related to tuition fees is likely to signal that high tuition fees are correlated to a high unobserved program quality, that we fail to capture through this specification. It might be particularly true in Chile, where full or partial scholarships have become widely available to low-income students thanks to recent reforms. If programs with high tuition fees are of good quality, they might then attract many applicants, driving the parameter to be positive. Note that this is not an issue regarding the counterfactual exercise I implement using these estimates. In order to be able to perform a valid counterfactual exercise, what is required is that the parameter estimates I use are policy-invariant, but not that they are causal. To the extent that I am not trying to capture the causal effect of tuition fees on application decisions, the parameter estimate corresponding to tuition fees can thus reflect a correlation rather than a causal effect. Columns (3) and (4) show that students from high income households have a higher preference for programs with high tuition fees compared to low income students. Table 3.A.5 presents results from an estimation including major fixed effects. The most popular majors seem to be medicine, law and architecture.

Parameter estimates do not change a lot across the different models. Yet, some interesting differences appear. For example, the distaste for distance increases when we take into account that some students are not informed about their actual probabilities of admission. This suggests that misinformed students apply to programs with binding requirements located far away from their home. When not taking heterogeneity in information into account, the only way to rationalize this for the model is for the distaste for distance to be lower than what it is in reality. Similarly, the taste for public universities decreases in the model taking into account students' heterogeneity in information. Again, it suggests that misinformed students are applying to programs in public universities for which they have a zero admission chance. The only way to rationalize this for a model where every students are perfectly informed was to artificially increase the students' taste for public universities.

One can also compare the results from a model taking into account that students are strategic with other popular models used to estimate students' preferences. Some studies have for example assumed that students rank programs truthfully under a DA algorithm (Abdulkadiroğlu et al. (2017)). Of course, such models do not take into account that students might strategize to form their ROL, which likely results in biased estimates. Results from such a comparison are quite striking (Table 3.A.7). For example, the parameter associated to private universities becomes positive when students are assumed to be truthful: the model does not take into account that many students might apply to private universities because they are less selective than public universities in Chile (Table 3.A.9). The model under the truth-telling assumption also predicts a lower taste for STEM programs, as it fails to take into account that students apply to STEM programs despite having lower probabilities of admission to them (Table 3.A.9). It is also interesting to note that the estimated distance coefficient is higher in the model taking into account that students are strategic: it captures the fact that students applying to programs close to their home might do so despite their lower probability of admission there. This might be explained by the fact that many students live in Santiago and apply to programs offered by the two most selective universities of the country, both located in the capital, despite having low probabilities of admission to them (Table 3.A.9).

Another popular estimation strategy in settings where the DA algorithm is used relies on the assumption that the outcome of the match is stable: all students are matched to their most preferred feasible program (Fack et al. (2019)). An underlying assumption in such cases is that students can perfectly anticipate programs' cutoffs. This is unlikely to be the case in Chile given that the number of seats per program is relatively low, so that programs' cutoffs vary across years. On top of this, programs' choice of weights and minimum scores requirements vary over time, making it difficult for students to anticipate the cutoffs from one year to the next. Table 3.A.8 contrasts estimates under the stability assumption with the ones taking into account that students are strategic, while some are not aware of the rules of the mechanism. Note that a comparison is difficult given the fact that only the sample of admitted students is used under the stability assumption. Any differences could thus come from the fact that the sample is not representative of the population of applicants. The model under stability captures the distaste for private universities, although the magnitude of the coefficient is substantially lower than when assuming that students are strategic. It also gives similar estimates than the model under truth-telling for the parameter associated to STEM.

3.8 Counterfactual Simulations

The fact that some students have incomplete information about the rules of the admission mechanism implies that they might make payoff-relevant mistakes in their ROL. In order to understand what is the impact of the complexity of the application system with respect to students' admission outcomes, it is necessary to simulate the ROL which would have been submitted by uninformed students in a counterfactual scenario where they would have been aware of the application rules. In this section, I thus use the parameter estimates obtained from the model of students' application decisions delineated above in order to implement this analysis.

3.8.1 Implementation

Through this counterfactual exercise, my goal is to compare the admission result of uninformed students in a scenario where they would have been informed about the platform's requirements to the one prevailing in the status quo scenario, i.e. when they are not aware of these rules. This exercise is implemented by removing all the programs to which a student has a zero-probability of admission from her choice set.[3] I describe the implementation of this counterfactual in what follows.

First, I use the parameters estimated in the previous section in order to obtain the ROL predicted by the model.[4] More specifically, I first randomly draw, for each student-program-rank combination, a shock from a Type I Extreme Value distribution. I then use this shock together with the parameter estimates to compute the students' conditional value from applying to each program, at each rank. From this, for each student-rank combination, I select the program with the highest conditional value among those which have not been ranked before. Doing this for each rank gives the students' ROL under the status quo scenario, i.e. the ROL predicted by the model before implementing the counterfactual exercise. I will refer to this ROL as the 'status quo' ROL. This will constitute the benchmark ROL to which I will compare the one obtained under the counterfactual scenario where students are aware of the rules for an application to be valid.

I then construct the ROL predicted by the model under the counterfactual scenario. To take into account that students are now aware of the platform's rules, I drop from each uninformed student's choice set all the programs to which they would not meet at least one of the requirements. Once this is done, I again construct the ROL as described above.

Note that, in order to compute students' expected utility in this counterfactual scenario, I correct their CCP of choosing the outside option at each rank, which appears in the penalty term of Equation (3.6), to take into account that they are now informed. I thus have to compute counterfactual conditional choice probabilities. In order to do that, I use the parameter estimates of the same flexible logit estimated in order to get the CCP to choose the outside option, but switching the dummy capturing whether a student is informed to be equal to one for previously uninformed students. All the other observable characteristics included in the above logit model remain fixed, such that a previously uninformed student will have the same conditional choice probability to apply to the outside option at a given rank than an

[3] Note that my objective for the future is to implement this counterfactual by directly correcting students' probabilities of admission. Yet, for now, the model I use predicts that the conditional choice probability of a student to apply to any program, even one where she has a zero-probability of admission, is positive. I thus plan to twist the model in order to make sure this CCP is equal to zero in the future. For now, I circumvent this issue by simply removing programs where students have a zero probability of admission from their choice sets: this ensures they will not apply to such a program under the counterfactual policy of interest.

[4] Note that in order to take into account that informed students will not submit invalid applications, I remove programs for which they do not meet the requirements from their choice sets.

informed student sharing the same observable characteristics. Overall, this gives rise to the ROL that would be selected by a student when being informed.

I am interested in knowing whether misinformation resulted in some students being rejected by all the programs they ranked, while they could have been admitted to a program they like had they been informed about the platform's requirements. Note that it amounts to focusing on the partial equilibrium effect of a policy preventing students from submitting invalid applications: would a previously uninformed student would be admitted to a program if we prevent him from submitting invalid applications, fixing other students' ROL? In order to determine the admission outcome of a student in this case, I use the program-specific score threshold observed in the data. Thus, a student will be admitted in the counterfactual scenario if:

$$t_{ij} > \bar{t}_j$$

where $\bar{t}_j$ is program j's threshold observed in the data, i.e. the score of the last admitted student in 2019.

Finally, note that I only use the sample of students with a math-language average above 450. Indeed, students who do not meet this criteria have a zero probability to be admitted to any program.

3.8.2 Results

Students who were previously making a mistake automatically have to change their ROL given the counterfactual scenario under study. Of course, the fact that students do change their ROL does not imply that their admission outcome will be impacted. Table 3.A.10 presents the results from this exercise, using estimates from the model presented in Table 3.A.5, Column (4). I select a random sample of 2,747 students. The predictions from the model match quite well the percentage of students with a math-language average above 450 who make at least one mistake in their ROL (19.91% in the data vs 21.82% predicted by the model) (Panel A, Table 3.A.10).

Panel B of Table 3.A.10 presents the students' admission outcome both under the status quo and counterfactual scenario. I find that 83.72% of the students with an average above 450 but who made at least one mistake in their ROL would be matched to a program in the status quo scenario. In contrast, 96.78% of these students would have been admitted to a program in the counterfactual scenario. The fraction of uninformed students that would have been admitted to a program through the centralized application platform in a scenario where they would have submitted a fully valid ROL is thus 13 percentage points higher than it is in the status quo scenario. Overall, more than 80% of previously unmatched uninformed students would have been assigned, had they been informed about the platform's application requirements. This suggests that the complexity of the application system negatively impacts students' admission prospects.

This is particularly true for low-income students, who are the most likely to be misinformed: more than 50% of the students who go from being assigned to being unassigned due to misinformation belong to one of the first three deciles of the income distribution in Chile (Table 3.A.11).

3.8.3 Counterfactual Policy: Preventing Students from Submitting Invalid Applications

The above results suggest that a simple policy, implemented directly on the online application platform, which would prevent students from submitting invalid applications, could significantly improve the admission prospects of students previously not aware of the programs' and universities' requirements. It would also have the potential to help closing the socio-economic gap in the access to higher education, as it would disproportionately affect low-income students.

Evaluating this policy would require to take into account its general equilibrium effects. In particular, students having rational expectations, they should update their beliefs over their probabilities of admission under the counterfactual scenario so that they are consistent with the new equilibrium. The policy considered here is for example likely to increase the share of applicants to programs with less stringent requirements, since previously uninformed students who applied to programs for which their application was invalid should now reshuffle their applications towards programs to which their probability of admission is not zero. Taking into account these general equilibrium effects is left for future work.

3.9 Extensions

3.9.1 Model with Unobserved Information Types

Making a mistake is a sufficient condition for the student to be uninformed. Yet, it might be that some students have incomplete information but are not considered as such because they fail to apply to programs with stringent rules, simply because they do not like them. Such missclassification might thus create a bias in the parameter estimates of the model.

In order to circumvent the limitation of assigning directly information types from what is observed in the data, a solution would be to estimate a mixture of informed and uninformed types. The joint log-likelihood writes:

$$\mathcal{L}^{unobs}(\theta) = \sum_i \ln[\pi \prod_r L_{ir}^I + (1-\pi) \prod_r L_{ir}^{NI}] \tag{3.11}$$

where L_{ir}^I (resp. L_{ir}^{NI}) writes as in Equation (3.9) (resp. (3.10)) and Π is the population probability to be informed.

In order to estimate θ, I will follow Arcidiacono and Miller (2011), who adapt the EM algorithm to CCP estimators. The identification proof of this model is left for future work.

3.9.2 Alternative Model

The main issue of the model presented in Section 3.5 is the presence of the unobserved preference shock which is independently distributed across students, programs and ranks, and which enters additively in the conditional value function. These assumptions imply 1) that there is no unobserved preference heterogeneity which can be correlated across ranks and 2) that students receive a utility flow from the mere fact of ranking a program, whatever their probability of admission to it.

Standard models in the literature (Agarwal and Somaini (2018), Calsamiglia et al. (2020)) do not include this additive shock at the student-program-rank level. My objective is thus to follow more closely their framework. In particular, I will consider the following model:

$$V_{ir} = \max_j \{q_{ijr}(u_{ij} + \epsilon_{ij}) + (1 - q_{ijr})V_{ir+1}\} \tag{3.12}$$

where q_{ijr} is the student i's beliefs over her probability to be admitted to program j at rank r, u_{ij} is the flow utility that student i would get from being matched to j and ϵ_{ij} is the unobserved preference of student i for program j. It is identically and independently distributed, following a Type I Extreme Value distribution. V_{ir+1} represents the continuation value for i at rank $r+1$.

Two differences can be noted between this model and the one presented in Section 3.5. First, in this model, the idiosyncratic shock is at the student-program level, which implies that the unobserved preference of a student i for a program j will be the same across all rank. This thus introduces unobserved preference heterogeneity which is correlated across ranks. Second, the taste shock does not enter the model additively anymore. This allows to interpret it as an unobserved preference shock, as it only realizes if the student is admitted to the program to which she applies.

Estimating such a model is burdensome. Agarwal and Somaini (2018) propose to use a Gibbs sampler, which can only be used when the number of possible ROL which can be submitted by a student is relatively small. In their application, for example, students can rank up to 3 schools out of 13. Similarly, Kapor et al. (2020a) can use this method as, in their application in New Haven, students can rank up to 4 schools out of 12. In contrast, in Chile, students can rank up to 10 programs out of almost 2,000.

Calsamiglia et al. (2020) propose a more tractable approach, by showing that the students' problem can be solved by Backward Induction. They then estimate their model using SMLE. However, given that there is a large number of programs that a student can rank, obtaining consistent estimates might require to simulate the model a large number of times for each student and each iteration of the maximization. It is thus intractable in practice. I leave for future work the provision of a tractable estimation method, which could be used to study large matching markets, as the one in Chile.

3.10 Conclusion

Centralized application systems have been presented as having the potential to close the existing gap in access to higher education with respect to students' socio-economic background. Yet, such centralized mechanisms usually come with their set of rules and requirements, which might make it difficult for some students to navigate it. Ultimately, the complexity of a platform might prevent students from submitting the list which maximizes their expected utility.

In this paper, I study the Chilean centralized application system, leveraging data over students' ROL, together with the programs' and universities' requirements made available publicly, in order to determine whether a student is informed about these rules or not. I find that, in Chile, 26% of the students submit a list containing at least one invalid application with respect to the platform's requirements, suggesting that misinformation is widespread. Such mistakes are correlated with a lower overall probability of being admitted to a program available within the centralized platform. Low-income students might be disproportionately affected by these lower admission prospects, as I find that they are significantly more likely to submit an invalid application than high-income students.

To study the implications of the complexity of the application mechanism for students' admission outcomes, I estimate a model of students' application behavior, taking into account that some students are not aware of the requirements potentially invalidating their application. Counterfactual policy simulations suggest that being misinformed about the application rules is costly. Indeed, more than 80% of previously unassigned uninformed students would have been assigned had they been informed about the platform's application requirements. This is particularly true for low-income students, who are the most likely to be misinformed: more than 50% of the students who go from being assigned to being unassigned due to misinformation belong to one of the first three deciles of the income distribution in Chile. Preventing students from submitting invalid applications to the online platform thus has the potential to close the socio-economic gap in the access to higher education. This policy would have general equilibrium effects which should be taken into account in order to evaluate it, which is thus left for future work.

Appendices

3.A Additional Figures and Tables

Table 3.A.1: Summary Statistics: Programs' Characteristics

	Programs' Characteristics	
	Mean	Standard Deviation
Private University	0.33	0.47
Number of Semesters	9.61	1.86
Annual Program Fee	3,521,691.94	1,061,780.61
STEM	0.51	0.50
Cutoff	536.36	73.99
Weight assigned to HS grades	14.07	5.87
Weight assigned to HS Rank	24.35	10.98
Weight assigned to Language	21.46	9.99
Weight assigned to Maths	26.85	11.16
Weight assigned to History and Social Sciences	9.37	7.64
Weight assigned to Sciences	11.47	6.50
Observations	1,806	

Notes. This table presents summary statistics of the characteristics of the programs available on the platform.

Table 3.A.2: Summary Statistics: Programs' Application Requirements

	Programs' Requirements	
	Mean	Standard Deviation
Choice across Optional Test	0.61	0.49
Sciences Test Required (if no choice)	0.72	0.45
Has a Minimum Weighted Score	0.82	0.39
Minimum Weighted Score (if has one)	503.71	52.38
Has a Minimum Math-Language Average Score	0.65	0.48
Minimum Math-Language Average Score (if has one)	489.14	22.89
Subject to a University Ranking Rule	0.08	0.26
Observations	1,806	

NOTES. This table presents summary statistics of the requirements imposed by the programs available on the platform.

Table 3.A.3: Number and Percentage of Students Making Different Types of Mistakes in their ROL

	(1) At least one mistake	(2) Average Math-Language below 450	Students with Math-Language>450: (3) Does not comply with programs' requirements	(4) Does not comply with universities' requirements
Nber Students	36,935	10,957	23,111	4,262
Percentage	26.39	29.66	88.96	16.41
Total	139,947	36,935	25,978	25,978

NOTES. This table presents summary statistics on the students making different types of mistakes in their submitted rank-ordered list.

Figure 3.A.1: Results from a Probit Regression: Probability to Make a Mistake

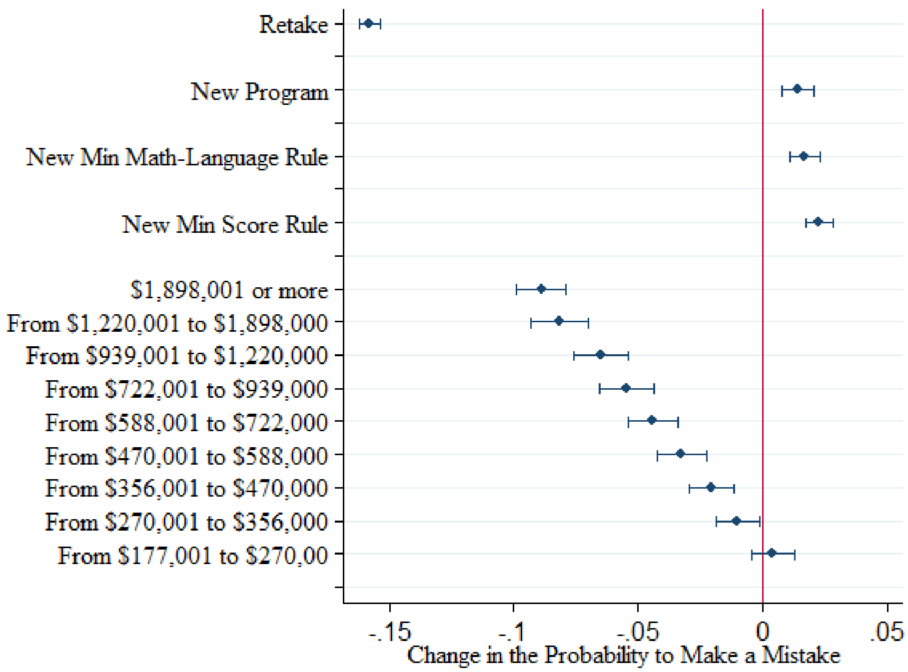

Notes. Average marginal effects from a probit regression. The dependent variable is a dummy equal to one if the student made at least one mistake in her application strategy. 'New Program' is a dummy equal to one when the student included at least one new program in her ROL. 'New min score rule' is a dummy equal to one when the student included at least one program which changed its weighted score requirement in 2019 compared to 2018 in her ROL. 'New min math-language rule' is a dummy equal to one when the student included at least one program which changed its average math-language average score requirement in 2019 compared to 2018 in her ROL. Retake is a dummy equal to one when the student already applied in the year before. Additional controls include the number of applications submitted, the student's high-school GPA, gender and mother's education.

Figure 3.A.2: Predicted Probability to Make a Mistake with Respect to Income

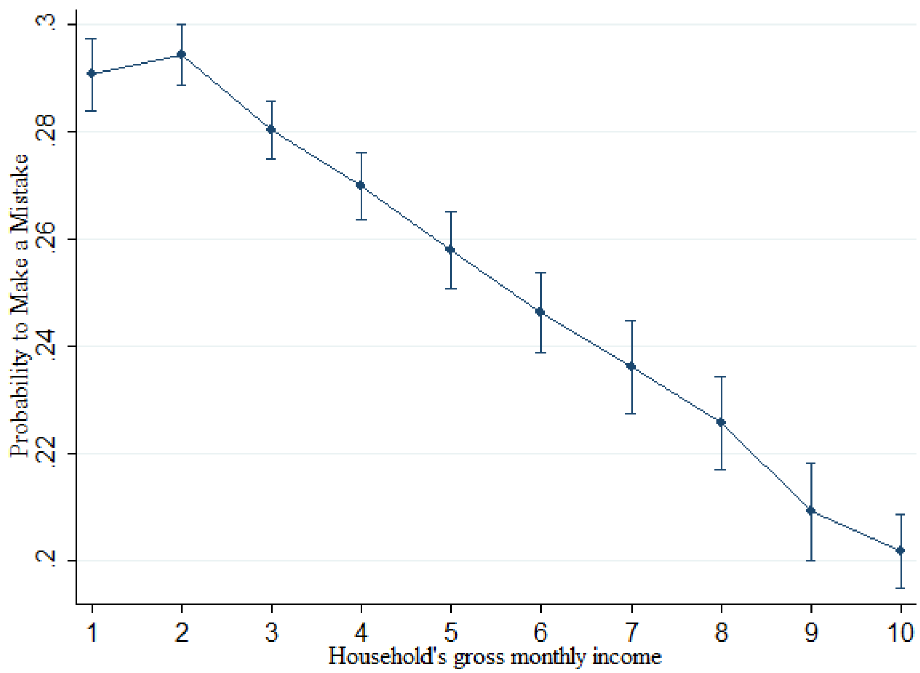

NOTES. Predicted probability that the student makes at least one mistake in her application list with respect to her income group. Group 1 corresponds to the bottom decile of the income distribution, while Group 10 corresponds to the top decile of the income distribution. Results use the parameter estimates from the probit regression presented in Figure 3.A.1. Other control variables are set to their average to compute these predicted probabilities.

Figure 3.A.3: Results from a Probit Regression: Probability to Submit an Invalid Application with respect to a Minimum Weighted Score Requirement.

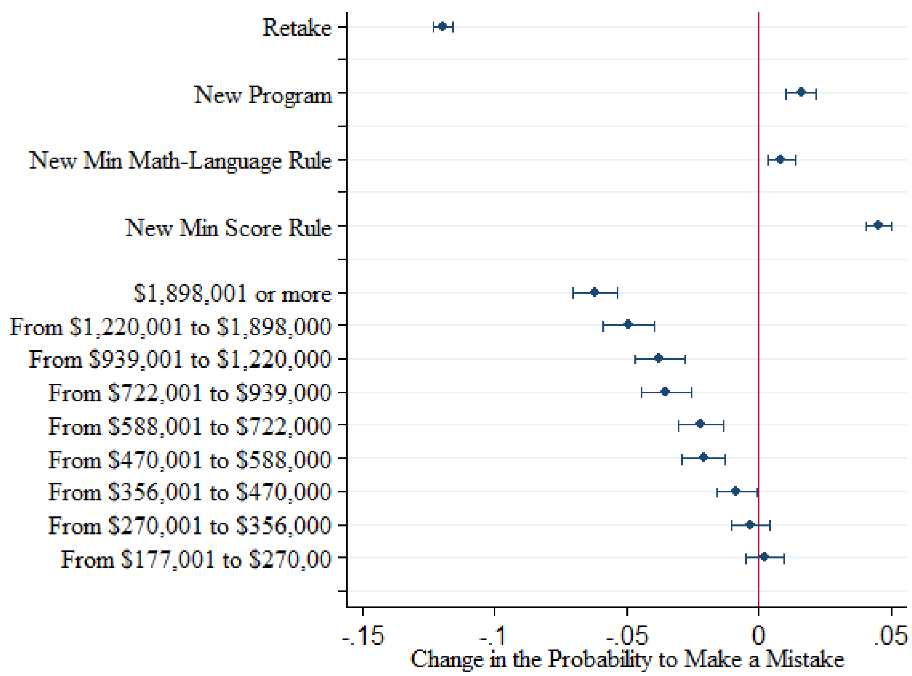

NOTES. Average marginal effects from a probit regression. The dependent variable is a dummy equal to one if the student made at least one application to a program to which her weighted score was lower than the one required. 'New Program' is a dummy equal to one when the student included at least one new program in her ROL. 'New min score rule' is a dummy equal to one when the student included at least one program which changed its weighted score requirement in 2019 compared to 2018 in her ROL. 'New min average math-language rule' is a dummy equal to one when the student included at least one program which changed its average math-language average score requirement in 2019 compared to 2018 in her ROL. Retake is a dummy equal to one when the student already applied in the year before. Additional controls include the number of applications submitted, the student's high-school GPA, gender and mother's education.

Figure 3.A.4: Results from a Probit Regression: Probability to Submit an Invalid Application with respect to a Minimum Math-Language Average Score Requirement.

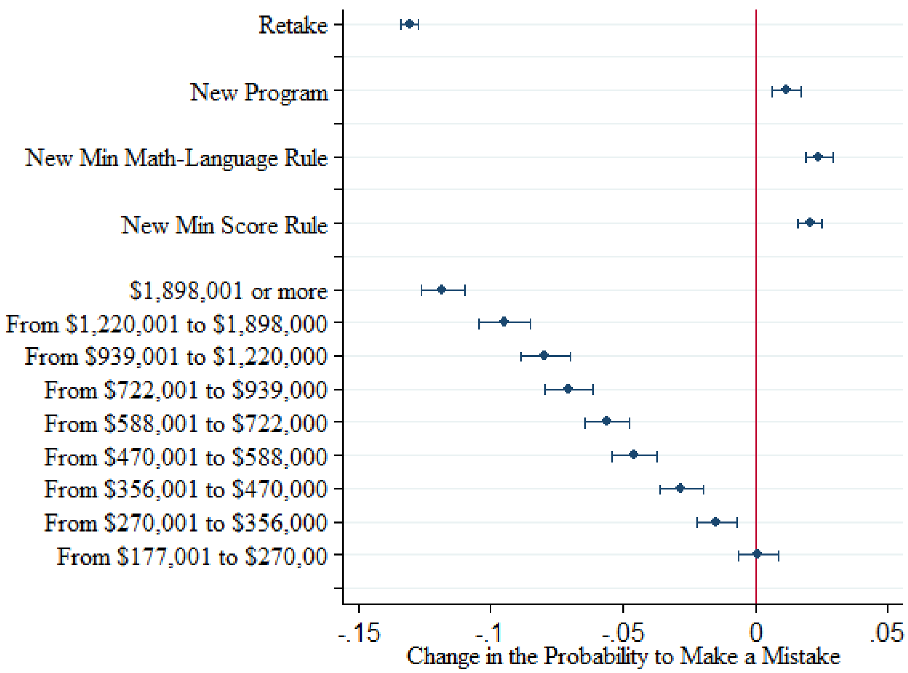

NOTES. Average marginal effects from a probit regression. The dependent variable is a dummy equal to one if the student made at least one application to a program to which her math-language average was lower than the one required. 'New Program' is a dummy equal to one when the student included at least one new program in her ROL. 'New min score rule' is a dummy equal to one when the student included at least one program which changed its weighted score requirement in 2019 compared to 2018 in her ROL. 'New min average math-language rule' is a dummy equal to one when the student included at least one program which changed its average math-language average score requirement in 2019 compared to 2018 in her ROL. Retake is a dummy equal to one when the student already applied in the year before. Additional controls include the number of applications submitted, the student's high-school GPA, gender and mother's education.

Figure 3.A.5: Predicted Probability to be Admitted to a Program within the Centralized System with respect to the Number of Mistakes

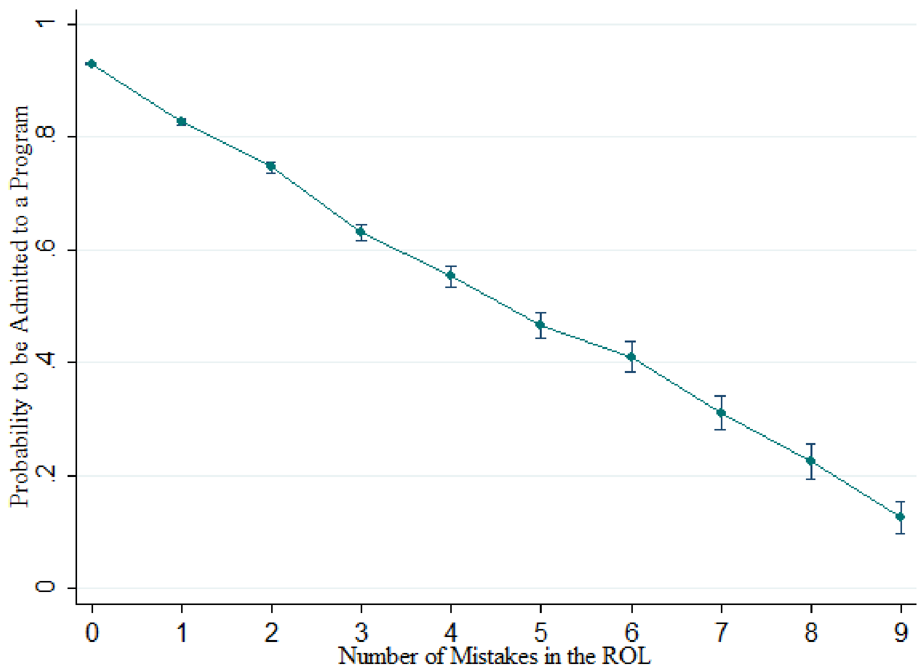

NOTES. This graph illustrates the probability that a student is admitted to a program available within the centralized system as a function of the number of invalid applications included in her ROL. The x-axis corresponds to the number of invalid applications contained in the student's ROL. 0 corresponds to a fully valid ROL. These predicted probabilities are constructed using a probit regression. The dependent variable is a dummy equal to one if the student was admitted to a program within the centralized admission mechanism. Controls include the number of applications submitted and the student's test scores. They are set to their average to compute predicted probabilities. Only students with an average math-language score above 450 are included.

Table 3.A.4: Summary Statistics: Students' Characteristics

	Students' Characteristics	
	Mean	Standard Deviation
Female	0.55	0.50
High-School Attended:		
Private	0.19	0.39
Private Subsidized	0.55	0.50
Public	0.26	0.44
Household's Income Group:		
From \$0 to \$177,000	0.11	0.31
From \$177,001 to \$270,000	0.15	0.36
From \$270,001 to \$356,000	0.16	0.37
From \$356,001 to \$470,000	0.11	0.32
From \$470,001 to \$588,000	0.09	0.28
From \$588,001 to \$722,000	0.08	0.27
From \$722,001 to \$939,000	0.06	0.24
From \$939,001 to \$1,220,000	0.06	0.24
From \$1,220,001 to \$1,898,000	0.06	0.23
\$1,898,001 or more	0.12	0.33
Mother's Education :		
Basic Education	0.15	0.35
Secondary Education	0.62	0.49
Higher Education	0.23	0.42
Language Test Score	508.50	178.61
Math Test Score	504.04	182.25
History, Geography and Social Sciences Test Score (if test taken)	554.69	97.01
Science Test Score (if test taken)	541.82	100.39
Retake	0.18	0.39
Observations	139,947	

Figure 3.A.6: CCP of the Outside Option

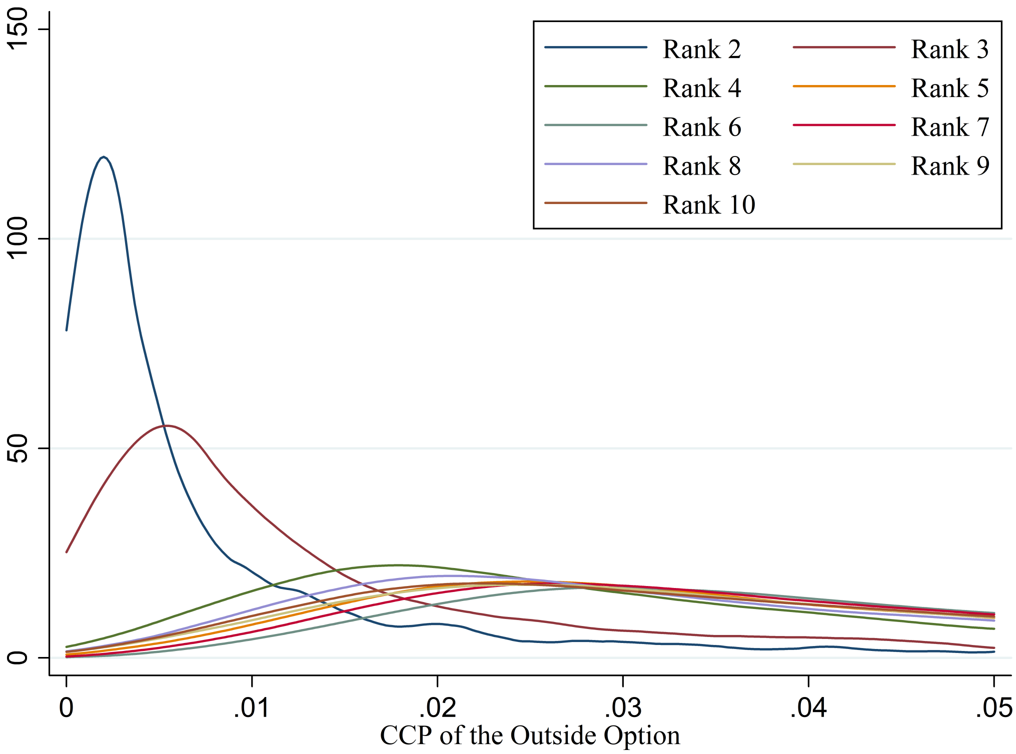

NOTES. Distribution of the CCP of choosing the outside option at different ranks.

Table 3.A.5: Preference Parameters

	(1) Rho	(2) Rho Two Types	(3) Rho	(4) Rho Two Types
Private	-0.765***	-0.739***	-0.779***	-0.749***
	(0.0319)	(0.0312)	(0.0319)	(0.0312)
Tuition Fees	0.384***	0.376***	0.264***	0.250***
	(0.0143)	(0.0140)	(0.0259)	(0.0254)
STEM	0.199***	0.217***	0.204***	0.223***
	(0.0333)	(0.0328)	(0.0333)	(0.0329)
Distance (100 km)	-0.565***	-0.592***	-0.570***	-0.596***
	(0.00871)	(0.00876)	(0.00876)	(0.00881)
Female X STEM	-0.478***	-0.508***	-0.487***	-0.516***
	(0.0437)	(0.0430)	(0.0437)	(0.0431)
Cutoff 2018	0.0014***	0.0013***	0.0025***	0.0024***
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Nber of Semesters	-0.0255***	-0.0227***	-0.0270***	-0.0242***
	(0.00609)	(0.00600)	(0.00608)	(0.00599)
Tuition Fees				
X Income Group 2			0.00915	0.0197
			(0.0329)	(0.0322)
X Income Group 3			0.119***	0.122***
			(0.0356)	(0.0349)
X Income Group 4			0.189***	0.190***
			(0.0380)	(0.0376)
X Income Group 5			0.312***	0.323***
			(0.0317)	(0.0311)
Cutoff 2018				
X Income Group 2			-0.0001	-0.0001
			(0.0003)	(0.0002)
X Income Group 3			-0.0009***	-0.0008**
			(0.0003)	(0.0003)
X Income Group 4			-0.0015***	-0.0015***
			(0.0003)	(0.0003)
X Income Group 5			-0.0034***	-0.0034***
			(0.0003)	(0.0003)
Observations	2,747	2,747	2,747	2,747

NOTES. Columns (1) and (3) present the coefficients from a model following Rho (2019) with all students being aware of the rules of the mechanisms, while columns (2) and (4) present the coefficients from a model with two types of students, according to their knowledge of programs' requirements. Groups correspond to quantiles of the income distribution. Standard errors in parentheses.

Table 3.A.6: Results: Preference Parameters

	(1) Rho		(2) Rho Two Types	
Private	-1.502***	(0.0340)	-1.472***	(0.0332)
Tuition Fees	0.247***	(0.0267)	0.231***	(0.0261)
Female X STEM	-0.569***	(0.0419)	-0.598***	(0.0412)
Distance (100 km)	-0.589***	(0.00867)	-0.611***	(0.00867)
Cutoff 2018	-0.00470***	(0.000218)	-0.00470***	(0.000213)
Nber of Semesters	-0.123***	(0.00581)	-0.122***	(0.00571)
Tuition Fees				
X Income Group 2	0.00587	(0.0327)	0.0152	(0.0319)
X Income Group 3	0.115**	(0.0354)	0.117***	(0.0347)
X Income Group 4	0.186***	(0.0380)	0.183***	(0.0375)
X Income Group 5	0.309***	(0.0317)	0.316***	(0.0311)
Cutoff 2018				
X Income Group 2	-0.00000291	(0.000252)	-0.0000500	(0.000246)
X Income Group 3	-0.000849**	(0.000280)	-0.000739**	(0.000274)
X Income Group 4	-0.00132***	(0.000306)	-0.00126***	(0.000302)
X Income Group 5	-0.00284***	(0.000259)	-0.00275***	(0.000254)
Major:				
Engineering	-0.902***	(0.0318)	-0.895***	(0.0314)
Medicine	0.256***	(0.0346)	0.288***	(0.0340)
Sciences	-0.491***	(0.0551)	-0.493***	(0.0542)
Law	0.504***	(0.0492)	0.436***	(0.0483)
Business	-0.650***	(0.0390)	-0.618***	(0.0379)
Education	-1.162***	(0.0384)	-1.150***	(0.0372)
Architecture	0.209**	(0.0688)	0.258***	(0.0676)
Psychology	0.0830	(0.0541)	0.123*	(0.0533)
Arts	-1.224***	(0.0685)	-1.212***	(0.0672)
Social Sciences	-0.613***	(0.0817)	-0.701***	(0.0814)
Humanities	-1.374***	(0.0664)	-1.362***	(0.0638)
Services	-0.699***	(0.0448)	-0.701***	(0.0440)
Observations	2,747		2,747	

Notes. Column (1) presents the coefficients from a model following Rho (2019) with all students being aware of the rules of the mechanisms, while column (2) present the coefficients from a model with two types of students, according to their knowledge of programs' requirements. Groups correspond to quantiles of the income distribution. Standard errors in parentheses.

Table 3.A.7: Preferences: Strategic vs Truthtelling

	(1) Rho Two Types	(2) Truth-telling	(3) Rho Two Types	(4) Truth-telling	(5) Rho Two Types	(6) Truth-telling
Private	-0.739***	0.151***	-0.749***	0.142***	-1.472***	0.0135
	(0.0312)	(0.0260)	(0.0312)	(0.0259)	(0.0332)	(0.0277)
Tuition Fees	0.376***	0.105***	0.250***	-0.0339	0.231***	-0.114***
	(0.0140)	(0.0109)	(0.0254)	(0.0185)	(0.0261)	(0.0195)
STEM	0.217***	0.0524	0.223***	0.0456		
	(0.0328)	(0.0273)	(0.0329)	(0.0274)		
Distance (100km)	-0.592***	-0.548***	-0.596***	-0.546***	-0.611***	-0.542***
	(0.00876)	(0.00612)	(0.00881)	(0.00609)	(0.00867)	(0.00605)
STEM X female	-0.508***	-0.222***	-0.516***	-0.198***	-0.598***	-0.199***
	(0.0430)	(0.0357)	(0.0431)	(0.0358)	(0.0412)	(0.0358)
Cutoff 2018	0.0013***	0.0035***	0.0024***	0.0032***	-0.0047***	0.0028***
	(0.0001)	(0.0002)	(0.0002)	(0.0003)	(0.0002)	(0.0003)
Nber Semesters	-0.0227***	0.0392***	-0.0242***	0.0348***	-0.122***	0.0173***
	(0.00600)	(0.00516)	(0.00599)	(0.00514)	(0.00571)	(0.00520)
Tuition Fees						
X Income Group 2			0.0197	0.0582*	0.0152	0.0612*
			(0.0322)	(0.0233)	(0.0319)	(0.0241)
X Income Group 3			0.122***	0.118***	0.117***	0.126***
			(0.0349)	(0.0259)	(0.0347)	(0.0268)
X Income Group 4			0.190***	0.210***	0.183***	0.227***
			(0.0376)	(0.0286)	(0.0375)	(0.0296)
X Income Group 5			0.323***	0.374***	0.316***	0.416***
			(0.0311)	(0.0240)	(0.0311)	(0.0247)
Cutoff 2018						
X Income Group 2			-0.0001	-0.0007	-0.0001	-0.0007
			(0.0002)	(0.0003)	(0.0002)	(0.0003)
X Income Group 3			-0.0008**	-0.0009*	-0.0007**	-0.0009*
			(0.0003)	(0.0004)	(0.0003)	(0.0004)
X Income Group 4			-0.0015***	-0.0006	-0.0013***	-0.0006
			(0.0003)	(0.0004)	(0.0003)	(0.0004)
X Income Group 5			-0.0034***	0.0035***	-0.0027***	0.0034***
			(0.0003)	(0.0004)	(0.0002)	(0.0004)
Major FE	No	No	No	No	Yes	Yes
Observations	2,747	2,747	2,747	2,747	2,747	2,747

NOTES. Columns (1), (3) and (5) present the coefficients from a model following Rho (2019), with two types of students, according to their knowledge of programs' requirements. Columns (2), (4) and (6) present the coefficients from a model under the assumption that students are truth-telling. Groups correspond to quantiles of the income distribution. Standard errors in parentheses.

Table 3.A.8: Preferences: Strategic vs Stable Equilibrium

	(1) Rho Two Types	(2) Stability	(3) Rho Two Types	(4) Stability	(5) Rho Two Types	(6) Stability
Private	-0.739***	-0.218***	-0.749***	-0.211**	-1.472***	-0.286***
	(0.0312)	(0.0659)	(0.0312)	(0.0658)	(0.0332)	(0.0713)
Tuition Fees	0.376***	0.339***	0.250***	0.213***	0.231***	0.105*
	(0.0140)	(0.0262)	(0.0254)	(0.0476)	(0.0261)	(0.0514)
STEM	0.217***	0.0666	0.223***	0.0685		
	(0.0328)	(0.0685)	(0.0329)	(0.0687)		
Distance (100 km)	-0.592***	-0.650***	-0.596***	-0.647***	-0.611***	-0.642***
	(0.00876)	(0.0183)	(0.00881)	(0.0182)	(0.00867)	(0.0181)
Female X STEM	-0.508***	-0.204*	-0.516***	-0.193*	-0.598***	-0.199*
	(0.0430)	(0.0934)	(0.0431)	(0.0937)	(0.0412)	(0.0934)
Cutoff 2018	0.0013***	0.0167***	0.0024***	0.0146***	-0.0047***	0.0148***
	(0.0001)	(0.0006)	(0.0002)	(0.0012)	(0.0002)	(0.00124)
Nber of Semesters	-0.0227***	0.00658	-0.0242***	0.00657	-0.122***	-0.0121
	(0.00600)	(0.0130)	(0.00599)	(0.0130)	(0.00571)	(0.0132)
Income Group:						
Tuition Fees						
X Income Group 2			0.0197	-0.0130	0.0152	-0.0142
			(0.0322)	(0.0606)	(0.0319)	(0.0638)
X Income Group 3			0.122***	0.0861	0.117***	0.0967
			(0.0349)	(0.0650)	(0.0347)	(0.0684)
X Income Group 4			0.190***	0.152*	0.183***	0.173*
			(0.0376)	(0.0698)	(0.0375)	(0.0737)
X Income Group 5			0.323***	0.389***	0.316***	0.445***
			(0.0311)	(0.0586)	(0.0311)	(0.0619)
Cutoff 2018						
X Income Group 2			-0.000104	0.000586	-0.0000500	0.000566
			(0.000252)	(0.00166)	(0.000246)	(0.00166)
X Income Group 3			-0.0008**	-0.0010	-0.0007**	-0.0011
			(0.0003)	(0.0018)	(0.0003)	(0.0018)
X Income Group 4			-0.0015***	-0.0009	-0.0012***	-0.0010
			(0.0003)	(0.0019)	(0.0003)	(0.0019)
X Income Group 5			-0.0034***	0.0096***	-0.0027***	0.0092***
			(0.0003)	(0.0018)	(0.0002)	(0.0017)
Major FE	No	No	No	No	Yes	Yes
Observations	2,747	2,747	2,747	2,747	2,747	2,747

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

NOTES. Columns (1), (3) and (5) present the coefficients from a model following Rho (2019), with two types of students, according to their knowledge of programs' requirements. Columns (2), (4) and (6) present the coefficients from a model under the assumption that the match is stable. Groups correspond to quintiles of the income distribution.

Table 3.A.9: Selectivity of Different Programs

	Cutoffs	Difference
Public	552.86	50.26^{***}
Private	502.60	
Santiago	529.74	-30.63^{***}
Outside Santiago	560.37	
Non-STEM	523.60	-24.18^{***}
STEM	547.78	

NOTES. This table summarizes the differences in programs' cutoffs according to their characteristics. Programs in public universities and located within Santiago have higher cutoffs, implying they are more selective. This is also true for STEM programs. Differences are significantly different from 0 at the 1% level.

Table 3.A.10: Model Fit and Counterfactual Results

Panel A		
	Data	Status Quo
% of students making at least one mistake in their ROL	19.91	21.86
Panel B		
	Status Quo	Counterfactual
% of students making at least one mistake in their status quo ROL who are admitted to a program	83.72	96.78

NOTES. Counterfactual simulations, for a random sample of 2% of the applicants with a math-language average score above 450. The determination of admission results is based on the programs' cutoff observed in the data.

Table 3.A.11: Counterfactual Results by Income Group

Household's Gross Monthly Income	% Uninformed Students Gaining Admission
From \$0 to \$177,000	12.00
From \$177,001 to \$270,000	18.67
From \$270,001 to \$356,000	21.33
From \$356,001 to \$470,000	8.00
From \$470,001 to \$588,000	16.00
From \$588,001 to \$722,000	13.33
From \$722,001 to \$939,000	1.33
From \$939,001 to \$1,220,000	2.67
From \$1,220,001 to \$1,898,000	2.67
\$1,898,001 or more	4.00
Total	100.00

NOTES. This table presents the share of uninformed students gaining admission in the counterfactual scenario. It restricts the sample to the sample of uninformed students who were unmatched in the status quo scenario, but admitted to a program in the counterfactual scenario.

Chapter 4

Robustness of Two-Way Fixed Effects Estimators to Heterogeneous Treatment Effects

Abstract. This paper provides necessary and sufficient conditions for the Two-Way Fixed Effects (TWFE) estimator to be robust to heterogeneous treatment effects. I decompose the TWFE estimator to show that it is a weighted sum of five different types of two-by-two comparisons, with positive weights. I show that parallel trends assumptions on either the untreated or treated potential outcomes must hold for each comparison to identify the Average Treatment Effect (ATE) of the group switching treatment status, when the effect of the treatment is contemporaneous. Both parallel trends assumptions are thus necessary and sufficient for the TWFE estimator to weigh each ATE positively, when allowing treatment effects to be heterogeneous across groups and periods. I further provide sufficient conditions under which the TWFE estimator remains valid even in the presence of dynamic treatment effects. Finally, I show how to exploit all available comparisons to build unbiased estimators of the ATT and ATE.

4.1 Introduction

Difference-in-differences is one of the most popular quasi-experimental methods to estimate the effect of a policy. Its core idea is to compare the evolution of the outcome of interest before and after a group receives a treatment to the one of a group which does not receive it. With two groups and two periods, such a comparison is an unbiased estimator of the Average Treatment Effect on the Treated (ATT). However, most empirical applications depart from this simple framework, and instead leverage the variation in treatment exposure across multiple groups and time periods. With several groups and periods, researchers will typically interpret as the ATT the parameter of the treatment dummy in a Two-Way Fixed Effects (TWFE) regression, also including group and time fixed effects.

Despite its popularity, the properties of this estimator remained under-studied until recently. The literature has now concluded that, under the standard common trends assumption, the TWFE estimator corresponds to a weighted sum of the Average Treatment Effects (ATEs) in each group and period, with weights that may be negative in the presence of heterogeneous treatment effects (de Chaisemartin and d'Haultfoeuille, 2020a). The TWFE estimator may then be negative even when all ATEs are positive.

This paper provides necessary and sufficient conditions for the TWFE estimator not to weigh any ATEs negatively, when they can be heterogeneous across groups and over time. I show that this requires not only a parallel trends assumption on the *untreated* potential outcomes, but also on the *treated* potential outcomes. Under these conditions, the TWFE estimator is thus heterogeneity-robust. Furthermore, I show that its weights can be corrected to build unbiased estimators of the ATT and of the ATE.

To derive these results, I consider a general framework where groups are allowed to enter and exit treatment over time. It is thus not restricted to the staggered case, which has received a lot of attention in spite of relatively few applications (de Chaisemartin and d'Haultfoeuille, 2020a). I first show that the TWFE estimator is a weighted sum of five different types of two-by-two difference-in-differences comparisons. It may include (i) standard, (ii) reverse, (iii) leaver, (iv) reverse leaver, and (v) double-switcher difference-in-differences. The four first comparisons contrast a group which switches treatment status (enters or leaves treatment) with a group keeping the same treatment status (either treated or untreated). Comparisons (v) contrast the evolution of outcomes of a group joining treatment to a group leaving it.

In a framework where the effect of the treatment can only be contemporaneous, I first establish that each of these comparisons is an unbiased estimator of the ATE of the group switching treatment status if and only if a parallel trends assumption either on the *untreated* or *treated* potential outcomes holds. While most researchers are only familiar with the standard difference-in-differences, these intermediary results broaden the panel of comparisons available to them. In particular, I show how the double-switcher difference-in-differences - comparing, across two time periods, a group switching from being treated to untreated, to a group switching from being untreated to treated - can help identify treatment effects of interest.

I then establish the implication of these results for the TWFE estimator: it is a weighted sum of ATEs in different groups and periods, all weighted positively. Each ATE enters proportionally to the number of comparison groups available to identify it. When allowing for heterogeneous treatment effects over time and across groups, necessary and sufficient conditions for this result are surprising: the trends in potential outcomes both when *untreated* and *treated* should evolve similarly across groups. Together, these common trends assumptions imply that ATEs should follow the same evolution over time across groups, but do not impose homogeneity.

I further investigate the robustness of the static TWFE estimator to the presence of dynamic treatment effects. On top of parallel trends assumptions, the absence of such dynamics is a sufficient, but not necessary, condition for the TWFE estimator to weigh all ATEs positively. I provide conditions, allowing for dynamics, for the TWFE estimator to be heterogeneity-robust. In particular, I consider a setting where common trends assumptions conditional on being never-treated and on joining treatment for the first time hold. I assume that the incremental effect of having been exposed to treatment for n periods is homogeneous across groups for a given time period, while ATEs of first exposure can be heterogeneous. In this framework, I show that there exist relevant settings in which all two-by-two comparisons, and thus the TWFE estimator, remain valid.

Finally, I conclude with the two main implications of this paper for applied researchers. First, the conditions for the TWFE estimator not to be under the threat of negative weights are such that testing for parallel trends on the untreated potential outcomes is not sufficient. Whenever using the TWFE estimator, one should also test for parallel trends of potential outcomes across groups when they are treated. I summarize tests available to the practitioner.[1] Second, I take stock of the fact that the TWFE estimator does not weigh each ATE by the corresponding sample size and propose alternative estimators. In particular, I exploit the existing valid comparisons which receive zero-weight in the heterogeneity-robust

[1]See Roth (2022) for recommended practices to perform such tests.

estimator suggested by de Chaisemartin and d'Haultfoeuille (2020a). I expand the latter by building an unbiased estimator of the ATT of all switching cells under less stringent assumptions. I further show that one can build, under adequate restrictions, unbiased estimators of the ATT - requiring a parallel trends assumption on the untreated potential outcomes only - and ATE.

Related Literature The difference-in-differences literature has recently put the TWFE estimator under increased scrutiny. de Chaisemartin and d'Haultfoeuille (2020a) study a general framework with several groups and time periods where groups can enter and leave treatment, without considering dynamic treatment effects. They conclude that the TWFE estimator is a weighted sum of ATEs in each group and period, with weights that may be negative when ATEs are heterogeneous over time or across groups. Borusyak et al. (2022) reach the same conclusion in the staggered case, where groups cannot exit treatment.

This paper revisits these results and provides necessary and sufficient conditions for the TWFE estimator to be robust to heterogeneity in treatment effects. I show that it is not required for ATEs to be homogeneous across groups, nor over time for the TWFE estimator to weigh all ATEs positively.

Both Goodman-Bacon (2021) and Strezhnev (2018) provide decompositions of the TWFE estimator in terms of two-by-two difference-in-differences comparisons to shed light on the origin of negative weights. Focusing on the staggered setting, Goodman-Bacon (2021) establishes that negative weights come from the comparisons of late treated units with already-treated units, which are biased estimators of the corresponding ATE in the presence of heterogeneous treatment effects over time. These comparisons are thus understood as 'forbidden comparisons'. Borusyak et al. (2022) conclude that 'these comparisons are only valid when the homogeneity assumption is true'. While Strezhnev (2018) expands the TWFE decomposition to the general case where groups can leave treatment, he then focuses on the staggered case and reaches the same conclusion.

The novel identification result of this paper builds on these decompositions. By carefully deriving the conditions which are both necessary and sufficient for the TWFE estimator to be heterogeneity-robust, I show that the so-called 'forbidden comparisons' are actually valid if the potential outcomes under treatment evolve similarly across groups. Importantly, and in contrast to the general conclusion of these papers, this does not require treatment effects to be homogeneous over time, nor across groups.

Finally, while the above papers have focused on the staggered case, I explore the properties of each of the five types of two-by-two difference-in-differences comparisons which enter the TWFE estimator. I establish the assumptions under which each comparison is an unbiased estimator of its corresponding ATE. On top of allowing to derive the main result of this paper, it also allows to expand the tool box of the applied economist. Apart from the standard difference-in-differences, only the reverse difference-in-differences has received some attention in the literature (Kim and Lee, 2019). Additionally, I highlight how the different comparisons can be combined to build unbiased estimators of the ATT and ATE, as well as to expand the heterogeneity-robust estimator suggested by de Chaisemartin and d'Haultfoeuille (2020a).

It should be noted that, following de Chaisemartin and d'Haultfoeuille (2020a), I do not consider cases with non-binary, continuous or several treatments (de Chaisemartin and d'Haultfoeuille, 2018; de Chaisemartin and d'Haultfoeuille, 2020b; Callaway et al., 2021). de Chaisemartin and D'Haultfoeuille (2022b) and Roth et al. (2023) provide rich reviews of this recent literature.

4.2 Framework

I use the same framework and notations as de Chaisemartin and d'Haultfoeuille (2020a). In particular, consider observations that are divided across G groups and T periods. For each group g in period t, we observe a number $N_{g,t}$ of individuals. Let us denote $D_{i,g,t}$ the binary treatment status of individual i in group g at period t, and $(Y_{i,g,t}(0), Y_{i,g,t}(1))$ the potential outcomes when untreated and when treated, respectively. The observed outcome of individual i in group g at period t is denoted $Y_{i,g,t}(D_{i,g,t})$. The following objects can be defined, for all $(g,t) \in \{1, ..., G\} \times \{1, ..., T\}$:

$$D_{g,t} = \frac{1}{N_{g,t}} \sum_{i=1}^{N_{g,t}} D_{i,g,t}, \qquad Y_{g,t}(0) = \frac{1}{N_{g,t}} \sum_{i=1}^{N_{g,t}} Y_{i,g,t}(0),$$

$$Y_{g,t}(1) = \frac{1}{N_{g,t}} \sum_{i=1}^{N_{g,t}} Y_{i,g,t}(1), \qquad \text{and} \qquad Y_{g,t} = \frac{1}{N_{g,t}} \sum_{i=1}^{N_{g,t}} Y_{i,g,t}.$$

where $D_{g,t}$, $Y_{g,t}(0)$, $Y_{g,t}(1)$, and $Y_{g,t}$ denote the average treatment, the average potential outcomes when untreated and when treated, and the average observed outcome in group g at period t, respectively. $\mathbf{\Omega}$ collects the potential outcomes and treatment status for all g and t, i.e. $\mathbf{\Omega} = \{Y_{g,t}(0), Y_{g,t}(1), D_{g,t}\}_{\forall g,t}$.

I also impose the same assumptions as de Chaisemartin and d'Haultfoeuille (2020a). Discussions of these assumptions can be found in their paper.

Assumption 1 *(Balanced Panel of Groups):* $\forall (g,t) \in \{1, ..., G\} \times \{1, ..., T\}, N_{g,t} > 0$.

Assumption 2 *(Sharp Design):* $\forall (g,t) \in \{1, ..., G\} \times \{1, ..., T\}$ *and* $i \in \{1, ..., N_{g,t}\}, D_{i,g,t} = D_{g,t}$.

Assumption 3 *(Independent Groups): The vectors* $(Y_{g,t}(0), Y_{g,t}(1), D_{g,t})_{1 \leq t \leq T}$ *are mutually independent.*

Assumption 4 *(Strong Exogeneity):* $\forall (g,t,d) \in \{1, ..., G\} \times \{2, ..., T\} \times \{0, 1\}$,
$E(Y_{g,t}(d) - Y_{g,t-1}(d) | D_{g,1}, ..., D_{g,T}) = E(Y_{g,t}(d) - Y_{g,t-1}(d))$.

These assumptions impose very few restrictions on $\mathbf{\Omega}$, and in particular do not restrict treatment effects to be homogeneous across groups, nor over time. It however rules out dynamic treatment effects, i.e. treatment effects cannot depend on the treatment history of the group. I relax this assumption in Section 4.4.3.

Finally, let us define some objects of interest. The ATE of group g in period t, $\Delta_{g,t}$, writes:

$$\Delta_{g,t} = \frac{1}{N_{g,t}} \sum_{i=1}^{N_{g,t}} [Y_{i,g,t}(1) - Y_{i,g,t}(0)]$$

Defining $N^{(1)} = \sum_{i,g,t} D_{i,g,t}$ as the number of treated units, the expected average treatment on the treated, δ^{TR}, writes:

$$\delta^{TR} = E\left[\sum_{g,t:D_{g,t}=1} \frac{N_{g,t}}{N^{(1)}} \Delta_{g,t}\right]$$

We consider the following TWFE regression:

$$Y_{g,t} = \alpha_g + \alpha_t + \beta_{fe} D_{g,t} + \epsilon_{g,t} \tag{4.1}$$

We let $\hat{\beta}_{fe}$ denote the OLS estimator of the coefficient of $D_{g,t}$, with $\beta_{fe} = E[\hat{\beta}_{fe}]$. de Chaisemartin and d'Haultfoeuille (2020a) show that, under Assumptions 1-4 and a common trends assumption on *untreated* potential outcomes, β_{fe} is equal to the expectation of a weighted sum of $\Delta_{g,t}$, with potentially negative weights when ATEs are heterogeneous across groups or over time.[2] This result implies that $\hat{\beta}_{fe}$ may be negative even if all ATEs are positive. The general intuition is that negative weights arise because the TWFE estimator includes so-called 'forbidden comparisons', comparing the outcome evolution of some groups to invalid control groups, such as always-treated units (de Chaisemartin and D'Haultfoeuille, 2022b). In the next sections, I revisit these results and provide sufficient and necessary conditions for the TWFE estimator to weigh all ATEs positively, while allowing ATEs to be heterogeneous across groups and over time.

4.3 A General Decomposition of the TWFE Estimator

What are the necessary and sufficient conditions for the TWFE estimator to weigh all ATEs positively? To answer this question, a crucial first step is to decompose the TWFE estimator as a weighted sum of standard two-by-two difference-in-differences comparisons. Such a decomposition will make it straightforward to understand what each difference estimates, and under which assumptions. Importantly, Strezhnev (2018) provides a decomposition of the TWFE estimator in the general case. He shows that the TWFE estimator can be written as a uniform average of difference-in-differences comparisons:

$$\hat{\beta}_{fe} = \frac{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t'\neq t}[(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t'\neq t}[1 - D_{g,t'} + D_{k,t'}]} \tag{4.2}$$

I now decompose this object further in order to clarify which comparisons enter the TWFE estimator.[3] In particular, I show that $\hat{\beta}_{fe}$ is a weighted sum of five different types of two-by-two comparisons. $\forall g \in \{1,..,G\}, \forall k \in \{1,..,G\} \setminus g$, $\forall t \in \{1,..,T\}, \forall t' \in \{1,..,T\} \setminus t$, we let $\hat{\beta}^{DD}_{g,k,t,t'}$ denote the two-by-two comparison of the outcomes of group g and k between periods

[2] Results in de Chaisemartin and d'Haultfoeuille (2020a) are derived when considering the regression of $Y_{i,g,t}$ on group fixed effects, period fixed effects and $D_{g,t}$, i.e. using more disaggregated outcome data. de Chaisemartin and d'Haultfoeuille (2022b) extend them to the aggregated version considered in this paper. The latter is equivalent to the one using individual-data, up to re-weighing by the population in each group. Re-weighing only matters for the variance of the estimator, not its unbiasedness and consistency: we can thus abstract from re-weighing, following de Chaisemartin and D'Haultfoeuille (2022b).

[3] Strezhnev (2018) interprets Equation (4.2) informally, and focusing on the staggered setting.

t and t':

$$\hat{\beta}^{DD}_{g,k,t,t'} = (Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})$$

We can now define the five objects entering the TWFE estimator:

Standard Difference-in-Differences, $\hat{\beta}^{S}_{g,k,t,t'}$:

$$\hat{\beta}^{S}_{g,k,t,t'} = \hat{\beta}^{DD}_{g,k,t,t'} \times \underbrace{\mathbb{1}\{D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 0, t > t'\}}_{\equiv \omega^{S}_{g,k,t,t'}} \tag{4.3}$$

Equation (4.3) corresponds to the standard difference-in-differences comparison, where the evolution of the outcome of a group, g, which becomes treated between period t' and t is compared to the one of a group, k, which is untreated in both periods.

Reverse Difference-in-Differences, $\hat{\beta}^{R}_{g,k,t,t'}$:

$$\hat{\beta}^{R}_{g,k,t,t'} = \hat{\beta}^{DD}_{g,k,t,t'} \times \underbrace{\mathbb{1}\{D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 1, D_{k,t'} = 1, t > t'\}}_{\equiv \omega^{R}_{g,k,t,t'}} \tag{4.4}$$

Equation (4.4) describes the reverse difference-in-differences comparison.[4] It contrasts the evolution of the outcome of a group, g, which becomes treated between period t' and t with the one of a group, k, which is treated in both periods.

Leaver Difference-in-Differences, $\hat{\beta}^{L}_{g,k,t,t'}$:

$$\hat{\beta}^{L}_{g,k,t,t'} = \hat{\beta}^{DD}_{g,k,t,t'} \times \underbrace{\mathbb{1}\{D_{g,t} = 1, D_{g,t'} = 1, D_{k,t} = 0, D_{k,t'} = 1, t > t'\}}_{\equiv \omega^{L}_{g,k,t,t'}} \tag{4.5}$$

Equation (4.5) describes the leaver difference-in-differences comparison. It compares the evolution of the outcome of a group, g, between period t' and t, which is treated in both periods with the one of a group, k, which is treated in period t' but has left treatment in period t. Units which remain treated are thus used as a control group for units leaving treatment. This comparison does not exist in a staggered design, where groups cannot leave

[4]Several studies have used such an empirical strategy, such as Rossi and Villar (2020) and Chabé-Ferret and Voia (2021).

treatment.

Reverse Leaver Difference-in-Differences, $\hat{\beta}^{RL}_{g,k,t,t'}$:

$$\hat{\beta}^{RL}_{g,k,t,t'} = \hat{\beta}^{DD}_{g,k,t,t'} \times \underbrace{\mathbb{1}\{D_{g,t}=0, D_{g,t'}=0, D_{k,t}=0, D_{k,t'}=1, t>t'\}}_{\equiv \omega^{RL}_{g,k,t,t'}} \tag{4.6}$$

Equation (4.6) describes the reverse leaver difference-in-differences comparison. It compares the evolution of the outcome of a group, g, between period t' and t, which is treated in neither of the two periods, with the one of a group which is treated in period t' but has left treatment in period t. Similarly, this comparison does not appear in staggered designs.

Double-Switcher Difference-in-Differences, $\hat{\beta}^{DS}_{g,k,t,t'}$:

$$\hat{\beta}^{DS}_{g,k,t,t'} = \hat{\beta}^{DD}_{g,k,t,t'} \times \underbrace{\mathbb{1}\{D_{g,t}=1, D_{g,t'}=0, D_{k,t}=0, D_{k,t'}=1, t>t'\}}_{\equiv \omega^{DS}_{g,k,t,t'}} \tag{4.7}$$

Equation (4.7) introduces a two-by-two comparison which has received seldom attention in the literature, the double-switcher difference-in-differences. It compares the evolution of outcomes of two groups: group g is not treated in period t' and joins treatment in period t, while group k is treated in period t' but leaves treatment in period t.

We can now rewrite the TWFE estimator as a sum, with positive weights, of these five different types of two-by-two comparisons.

Theorem 1 *$\hat{\beta}_{fe}$ is a weighted sum of five different types of two-by-two difference-in-differences:*

$$\hat{\beta}_{fe} = \frac{\sum_t \sum_{t' \neq t} \sum_g \sum_{k \neq g} [\hat{\beta}^{S}_{g,k,t,t'} + \hat{\beta}^{R}_{k,g,t',t} + \hat{\beta}^{L}_{g,k,t,t'} + \hat{\beta}^{RL}_{k,g,t',t} + 2\hat{\beta}^{DS}_{g,k,t,t'}]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t} [1 - D_{g,t'} + D_{k,t'}]}$$

where, for $g \neq k$, $t \neq t'$ and $c \in \{S, L, R, RL, DS\}$:

$$\hat{\beta}^{c}_{g,k,t,t'} = \omega^{c}_{g,k,t,t'} \hat{\beta}^{DD}_{g,k,t,t'}$$

Details of the proof are provided in Appendix 4.A.1. The above decomposition shows that, in the general case, the TWFE estimator includes five types of two-by-two difference-in-differences comparisons. Each weight simply corresponds to a dummy equal to one each time a relevant comparison group exists. Next section clarifies what each of these comparisons identifies, and under which assumptions.

4.4 Identification

Section 4.4.1 studies conditions under which each of the five two-by-two comparisons defined above identifies the ATE of the group switching treatment status. Section 4.4.2 highlights the implication of this result for the TWFE estimator. When allowing for heterogeneous treatment effects over time and across groups, parallel trends assumptions on both *untreated* and *treated* potential outcomes are necessary and sufficient for the latter to be a weighted sum of ATEs, with positive weights. Finally, Section 4.4.3 allows for dynamic treatment effects, and provides conditions under which the static TWFE estimator remains valid.

4.4.1 The Five Difference-in-Differences Comparisons: Identification

Let us first study separately each of the two-by-two difference-in-differences comparisons that may be included in the TWFE estimator. In what follows, I consider the set of $\mathbf{\Omega}$ such that Assumptions 1-4 hold.

Standard Difference-in-Differences, $\hat{\beta}^S_{g,k,t,t'}$: Consider two groups g and k, two periods t, t', such that $t > t'$, $D_{g,t} = 1$, $D_{g,t'} = 0$, $D_{k,t} = 0$ and $D_{k,t'} = 0$. We are thus in the standard case where we observe a group which remains untreated between period t' and t, and a group which becomes treated. It is well-established that the following common trends assumption is required for the standard difference-in-differences estimator to identify the ATT, i.e. the ATE of group g in time t.

Assumption 5 *(Common Trends of the Potential Outcome Without Treatment): For $t \geq 2$, $E(Y_{g,t}(0) - Y_{g,t-1}(0))$ does not vary across g.*

We can write:

Assumption 5 holds if and only if $\forall g, k \neq g, t, t' \neq t, \forall \mathbf{\Omega}, E[\hat{\beta}^S_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t}|\mathbf{D}] \times \omega^S_{g,k,t,t'}$

where $\mathbf{D}$ is the vector of treatment status history for every group. Details can be found in Appendix 4.A.2.

Reverse Difference-in-Differences, $\hat{\beta}^R_{g,k,t,t'}$**:** Consider two groups g and k, two periods t, t', such that $t > t'$, $D_{g,t} = 1$, $D_{g,t'} = 0$, $D_{k,t} = 1$ and $D_{k,t'} = 1$. We are in a case where always-treated units are used as a control group for late-treated units. These comparisons are shown to be at the origin of negative weights in the TWFE estimator (Goodman-Bacon, 2021), and are presented as 'forbidden comparisons'. In particular, under Assumption 5, we have:

$$E[\hat{\beta}^R_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t} - (\Delta_{k,t} - \Delta_{k,t'})|\mathbf{D}] \times \omega^R_{g,k,t,t'}$$

Thus, if $E[\Delta_{k,t}|\mathbf{D}] \neq E[\Delta_{k,t'}|\mathbf{D}]$ and $E[\Delta_{k,t}|\mathbf{D}] \neq E[\Delta_{g,t}|\mathbf{D}]$, i.e. in the presence of heterogeneous treatment effects over time and across groups, these comparisons are biased estimators of $E[\Delta_{g,t}|\mathbf{D}]$.[5] This is why they are typically understood as forbidden.

Yet, Kim and Lee (2019) show that, under a common trends assumption on the potential outcomes under treatment, these two-by-two comparisons are unbiased estimators of the ATE of group g in the pre-treatment period, t'. We thus consider the following assumption:

Assumption 6 *(Common Trends of the Potential Outcome With Treatment): For* $t \geq 2$, $E(Y_{g,t}(1) - Y_{g,t-1}(1))$ *does not vary across* g.

In particular, adding and subtracting $E[Y_{g,t'}(1)|D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 1, D_{k,t'} = 1]$, we have:

$$E[\hat{\beta}^{DD}_{g,k,t,t'}|D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 1, D_{k,t'} = 1]$$

[5]Note that this qualifies the statement of Goodman-Bacon (2021) who notes that, when effects vary over time but not across units, time-varying effects bias estimates away from their corresponding ATE. However, in a static framework, when treatment effects are homogeneous across groups in a given time period but not over time, we would also have an unbiased estimator of a relevant ATE:

$$E[\hat{\beta}^R_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{k,t'}|\mathbf{D}] \times \omega^R_{g,k,t,t'} = E[\Delta_{g,t'}|\mathbf{D}] \times \omega^R_{g,k,t,t'}$$

Thus, under a common trends assumption on the untreated potential outcomes only, reverse difference-in-differences are forbidden comparisons if treatment effects vary over time *and* across groups, and not merely over time.

$$=E[\Delta_{g,t'} + (Y_{g,t}(1) - Y_{g,t'}(1) - (Y_{k,t}(1) - Y_{k,t'}(1)))|D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 1, D_{k,t'} = 1]$$

In a framework allowing for heterogeneous treatment effects, Assumption 6 is thus both necessary and sufficient for the reverse difference-in-differences to identify the ATE of the switching group in period t'. We can write:

Assumption 6 holds if and only if $\forall g, k \neq g, t, t' \neq t, \forall \mathbf{\Omega}, E[\hat{\beta}^{R}_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t'}|\mathbf{D}] \times \omega^{R}_{g,k,t,t'}$

Overall, these comparisons should not be systematically understood as forbidden when ATEs are allowed to be heterogeneous: they require a different common trends assumption to be valid.

Leaver Difference-in-Differences, $\hat{\beta}^{L}_{g,k,t,t'}$: Consider two groups g and k, two periods t, t', such that $t > t'$, $D_{g,t} = 1$, $D_{g,t'} = 1$, $D_{k,t} = 0$ and $D_{k,t'} = 1$. We are in a case where always-treated units are used as a control group for a group which is initially treated, and leaves treatment in period t. Taking the expectation and adding and subtracting $E(Y_{k,t}(1)|D_{g,t} = 1, D_{g,t'} = 1, D_{k,t} = 0, D_{k,t'} = 1)$, we find:

$$E[\hat{\beta}^{DD}_{g,k,t,t'}|D_{g,t} = 1, D_{g,t'} = 1, D_{k,t} = 0, D_{k,t'} = 1]$$
$$=E[\Delta_{k,t} + (Y_{g,t}(1) - Y_{g,t'}(1) - (Y_{k,t}(1) - Y_{k,t'}(1)))|D_{g,t} = 1, D_{g,t'} = 1, D_{k,t} = 0, D_{k,t'} = 1]$$

That is:

Assumption 6 holds if and only if $\forall g, k \neq g, t, t' \neq t, \forall \mathbf{\Omega}, E[\hat{\beta}^{L}_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{k,t}|\mathbf{D}] \times \omega^{L}_{g,k,t,t'}$

Reverse Leaver Difference-in-Differences, $\hat{\beta}^{RL}_{g,k,t,t'}$: Consider two groups g and k, two periods t, t', such that $t > t'$, $D_{g,t} = 0$, $D_{g,t'} = 0$, $D_{k,t} = 0$ and $D_{k,t'} = 1$. In this case, never-treated units are used as a control group for a group which is initially treated, and leaves treatment in period t. Taking the expectation and adding and subtracting $E(Y_{k,t'}(0)|D_{g,t} = 0, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 1)$, we can show that this comparison identifies the ATE of group k in period t' under Assumption 5:

$$E[\hat{\beta}^{DD}_{g,k,t,t'}|D_{g,t} = 0, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 1]$$

$$=E[\Delta_{k,t'} + (Y_{g,t}(0) - Y_{g,t'}(0) - (Y_{k,t}(0) - Y_{k,t'}(0)))|D_{g,t} = 0, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 1]$$

We thus have:

Assumption 5 holds if and only if $\forall g, k \neq g, t, t' \neq t, \forall \mathbf{\Omega}, E[\hat{\beta}^{RL}_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{k,t'}|\mathbf{D}] \times \omega^{RL}_{g,k,t,t'}$

Double-Switcher Difference-in-Differences, $\hat{\beta}^{DS}_{g,k,t,t'}$: Consider two groups g and k, two periods t, t', such that $t > t'$, $D_{g,t} = 1$, $D_{g,t'} = 0$, $D_{k,t} = 0$ and $D_{k,t'} = 1$. We are in a case where we compare the outcomes of two groups which treatment status change over time in different directions. Group g is initially not treated, and becomes treated in period t. In contrast, group k is initially treated, and leaves treatment in period t. In this case, the sum of the ATE of group g in period t and of group k in period t' can be identified if the standard common trends assumption on potential outcomes when untreated holds:

$$\begin{aligned}
&E[\hat{\beta}^{DD}_{g,k,t,t'}|D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 1] \\
&= E[(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'}) \\
&+ (Y_{g,t}(0) - Y_{g,t}(0)) + (Y_{k,t'}(0) - Y_{k,t'}(0))|D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 1] \\
&= E[\Delta_{g,t} + \Delta_{k,t'}|D_{g,t} = 1, D_{g,t'} = 0, D_{k,t} = 0, D_{k,t'} = 1]
\end{aligned}$$

Overall, we thus have:[6]

Assumption 5 if and only if $\forall g, k \neq g, t, t' \neq t, \forall \mathbf{\Omega}, E[\hat{\beta}^{DS}_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t} + \Delta_{k,t'}|\mathbf{D}] \times \omega^{DS}_{g,k,t,t'}$

Note that this result is of interest in itself. It implies that, in such a setting with two groups and two periods, if one is willing to assume that ATE are homogeneous across time and across groups, a simple two-by-two difference-in-differences would allow to recover the ATT under the standard common trends assumption even in cases where we observe groups changing treatment status in opposite directions over time.[7]

[6]Note that, alternatively, we can show:

Assumption 6 holds if and only if $\forall g, k \neq g, t, t' \neq t, \forall \mathbf{\Omega}, E[\hat{\beta}^{DS}_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t'} + \Delta_{k,t}|\mathbf{D}] \times \omega^{DS}_{g,k,t,t'}$

[7]Note that assuming that ATE are homogeneous will not be necessary to show that under Assumptions

On top of this, this result implies that one can recover an additional treatment effect which may be of interest in certain settings, even under heterogeneous treatment effects. For example, consider two groups and three periods, such that $D_{1,1} = D_{1,2} = 0$, $D_{1,3} = 1$, $D_{2,1} = 0$, $D_{2,2} = 1$, and $D_{2,3} = 0$. The observations of the two first periods can be used to recover $\Delta_{2,2}$ under Assumption 5. The information contained in the last period would usually be lost. Yet, comparing the changes in outcomes of the two groups in periods 2 and 3 allows to identify the sum of $\Delta_{2,2}$ and $\Delta_{1,3}$. Subtracting the first from the second comparison would thus allow to identify $\Delta_{1,3}$.

4.4.2 A General Decomposition Result

Section 4.3 shows that the TWFE estimator is a weighted sum of five different objects. Section 4.4.1 establishes the assumptions under which each of these objects is an unbiased estimator of its corresponding ATE. Overall, we obtain the following decomposition:

Theorem 2

$$E\left[\hat{\beta}_{fe}|\mathbf{D}\right] = \frac{\sum_t \sum_g \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \sum_{t' \neq t} \sum_{k \neq g} \left[\omega^S_{g,k,t,t'} + \omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'} + 2\omega^{DS}_{k,g,t',t}\right]\right]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t}[1 - D_{g,t'} + D_{k,t'}]}$$

$\forall \mathbf{\Omega}$ *such that Assumptions 1-4 hold, if and only if, Assumptions 5 and 6 hold.*

Theorem 2 is obtained by taking the expectation of the expression of $\hat{\beta}_{fe}$ provided in Theorem 1, and plugging in each term derived in Section 4.4.1. Details of the proof are provided in Appendix 4.A.2. Theorem 2 establishes that, in a general framework where there exist groups switching on and off treatment and where treatment effects are not restricted to be homogeneous across groups nor over time, common trends conditions on both treated and untreated potential outcomes are necessary and sufficient for the TWFE estimator to always identify a convex combination of ATEs.[8] In particular, the TWFE estimator will never be negative when all ATEs are positive. Note that Assumptions 5 and 6 are also both imposed by de Chaisemartin and d'Haultfoeuille (2020a) for the estimator they propose as an alternative to the TWFE estimator to be unbiased.

5 and 6 none of the ATE that enter the TWFE estimator are weighted negatively.

[8]Note indeed that $\frac{\sum_t \sum_g \sum_{t' \neq t} \sum_{k \neq g}\left[\omega^S_{g,k,t,t'} + \omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'} + 2\omega^{DS}_{k,g,t',t}\right]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t}[1 - D_{g,t'} + D_{k,t'}]} = 1$.

The weights derived in Theorem 2 are very intuitive: they correspond to dummies equal to one each time a relevant comparison group, as defined by the five types of two-by-two comparisons in Section 4.3, exists. Each of the ATEs thus enters proportionally to the number of comparison groups available to identify it.

Note that if we consider a framework where there exist no groups entering nor leaving treatment while some remain treated, then conditions for the TWFE estimator not to weigh any ATEs negatively are weaker. In particular, only the common trends assumption on the untreated potential outcomes must hold. This however rules out the staggered difference-in-differences framework. When, in contrast, there exist no groups entering nor leaving treatment while some remain untreated, the common trends assumption on the untreated potential outcomes can be violated as long as its counterpart for treated potential outcomes holds.

4.4.3 Extension: Introducing Dynamic Treatment Effects

The above decomposition is derived in the case where treatment can only have a contemporaneous effect on the outcome. Let us now introduce the dynamic potential outcome framework, following Robins (1986) and de Chaisemartin and D'Haultfoeuille (2022a). In particular, we denote $Y_{g,t}(\mathbf{d})$ the average potential outcome for group g in period t when facing the treatment sequence $\mathbf{d} \in \{0,1\}^T$. The realized outcome of group g writes $Y_{g,t} = Y_{g,t}(D_{g,1}, D_{g,2}, ..., D_{g,T}) = Y_{g,t}(\mathbf{D}_{g,T})$, where $\mathbf{D}_{g,t}$ is the vector of treatment status for group g up to period t. Now, $\mathbf{\Omega}$ is the vector $\{(Y_{g,t}(\mathbf{d}))_{d\in\{0,1\}^T}, \mathbf{D}_{g,t}\}_{\forall g,t}$. We only maintain Assumptions 1 and 2. Let us also follow de Chaisemartin and D'Haultfoeuille (2022a) and impose that a group's current outcome does not depend on its future treatment status, as well as a common trends assumption on the never-treated potential outcomes:

Assumption 7 *For all g, for all $\mathbf{d} \in \{0,1\}^T$, $Y_{g,t}(\mathbf{d}) = Y_{g,t}(d_1, ..., d_t)$.*

Assumption 8 $\forall t \geq 2, \forall g, E(Y_{g,t}(\mathbf{0}_t) - Y_{g,t-1}(\mathbf{0}_{t-1})|\mathbf{D})$*, where $\mathbf{0}_t$ corresponds to a vector of zeros of length t, does not vary across g.*

Potential outcomes can now depend on past treatments: $Y_{g,t}(\mathbf{D}_{t-1}, d_t)$ may be different from $Y_{g,t}(\mathbf{D}'_{t-1}, d_t)$ where $\mathbf{D}_{t-1} \neq \mathbf{D}'_{t-1}$. Can the static TWFE estimator (Equation (4.1)) be robust to dynamic treatment effects? I now use the decomposition provided in Theorem 1[9] to show that it is the case in some relevant scenarios. Let us assume that untreated potential outcomes do not depend on the treatment history of a group, that is:

Assumption 9 *(No Dynamics for Untreated Potential Outcomes):*

$$Y_{g,t}(\mathbf{D}_{t-1}, 0) = Y_{g,t}(\mathbf{D}^{'}_{t-1}, 0), \ \forall \mathbf{D}_{t-1}, \mathbf{D}^{'}_{t-1} \in \{0,1\}^{t-1}$$

A General Decomposition Result in the Dynamic Framework Given Assumptions 8 and 9, the standard, reverse leaver and double-switcher difference-in-differences remain unbiased estimators of their corresponding ATEs even in the presence of dynamics. Let us now consider the reverse difference-in-differences, where we compare group g switching from untreated to treated between period t and t', to group k remaining treated across the two periods. Adding and subtracting $E[Y_{g,t'}(\mathbf{D}_{g,t'-1}, 1)|\mathbf{D}]$, we have:

$$E[\hat{\beta}^{DD}_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t'}(\mathbf{D}_{g,t'-1}) + Y_{g,t}(\mathbf{D}_{g,t-1}, 1) - Y_{g,t'}(\mathbf{D}_{g,t'-1}, 1) - (Y_{k,t}(\mathbf{D}_{k,t-1}, 1) - Y_{k,t'}(\mathbf{D}_{k,t'-1}, 1))|\mathbf{D}]$$

where $\Delta_{g,t'}(\mathbf{D}_{g,t'-1}) = Y_{g,t'}(\mathbf{D}_{g,t'-1}, 1) - Y_{g,t'}(\mathbf{D}_{g,t'-1}, 0)$, the treatment effect of group g in time t' conditional on having faced treatment history $\mathbf{D}_{g,t'-1}$. A necessary and sufficient condition for the reverse two-by-two comparisons included in the TWFE estimator is thus that, conditional on the treatment history of each group, their treated potential outcomes would follow the same trend. The same assumption is required for the leaver difference-in-differences. The counterpart of Theorem 2 in this framework is thus as follows:

Theorem 3

$$E\left[\hat{\beta}_{fe}|\mathbf{D}\right] = \frac{\sum_t \sum_g \left[E\left[\Delta_{g,t}(\mathbf{D}_{g,t-1})|\mathbf{D}\right] \sum_{t' \neq t} \sum_{k \neq g} \left[\omega^S_{g,k,t,t'} + \omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'} + 2\omega^{DS}_{k,g,t',t}\right]\right]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t} [1 - D_{g,t'} + D_{k,t'}]}$$

$\forall \mathbf{\Omega}$ such that Assumptions 1-2, 7-9 hold, if and only if, $\forall t > 2$, $E[Y_{g,t}(\mathbf{D}_{g,t-1}, 1) - Y_{g,t-1}(\mathbf{D}_{g,t-2}, 1)|\mathbf{D}]$ does not vary across g.

The proof follows the same argument as for Theorem 2. Note that, together with Assumption 6, the absence of dynamic treatment effects is thus a sufficient condition for the TWFE

[9] Note indeed that this decomposition does not rely on any assumptions about the true data generating process for $Y_{g,t}$, and thus remains valid in the presence of dynamic treatment effects.

estimator not to weigh any ATEs negatively. However, it is not necessary, and I now show that the condition in Theorem 3 can be satisfied even in the presence of dynamic treatment effects.

Sufficient Conditions in the Presence of Dynamic Treatment Effects Let us now derive sufficient conditions for the TWFE estimator to be heterogeneity-robust even in the presence of dynamics. I impose the following assumption:

Assumption 10

1. *(Common Trends on the Potential Outcome of First Exposure to Treatment:) For $t \geq 2$, $E(Y_{g,t}(\mathbf{0}_{t-1}, 1) - Y_{g,t-1}(\mathbf{0}_{t-2}, 1)|\mathbf{D})$ does not vary across g.*

2. *(Homogeneous Effect from n-period Treatment Exposure:) $\forall g, t$, the ATE of group g in period t conditional on history $D_{g,t-1}$ writes: $\Delta_{g,t}(D_{g,t-1}) = \Delta_{g,t}(\mathbf{0}_{t-1}) + \tau_t(\sum_{\ell=1}^{t-1} D_{g,\ell})$, with $\tau_t(0) = 0$.*

Assumption 10.1. is a parallel trends assumption on *treated* potential outcomes, conditional on not having been treated before. This simply means that the difference in outcomes when being treated for the first time in period t or in period $t-1$ would be the same across groups.

Assumption 10.2. defines the ATE of group g in period t as being equal to the sum of the ATE of the group if it was first exposed to treatment in t, and the effect of having been exposed to treatment for $\sum_{\ell=1}^{t-1} D_{g,\ell}$ periods prior to t, $\tau_t(\sum_{\ell=1}^{t-1} D_{g,\ell})$. The latter is a period-t specific function, taking as argument the group's number of periods of exposure to treatment prior to t. This implies that the incremental effect of having been exposed to treatment for n periods is homogeneous across groups for a given t. Note that this assumption does not rule out treatment effect heterogeneity across groups, nor over time. Indeed, the first-exposure ATE is g, t-specific. Further, the incremental effect $\tau_t(n)$ is also allowed to vary across time for a given n.

Overall, Assumptions 8 to 10 imply that $E[Y_{g,t}(\mathbf{D}_{k,t-1}, 1) - Y_{g,t'}(\mathbf{D}_{k,t'-1}, 1) - (Y_{k,t}(\mathbf{D}_{k,t-1}, 1) - Y_{k,t'}(\mathbf{D}_{k,t'-1}, 1)|\mathbf{D}]$ is equal to zero.[10] Under these assumptions, we can derive the following theorem:

Theorem 4 *Suppose Assumptions 1-2, 7-10 hold. If:*

[10] Details can be found in Appendix 4.A.2.

1. $\nexists g, k \neq g, t, t' \neq t$ *such that* $\omega^R_{g,k,t',t} = 1$*, or if* $g, k \neq g, t, t' \neq t$ *for which* $\omega^R_{g,k,t',t} = 1$ *are such that* $\sum_{\ell=1}^{t'-1} D_{g,\ell} = \sum_{\ell=1}^{t'-1} D_{k,\ell}$ *and,*

2. $\nexists g, k \neq g, t, t' \neq t$ *such that* $\omega^L_{k,g,t,t'} = 1$*, or if* $g, k \neq g, t, t' \neq t$ *for which* $\omega^L_{k,g,t,t'} = 1$ *are such that* $\sum_{\ell=1}^{t'-1} D_{g,\ell} = \sum_{\ell=1}^{t'-1} D_{k,\ell}$*,*

then,

$$E\left[\hat{\beta}_{fe}|\mathbf{D}\right] = \frac{1}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t} [1 - D_{g,t'} + D_{k,t'}]} \times$$
$$\left[\sum_t \sum_g \left[E\left[\Delta_{g,t}(D_{g,t-1})|\mathbf{D}\right] \sum_{t' \neq t} \sum_{k \neq g} \left[\omega^S_{g,k,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'} + 2\omega^{DS}_{k,g,t',t}\right]\right] \right.$$
$$\left. + \sum_t \sum_g \left[E\left[\Delta_{g,t}(D_{k,t-1})|\mathbf{D}\right] \sum_{t' \neq t} \sum_{k \neq g} \left[\omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'}\right]\right]\right]$$

Details can be found in Appendix 4.A.2. This result stems from the fact that, first, reverse or leaver difference-in-differences between t' and t when both groups have been exposed to treatment the same number of periods up to $t'-1$ and up to period $t-1$ identifies the ATE of the group switching treatment status under its own history. This would for example happen when comparing periods 2 and 4 of group g with $\mathbf{D}_{g,t} = (0, 0, 1, 1)$ and group k with $\mathbf{D}_{k,t} = (0, 1, 0, 1)$. Second, it relies on the fact that when the number of periods under treatment prior to the treated period of the switching group is the same as in the comparison group, reverse and leaver difference-in-differences identify the ATE of the switching group under the history of the comparison group. For example, comparing a group g such that $\mathbf{D}_{g,t} = (1, 0, 0, 1)$ to a group k with $\mathbf{D}_{k,t} = (0, 0, 1, 1)$ in periods 3 and 4 allows to identify the ATE of group g in period 3 when being exposed to treatment for the first time.

Overall, the absence of dynamic treatment effects over time is a sufficient condition for the TWFE estimator to be heterogeneity-robust. This may be a plausible assumption when individuals stay within a group g only one time period. For example, when evaluating educational policies, the treatment unit is often a grade within a school: students within this group will be treated a single period, limiting the potential for dynamic treatment effects.

However, I also show that the TWFE estimator remains valid in relevant settings, even in the presence of dynamic treatment effects. The framework provided should help researchers evaluate whether their setting contains comparisons threatening the validity of the static TWFE estimator. Under the dynamics described in Assumption 10 and when reverse and leaver difference-in-differences exist, the potential for the TWFE estimator to weigh biased estimators of ATEs is null under two conditions. First, the existing reverse difference-in-differences should be such that the cumulative number of treated periods is the same in the two groups prior to the last period of the comparison. Second, the existing leaver difference-in-differences should be such that the cumulative number of treated periods is the same in the two groups prior to the first period of the comparison.

4.5 Implications for Empirical Analyses

I now derive the implications of Section 4.4.2 for researchers. First, when using the TWFE estimator, researchers should test for both *untreated* and *treated* parallel trends. Second, I provide alternative estimators exploiting the two-by-two comparisons highlighted above.

4.5.1 Testing for Common Trends on *Treated* Potential Outcomes

To the extent that the static TWFE estimator remains broadly used in empirical studies, it is crucial to establish under which assumptions all ATEs it includes are weighted positively. Suppose the researcher is in a framework where dynamic treatment effects can safely be discarded. As shown in Section 4.4.2, the TWFE estimator will then be heterogeneity-robust under two common trends assumption, Assumptions 5 and 6.

Overall, in a setting with several groups and time periods, when using the TWFE estimator, one should test not only for common trends on the potential outcomes when untreated, but also when treated.[11] While Assumption 6 is not directly testable, it has a natural testable counterpart, suggested by Kim and Lee (2019) who study the reverse difference-in-differences. In particular, one can use groups which stay treated over at least two periods in order to compare the evolution of their outcomes. In particular, one should test whether the following relation holds:

$$E\left(Y_{g,t}-Y_{g,t'} \mid D_{g,t}=1, D_{g,t'}=1\right)=E\left(Y_{k,t}-Y_{k,t'} \mid D_{k,t}=1, D_{k,t'}=1\right)$$

[11] Note that one should perform a similar test when using the heterogeneity-robust estimator provided by de Chaisemartin and d'Haultfoeuille (2020a).

4.5.2 Alternative Estimators

Even if non-negative, the weights derived in Section 4.4.2 do not correspond to each group's sample size. This might make the TWFE estimator difficult to interpret, and implies that it is generally a biased estimator of δ^{TR}. This may motivate the use of other heterogeneity-robust estimators, such as the one suggested by de Chaisemartin and d'Haultfoeuille (2020a), which estimates a quantity which may be of interest: the ATE of switching cells.[12] This paper highlights additional comparisons which can be exploited to construct alternative estimators. I first show that one can build an augmented, unbiased estimator of the ATE of switching cells, while relaxing some of the assumptions de Chaisemartin and d'Haultfoeuille (2020a) impose. Second, I provide an unbiased estimator of the ATT, δ^{TR}, and of the ATE.

Unbiased Estimator of the ATE of Switching Cells

Let us consider the following object, the ATE of all switching cells:

$$\delta^S = E\left[\frac{1}{N_S}\sum_{(i,g,t):t\geq 2, D_{g,t}\neq D_{g,t-1}}[Y_{i,g,t}(1) - Y_{i,g,t}(0)]\right]$$

where $N_S = \sum_{(g,t):t\geq 2, D_{g,t}\neq D_{g,t-1}} N_{g,t}$.

de Chaisemartin and d'Haultfoeuille (2020a) provide an unbiased estimator of this object, under assumptions ensuring the existence of relevant comparison groups. I now show that an unbiased estimator of δ^S can be built while relaxing these requirements. This estimator exploits the additional comparisons highlighted above.

Assumption 11 *(Mean Independence between a Group's Outcome and Other Group Treatments): For all g and t, $E[Y_{g,t}(0)|\mathbf{D}] = E[Y_{g,t}(0)|\mathbf{D}_g]$ and $E[Y_{g,t}(1)|\mathbf{D}] = E[Y_{g,t}(1)|\mathbf{D}_g]$.*

Assumption 12 *(Existence of Stable Groups): For all $t \geq 2$:*

(i) If there is at least one $g \in \{1, ..., G\}$ such that $D_{g,t-1} = 0$, $D_{g,t} = 1$, then 1) there exists at least one $g' \neq g$, $g' \in \{1, ..., G\}$ such that either $D_{g',t-1} = D_{g',t} = 0$ or 2) g is such that $D_{g,t+1} = 0$ and there exists at least one g' such that $D_{g',t} = D_{g',t+1} = 0$.

[12] Note, however, that it is not an unbiased estimator of the ATT.

(ii) If there is at least one $g \in \{1, ..., G\}$ such that $D_{g,t-1} = 1$, $D_{g,t} = 0$, then 1) there exists at least one $g' \neq g$, $g' \in \{1, ..., G\}$ such that either $D_{g',t-1} = D_{g',t} = 1$ or 2) g is such that $D_{g,t+1} = 1$ and there exists at least one g' such that $D_{g',t} = D_{g',t+1} = 1$.

I follow de Chaisemartin and d'Haultfoeuille (2020a) in imposing Assumption 11. However, I relax the assumptions they impose with respect to the existence of 'stable groups', which correspond to Assumption 12(i)1) and Assumption 12(ii)1). They require the existence of a group which stays untreated (treated, respectively) over two consecutive periods if there exists a group which joins (leaves, respectively) treatment over these periods. When such assumptions hold, one can identify the ATE in period t for each group switching treatment status between $t-1$ and t, and hence δ^S.

Yet, if these assumptions fail to hold, I now show that one can exploit other comparisons to identify the ATE in period t for each group switching treatment status between $t-1$ and t. In particular, while the estimator suggested by de Chaisemartin and d'Haultfoeuille (2020a) comprises only two kinds of comparisons, at least four two-by-two comparisons could be included.[13] First, one could use the two-by-two comparisons initially thought of as 'forbidden comparisons', the reverse difference-in-differences. Second, one could include the reverse leaver difference-in-differences.

Assumptions 12(i)2) and 12(ii)2) define the alternative stable groups which are needed to derive an unbiased estimator of δ^S when Assumptions 12(i)1) or 12(ii)1) are not satisfied. Let us now define the augmented estimator. For all $t \in \{2, ..., T\}$ and for all $(d, d') \in \{0, 1\}^2$, let

$$N_{d,d',t} = \sum_{g:D_{g,t}=d,D_{g,t-1}=d'} N_{g,t}$$

denote the number of observations with treatment d' at period $t-1$ and d at period t. The following objects will be included in the augmented estimator:

$$DID_{+,g,t} = \mathbb{1}\{D_{g,t} = 1, D_{g,t-1} = 0\} \left[(Y_{g,t} - Y_{g,t-1}) - \sum_{k:D_{k,t}=D_{k,t-1}=0} \frac{N_{k,t}}{N_{0,0,t}} (Y_{k,t} - Y_{k,t-1}) \right]$$

[13] For simplicity, I do not consider using the double-switcher difference-in-differences comparison.

$$DID_{-,g,t} = \mathbb{1}\{D_{g,t} = 0, D_{g,t-1} = 1\} \left[\sum_{k:D_{k,t}=1,D_{k,t-1}=1} \frac{N_{k,t}}{N_{1,1,t}} (Y_{k,t} - Y_{k,t-1}) - (Y_{g,t} - Y_{g,t-1}) \right]$$

$$DID^R_{+,g,t} = \mathbb{1}\{D_{g,t} = 1, D_{g,t-1} = 0\} \left[(Y_{g,t} - Y_{g,t-1}) - \sum_{k:D_{k,t}=D_{k,t-1}=1} \frac{N_{k,t}}{N_{1,1,t}} (Y_{k,t} - Y_{k,t-1}) \right]$$

$$DID^R_{-,g,t} = \mathbb{1}\{D_{g,t} = 0, D_{g,t-1} = 1\} \left[\sum_{k:D_{k,t}=0,D_{k,t-1}=0} \frac{N_{k,t}}{N_{0,0,t}} (Y_{k,t} - Y_{k,t-1}) - (Y_{g,t} - Y_{g,t-1}) \right]$$

$DID_{+,g,t}$ and $DID_{-,g,t}$ are defined in a similar fashion as in de Chaisemartin and d'Haultfoeuille (2020a). Following them, we let $DID_{+,g,t} = 0$ if there is no group such that $D_{g,t} = 1$ and $D_{g,t-1} = 0$ or no group such that $D_{g,t} = D_{g,t-1} = 0$. Similarly, we let $DID_{-,g,t} = 0$ if there is no group such that $D_{g,t} = 0$ and $D_{g,t-1} = 1$ or no group such that $D_{g,t} = D_{g,t-1} = 1$. We follow the same rule for the two additional objects. We let $DID^R_{+,g,t} = 0$ if there is no group such that $D_{g,t} = 1$ and $D_{g,t-1} = 0$ or no group such that $D_{g,t} = D_{g,t-1} = 1$. Similarly, we let $DID_{-,g,t} = 0$ if there is no group such that $D_{g,t} = 0$ and $D_{g,t-1} = 1$ or no group such that $D_{g,t} = D_{g,t-1} = 0$.

Finally, let us define the augmented estimator of δ^S:

$$\begin{aligned} DID^A_M = \sum_{t=2}^{T} \Bigg[\frac{1}{N_S} \sum_{g:D_{g,t}=1,D_{g,t-1}=0} & N_{g,t}(\mathbb{1}\{\exists k \neq g, D_{k,t} = D_{k,t-1} = 0\} DID_{+,g,t} \\ & +.\mathbb{1}\{\nexists k \neq g, D_{k,t} = D_{k,t-1} = 0\} DID^R_{-,g,t+1}) \\ + \frac{1}{N_S} \sum_{g:D_{g,t}=0,D_{g,t-1}=1} & N_{g,t}(\mathbb{1}\{\exists k \neq g, D_{k,t} = D_{k,t-1} = 1\} DID_{-,g,t} \\ & +\mathbb{1}\{\nexists k \neq g, D_{k,t} = D_{k,t-1} = 1\} DID^R_{+,g,t+1}) \Bigg] \end{aligned}$$

Theorem 5 *If Assumption 1, 2, 4-6, 11-12 hold, then* $E[DID^A_M] = \delta^S$.

The proof, following closely the steps of de Chaisemartin and d'Haultfoeuille (2020a), is provided in Appendix 4.A.3. This estimator uses the same comparisons as the one suggested by de Chaisemartin and d'Haultfoeuille (2020a). However, when a given comparison is equal to zero due to the absence of a stable group, it augments it by using either the reverse or reverse leaver difference-in-differences. Note that these two comparisons contrast outcomes between $t+1$ and t, as they identify the ATE of the switching group in the period before it switches treatment status.

Unbiased Estimators of the ATT and ATE

Each two-by-two comparison highlighted in Section 4.3 is an unbiased estimator of the ATE of the group switching treatment status. Yet, most of them receive a weight equal to zero in the heterogeneity-robust estimator provided by de Chaisemartin and d'Haultfoeuille (2020a). Moreover, the latter is not an unbiased estimator of the ATT, δ^{TR}. I now show that one can exploit the comparisons identifying the ATE of the switching group when it is treated, the standard and reverse leaver difference-in-differences, to build an estimator of δ^{TR}.

Assumption 13 *(Existence of Stable Groups): For all g, t such that $D_{g,t} = 1$, there exists at least one group k and time period t' such that $\omega^{S}_{g,k,t,t'} = 1$ or $\omega^{RL}_{k,g,t',t} = 1$.*

Assumption 13 simply specifies that, for each group g treated in period t, there exists at least one group and time period such that a comparison allowing to identify its ATE can be performed. We can then re-weigh appropriately the existing standard and reverse leaver difference-in-differences identifying the ATE corresponding to group g in period t when $D_{g,t} = 1$ in order to build an estimator of δ^{TR}. We obtain the following estimator:

$$DID^{TR} = \frac{1}{N^{(1)}} \sum_{g,t:D_{g,t}=1} N_{g,t} \left[\frac{\sum_{k\neq g} \sum_{t'\neq t} \left[\hat{\beta}^{S}_{g,k,t,t'} + \hat{\beta}^{RL}_{k,g,t',t}\right]}{\sum_{k\neq g} \sum_{t'\neq t} \left[\omega^{S}_{g,k,t,t'} + \omega^{RL}_{k,g,t',t}\right]} \right]$$

Theorem 6 *If Assumption 1, 2, 4-6, 11 and 13 hold, then $E[DID^{TR}] = \delta^{TR}$.*

The proof is immediate, as Section 4.4.1 establishes that each of the two-by-two comparisons incorporated in DID^{TR} identifies the ATE of group g in t. It thus suffices to re-weigh each available comparison appropriately to form an unbiased estimator of the ATT.[14]

[14]Note that the comparisons included in this estimator are also unbiased in the dynamic framework considered in Section 4.4.3.

Finally, one could additionally use the reverse and leaver difference-in-differences which identify the ATE of the group switching status at the time where it is untreated in order to identify the ATE, δ^T:

$$\delta^T = E\left[\frac{1}{N}\sum_{(i,g,t)}[Y_{i,g,t}(1) - Y_{i,g,t}(0)]\right]$$

Building an unbiased estimator of this object requires to assume that relevant comparisons can be performed.

Assumption 14 *(Existence of Stable Groups): For all g, t, there exists at least one group k and time period t' such that $\omega^S_{g,k,t,t'} = 1$ or $\omega^R_{g,k,t',t} = 1$ or $\omega^{RL}_{k,g,t',t} = 1$ or $\omega^L_{k,g,t,t'} = 1$.*

We can then define:

$$DID^T = \frac{1}{N}\sum_{g,t} N_{g,t}\left[\frac{\sum_{k\neq g}\sum_{t'\neq t}\left[\hat{\beta}^S_{g,k,t,t'} + \hat{\beta}^{RL}_{k,g,t',t} + \hat{\beta}^L_{k,g,t,t'} + \hat{\beta}^R_{g,k,t',t}\right]}{\sum_{k\neq g}\sum_{t'\neq t}\left[\omega^S_{g,k,t,t'} + \omega^{RL}_{k,g,t',t} + \omega^L_{k,g,t,t'} + \omega^R_{g,k,t',t}\right]}\right]$$

where N is the total number of observations in the sample.

Theorem 7 *If Assumption 1, 2, 4-6, 11 and 14 hold, then $E[DID^T] = \delta^T$.*

Again, the proof is immediate and stems from the fact that each of the comparisons included in DID^T identifies the ATE of group g in period t, irrespective of its treatment status.[15]

4.6 Conclusion

Difference-in-differences is one of the most popular quasi-experimental methods to estimate causal effects. Most empirical applications have yet departed from the traditional two-group two-period setting, for which it is established that comparing the evolution of the outcome of interest before and after a group receives a treatment to the one of a never-treated group identifies the ATT. With several groups and periods, researchers will typically interpret the parameter of the treatment dummy in a TWFE regression as the ATT. Yet, recent developments in the difference-in-differences literature have concluded that, under the standard

[15]Note that I derive this estimator within a static framework. It could be easily adapted to the framework of Section 4.4.3 if one is interested in the sample ATE even in the presence of dynamic treatment effects. It would require to consider either frameworks in which all comparisons are unbiased estimators of their corresponding ATE, or by including only valid comparisons in the estimator.

common trends assumption, the TWFE estimator may weigh negatively some ATEs in the presence of heterogeneous treatment effects.

This paper provides necessary and sufficient assumptions for the TWFE estimator to weigh all the ATEs it includes positively. When only contemporaneous treatment effects are considered, I show that it requires a common trends assumptions on both the *treated* and *untreated* potential outcomes. I further consider the presence of dynamic treatment effects, and show that the TWFE estimator remains valid under plausible conditions.

To derive this result, I decompose the TWFE estimator and show that it may include five different types of standard two-by-two comparisons, all entering positively. I then study these comparisons separately and find that each is an unbiased estimator of the ATE of the group switching treatment status under either a common trends assumption on potential outcomes when treated or when untreated. Under these assumptions, I show that the TWFE estimator weighs each ATE proportionally to the number of comparison groups available to identify it. Finally, I show how to combine the highlighted comparisons in order to construct unbiased estimators of the ATT and ATE. I also use them to build an unbiased estimator of the ATE of all switching cells under less stringent assumptions than de Chaisemartin and d'Haultfoeuille (2020a).

As noted by de Chaisemartin and D'Haultfoeuille (2022b), 'understanding the circumstances where TWFE and heterogeneity-robust difference-in-differences estimators are more likely to differ is an important question'. Results derived above are key to understand why the TWFE and heterogeneity-robust difference-in-differences estimators may be very similar in practice. The valid comparisons highlighted in the paper may also open the way to developing heterogeneity-robust estimators exploiting the variation present in the data in a more comprehensive manner.

Appendices

4.A Proof

4.A.1 Proof of Theorem 1

We take as a point of departure the decomposition of Strezhnev (2018):

$$\hat{\beta}_{fe} = \frac{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t} [(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t} [1 - D_{g,t'} + D_{k,t'}]} \tag{4.8}$$

Let us focus on the numerator:

$$\begin{aligned}
&\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t} [(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})] \\
=& \underbrace{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' < t} [(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})]}_{A} - \underbrace{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' > t} [(Y_{g,t'} - Y_{g,t}) - (Y_{k,t'} - Y_{k,t})]}_{B}
\end{aligned}$$

Let us start by decomposing A:

$$\begin{aligned}
A = \sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' < t} \Bigg[& \sum_{g:D_{g,t'=1}} \sum_{k:D_{k,t'=1}} \left[(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})\right] \\
& + \sum_{g:D_{g,t'=1}} \sum_{k:D_{k,t'=0}} \left[(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})\right] \\
& + \sum_{g:D_{g,t'=0}} \sum_{k:D_{k,t'=0}} \left[(Y_{g,t} - Y_{g,t'}) - (Y_{k,t} - Y_{k,t'})\right]
\end{aligned}$$

$$+\sum_{g:D_{g,t'}=0}\sum_{k:D_{k,t'}=1}\big[(Y_{g,t}-Y_{g,t'})-(Y_{k,t}-Y_{k,t'})\big]\Bigg]$$

Similarly, we can decompose B:

$$\begin{aligned}B=\sum_t\sum_{g:D_{g,t}=1}\sum_{k:D_{k,t}=0}\sum_{t'>t}\Bigg[&\sum_{g:D_{g,t'}=1}\sum_{k:D_{k,t'}=1}\big[(Y_{g,t'}-Y_{g,t})-(Y_{k,t'}-Y_{k,t})\big]\\&+\sum_{g:D_{g,t'}=1}\sum_{k:D_{k,t'}=0}\big[(Y_{g,t'}-Y_{g,t})-(Y_{k,t'}-Y_{k,t})\big]\\&+\sum_{g:D_{g,t'}=0}\sum_{k:D_{k,t'}=0}\big[(Y_{g,t'}-Y_{g,t})-(Y_{k,t'}-Y_{k,t})\big]\\&+\sum_{g:D_{g,t'}=0}\sum_{k:D_{k,t'}=1}\big[(Y_{g,t'}-Y_{g,t})-(Y_{k,t'}-Y_{k,t})\big]\Bigg]\end{aligned}$$

The second term of A and B are the same, they will thus disappear when computing A-B. We thus have, using the notations defined in Section 4.3:

$$\begin{aligned}&\sum_t\sum_{g:D_{g,t}=1}\sum_{k:D_{k,t}=0}\sum_{t'\neq t}[(Y_{g,t}-Y_{g,t'})-(Y_{k,t}-Y_{k,t'})]\\=&\sum_t\sum_{t'\neq t}\sum_g\sum_{k\neq g}[\hat{\beta}^S_{g,k,t,t'}+\hat{\beta}^R_{k,g,t',t}+\hat{\beta}^L_{g,k,t,t'}+\hat{\beta}^{RL}_{k,g,t',t}+2\hat{\beta}^{DS}_{g,k,t,t'}]\end{aligned}$$

Hence, we can write:

$$\hat{\beta}_{fe}=\frac{\sum_t\sum_{t'\neq t}\sum_g\sum_{k\neq g}[\hat{\beta}^S_{g,k,t,t'}+\hat{\beta}^R_{k,g,t',t}+\hat{\beta}^L_{g,k,t,t'}+\hat{\beta}^{RL}_{k,g,t',t}+2\hat{\beta}^{DS}_{g,k,t,t'}]}{\sum_t\sum_{g:D_{g,t}=1}\sum_{k:D_{k,t}=0}\sum_{t'\neq t}[1-D_{g,t'}+D_{k,t'}]}\tag{4.9}$$

4.A.2 Section 4.4: Proofs

The Five Difference-in-Differences Comparisons: Identification

Let us focus on the standard difference-in-differences. We want to prove that Assumption 5 holds if and only if, $\forall g,k\neq g,t,t'\neq t,\forall\mathbf{\Omega},E[\hat{\beta}^S_{g,k,t,t'}|\mathbf{D}]=E[\Delta_{g,t}|\mathbf{D}]\times\omega^S_{g,k,t,t'}$.

First, it is well known that Assumption 5 is a sufficient condition for the standard difference-in-differences to identify the ATE of the group joining treatment. Let us now show that when Assumption 5 does not hold, it is not true that $E[\hat{\beta}^S_{g,k,t,t'}|\mathbf{D}] = E[\Delta_{g,t}|\mathbf{D}] \times \omega^S_{g,k,t,t'}$, $\forall g, k \neq g, t, t' \neq t$ and $\forall \mathbf{\Omega}$.

Let us assume that Assumption 5 does not hold. Then, $\exists g, k \neq g, t, t' \neq t$ such that $E[Y_{g,t}(0) - Y_{g,t'}(0)|\mathbf{D}] \neq E[Y_{k,t}(0) - Y_{k,t'}(0)|\mathbf{D}]$. Let us consider $\mathbf{\Omega}$ such that $\omega_{g,k,t,t'} = 1$. Then, we have:

$$
\begin{aligned}
E[\hat{\beta}^S_{g,k,t,t'}|\mathbf{D}] &= E[\Delta_{g,t} + (Y_{g,t}(0) - Y_{g,t'}(0) - (Y_{k,t}(0) - Y_{k,t'}(0)))|\mathbf{D}] \times \omega^S_{g,k,t,t'} \\
&\neq E[\Delta_{g,t}|\mathbf{D}] \times \omega^S_{g,k,t,t'}
\end{aligned}
$$

The proofs corresponding to the other two-by-two comparisons follow the same argument.

Proof of Theorem 2

Let us impose Assumptions 1-6. Taking the expectation of $\hat{\beta}_{fe}$ conditional on $\mathbf{D}$, we have:

$$
\begin{aligned}
E\left[\hat{\beta}_{fe}|\mathbf{D}\right] &= \frac{\sum_t \sum_{t' \neq t} \sum_g \sum_{k \neq g} E\left[\hat{\beta}^S_{g,k,t,t'} + \hat{\beta}^R_{k,g,t',t} + \hat{\beta}^L_{g,k,t,t'} + \hat{\beta}^{RL}_{k,g,t',t} + 2\hat{\beta}^{DS}_{g,k,t,t'}|\mathbf{D}\right]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t}[1 - D_{g,t'} + D_{k,t'}]} \\
&= \frac{\sum_t \sum_{t' \neq t} \sum_g \sum_{k \neq g} \left[\beta^S_{g,k,t,t'} + \beta^R_{k,g,t',t} + \beta^L_{g,k,t,t'} + \beta^{RL}_{k,g,t',t} + 2\beta^{DS}_{g,k,t,t'}\right]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t' \neq t}[1 - D_{g,t'} + D_{k,t'}]}
\end{aligned}
$$

where $\beta^C_{g,k,t,t'} \equiv E\left[\hat{\beta}^C_{g,k,t,t'}|\mathbf{D}\right]$,

We can rewrite the numerator:

$$
\begin{aligned}
&\sum_t \sum_{t' \neq t} \sum_g \sum_{k \neq g} \left[\beta^S_{g,k,t,t'} + \beta^R_{k,g,t',t} + \beta^L_{g,k,t,t'} + \beta^{RL}_{k,g,t',t} + 2\beta^{DS}_{g,k,t,t'}\right] \\
=&\sum_t \sum_{t' \neq t} \sum_g \sum_{k \neq g} \Big[E\big[\Delta_{g,t}|\mathbf{D}\big] \times \big[\omega^S_{g,k,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'}\big] \\
&+ E\big[\Delta_{k,t}|\mathbf{D}\big] \times \big[\omega^R_{k,g,t',t} + \omega^L_{g,k,t,t'}\big] + E\big[\Delta_{k,t'}|\mathbf{D}\big] \times \big[2\omega^{DS}_{g,k,t,t'}\big]\Big]
\end{aligned}
$$

The fist term of the sum rewrites:

$$
\sum_t \sum_{t' \neq t} \sum_g \sum_{k \neq g} \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \times \left[\omega^S_{g,k,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'}\right]\right]
$$

$$= \sum_t \sum_g \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \sum_{t'\neq t} \sum_{k\neq g} \left[\omega^S_{g,k,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'}\right] \right]$$

Let us focus on the term $\sum_t \sum_{t'\neq t} \sum_g \sum_{k\neq g} \left[E\left[\Delta_{k,t}|\mathbf{D}\right] \times \left[\omega^R_{k,g,t',t} + \omega^L_{g,k,t,t'}\right]\right]$:

$$\sum_t \sum_{t'\neq t} \sum_g \sum_{k\neq g} \left[E\left[\Delta_{k,t}|\mathbf{D}\right] \times \left[\omega^R_{k,g,t',t} + \omega^L_{g,k,t,t'}\right]\right]$$
$$= \sum_t \sum_g \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \times \sum_{t'\neq t} \sum_{k\neq g} \left[\omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'}\right] \right]$$

And, focusing on the term $\sum_t \sum_{t'\neq t} \sum_g \sum_{k\neq g} \left[E\left[\Delta_{k,t'}|\mathbf{D}\right] \times 2\omega^{DS}_{g,k,t,t'}\right]$:

$$\sum_t \sum_{t'\neq t} \sum_g \sum_{k\neq g} \left[E\left[\Delta_{k,t'}|\mathbf{D}\right] \times 2\omega^{DS}_{g,k,t,t'}\right]$$
$$= \sum_t \sum_g \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \times \sum_{t'\neq t} \sum_{k\neq g} 2\omega^{DS}_{k,g,t',t} \right]$$

The numerator thus writes:

$$\sum_t \sum_{t'\neq t} \sum_g \sum_{k\neq g} \left[\beta^S_{g,k,t,t'} + \beta^R_{k,g,t',t} + \beta^L_{g,k,t,t'} + \beta^{RL}_{k,g,t',t} + 2\beta^{DS}_{g,k,t,t'}\right]$$
$$= \sum_t \sum_g \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \sum_{t'\neq t} \sum_{k\neq g} \left[\omega^S_{g,k,t,t'} + \omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'} + 2\omega^{DS}_{k,g,t',t}\right] \right]$$

We thus have, $\forall \mathbf{\Omega}$:

$$E\left[\hat{\beta}_{fe}|\mathbf{D}\right] = \frac{\sum_t \sum_g \left[E\left[\Delta_{g,t}|\mathbf{D}\right] \sum_{t'\neq t} \sum_{k\neq g} \left[\omega^S_{g,k,t,t'} + \omega^R_{g,k,t',t} + \omega^L_{k,g,t,t'} + \omega^{RL}_{k,g,t',t} + 2\omega^{DS}_{g,k,t,t'} + 2\omega^{DS}_{k,g,t',t}\right] \right]}{\sum_t \sum_{g:D_{g,t}=1} \sum_{k:D_{k,t}=0} \sum_{t'\neq t}[1 - D_{g,t'} + D_{k,t'}]}$$

Following Section 4.4.1, if Assumption 5 or 6 does not hold, then we would be able to find an $\mathbf{\Omega}$ such that there would exist a two-by-two comparison between two groups and two time periods $g, k \neq g, t, t' \neq t$ entering with a positive weight in $E[\hat{\beta}_{fe}|\mathbf{D}]$, while being a biased estimator of its corresponding ATE. Hence, Assumption 5 and 6 are both necessary and sufficient for the above statement to hold.

Introduction of Dynamics: Details

A General Decomposition Result in the Dynamic Framework First, note that Assumptions 8 and 9 are such that the standard, reverse leaver and double-switcher difference-and-differences are unbiased estimators of their corresponding ATEs. Focusing on the standard difference-in-differences, we have:

$$
\begin{aligned}
&E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},0))|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},0)) + Y_{g,t}(\mathbf{D}_{g,t-1},0) - Y_{g,t}(\mathbf{D}_{g,t-1},0)|\mathbf{D} \\
=&E[\Delta_{g,t}(\mathbf{D}_{g,t-1}) + Y_{g,t}(\mathbf{D}_{g,t-1},0) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},0))|\mathbf{D}] \\
=&E[\Delta_{g,t}(\mathbf{D}_{g,t-1})|\mathbf{D}]
\end{aligned}
$$

Similarly, for the reverse leaver difference-in-differences, under Assumptions 8 and 9, we have:

$$
\begin{aligned}
&E[Y_{g,t}(\mathbf{D}_{g,t-1},0) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1))|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{D}_{g,t-1},0) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1)) + Y_{k,t'}(\mathbf{D}_{k,t'-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},0)| \\
=&E[\Delta_{k,t'}(\mathbf{D}_{k,t'-1}) + Y_{g,t}(\mathbf{D}_{g,t-1},0) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},0))|\mathbf{D}] \\
=&E[\Delta_{k,t'}(\mathbf{D}_{k,t'-1})|\mathbf{D}]
\end{aligned}
$$

While, for the double switcher difference-in-difference, we have, under Assumptions 8 and 9:

$$
\begin{aligned}
&E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1)|\mathbf{D}] \\
=&E[\Delta_{g,t}(\mathbf{D}_{g,t-1}) + \Delta_{k,t'}(\mathbf{D}_{k,t'-1})|\mathbf{D}]
\end{aligned}
$$

Finally, while details for the reverse difference-in-differences are given in the body of the paper, the leaver difference-in-differences is such that:

$$
\begin{aligned}
&E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1)|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1) + Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t}(\mathbf{D}_{k,t-1},1)|\mathbf{D}] \\
=&E[\Delta_{k,t}(\mathbf{D}_{k,t-1}) + Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1))|\mathbf{D}]
\end{aligned}
$$

It is thus necessary and sufficient for $E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1))|\mathbf{D}]$ to be zero for the reverse difference-in-difference to be an unbiased estimator of the ATE of group k in period t.

Sufficient Conditions in the Presence of Dynamic Treatment Effects Let us focus on the reverse difference-in-differences and derive sufficient conditions for those to estimate their corresponding ATE without bias. For that, we would need the following equation to be equal to zero:

$$\begin{aligned}
&E[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1))|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{D}_{k,t-1},1) - Y_{g,t'}(\mathbf{D}_{k,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1)) \\
&+ [Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t}(\mathbf{D}_{k,t-1},1)] - [Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - Y_{g,t'}(\mathbf{D}_{k,t'-1},1)]\,|\mathbf{D}]
\end{aligned} \tag{4.10}$$

Let us first show that, under Assumptions 8-10 we have that $E[Y_{g,t}(\mathbf{D}_{k,t-1},1) - Y_{g,t'}(\mathbf{D}_{k,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1)|\mathbf{D}] = 0$. Indeed, we have:

$$\begin{aligned}
&E[Y_{g,t}(\mathbf{D}_{k,t-1},1) - Y_{g,t'}(\mathbf{D}_{k,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1)|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{D}_{k,t-1},1) - Y_{g,t'}(\mathbf{D}_{k,t'-1},1) - (Y_{k,t}(\mathbf{D}_{k,t-1},1) - Y_{k,t'}(\mathbf{D}_{k,t'-1},1) \\
&+ Y_{g,t}(\mathbf{D}_{k,t-1},0) - Y_{g,t}(\mathbf{D}_{k,t-1},0) + Y_{g,t'}(\mathbf{D}_{k,t'-1},0) - Y_{g,t'}(\mathbf{D}_{k,t'-1},0) \\
&+ Y_{k,t}(\mathbf{D}_{k,t-1},0) - Y_{k,t}(\mathbf{D}_{k,t-1},0) + Y_{k,t'}(\mathbf{D}_{k,t'-1},0) - Y_{k,t'}(\mathbf{D}_{k,t'-1},0)|\mathbf{D}] \\
=&E[\Delta_{g,t}(D_{k,t-1}) - \Delta_{g,t'}(D_{k,t'-1}) - \Delta_{k,t}(D_{k,t-1}) + \Delta_{k,t'}(D_{k,t'-1})) \\
&+ (Y_{g,t}(D_{k,t-1},0) - Y_{g,t'}(D_{k,t'-1},0) - (Y_{k,t}(D_{k,t-1},0) - Y_{k,t'}(D_{k,t'-1},0)))|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{0}_{t-1},1) - Y_{g,t}(\mathbf{0}_{t-1},0) + \tau_t(\sum_{\ell=1}^{t-1} D_{k,\ell}) - (Y_{g,t'}(\mathbf{0}_{t'-1},1) - Y_{g,t'}(\mathbf{0}_{t'-1},0) + \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell})) \\
&- (Y_{k,t}(\mathbf{0}_{t-1},1) - Y_{k,t}(\mathbf{0}_{t-1},0) + \tau_t(\sum_{\ell=1}^{t-1} D_{k,\ell})) + (Y_{k,t'}(\mathbf{0}_{t'-1},1) - Y_{k,t'}(\mathbf{0}_{t'-1},0) + \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell}))|\mathbf{D}] \\
=&E[Y_{g,t}(\mathbf{0}_{t-1},1) - Y_{g,t'}(\mathbf{0}_{t'-1},1) - (Y_{k,t}(\mathbf{0}_{t-1},1) - Y_{k,t'}(\mathbf{0}_{t'-1},1)) \\
&- [Y_{g,t}(\mathbf{0}_{t-1},0) - Y_{g,t'}(\mathbf{0}_{t'-1},0) - (Y_{k,t}(\mathbf{0}_{t-1},0) - Y_{k,t'}(\mathbf{0}_{t'-1},0))]|\mathbf{D}] \\
=&0
\end{aligned}$$

where the second equality follows from Assumptions 8 and 9, and the third equality from Assumption 10.2. The first term of the fourth equality cancels out following from Assumption 10.1, while the second term does following Assumptions 8 and 9.

Let us now examine the last two terms of Equation (4.10) to understand what does the reverse difference-in-differences identifies under Assumption 10. In particular, we have:

$$
\begin{aligned}
& E\left[Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t}(\mathbf{D}_{k,t-1},1)\right] - \left[Y_{g,t'}(\mathbf{D}_{g,t'-1},1) - Y_{g,t'}(\mathbf{D}_{k,t'-1},1)\right]|\mathbf{D}] \\
&= E\Bigg[[\Delta_{g,t}(D_{g,t-1}) + Y_{g,t}(\mathbf{D}_{g,t-1},0) - \Delta_{g,t}(D_{k,t-1}) - Y_{g,t}(\mathbf{D}_{k,t-1},0)] \\
& -[\Delta_{g,t'}(D_{g,t'-1}) + Y_{g,t'}(\mathbf{D}_{g,t'-1},0) - \Delta_{g,t'}(D_{k,t'-1}) - Y_{g,t'}(\mathbf{D}_{k,t-1},0)]|\mathbf{D}\Bigg] \\
&= E\Bigg[[Y_{g,t}(\mathbf{0}_{t-1},1) + \tau_t(\sum_{\ell=1}^{t-1} D_{g,\ell}) - Y_{g,t}(\mathbf{0}_{t-1},1) - \tau_t(\sum_{\ell=1}^{t-1} D_{k,\ell})] \\
& -[Y_{g,t'}(\mathbf{0}_{t'-1},1) + \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{g,\ell}) - Y_{g,t'}(\mathbf{0}_{t'-1},1) - \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell})]|\mathbf{D}\Bigg] \\
&= E\Bigg[[\tau_t(\sum_{\ell=1}^{t-1} D_{g,\ell}) - \tau_t(\sum_{\ell=1}^{t-1} D_{k,\ell})] - [\tau_{t'}(\sum_{\ell=1}^{t'-1} D_{g,\ell}) - \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell})]|\mathbf{D}\Bigg]
\end{aligned}
\tag{4.11}
$$

where the first equality uses the definition of $\Delta_{g,t}(D_{g,t-1}) = Y_{g,t}(\mathbf{D}_{g,t-1},1) - Y_{g,t}(\mathbf{D}_{g,t-1},0)$. The second equality follows from Assumptions 10.2, according to which $\Delta_{g,t}(D_{g,t-1}) = Y_{g,t}(\mathbf{0}_{t-1},1) - Y_{g,t}(\mathbf{0}_{t-1},0) + \tau_t(\sum_{\ell=1}^{t-1} D_{g,t-1})$ and Assumption 9. The last equality stems from Assumption 10.1.

This implies that if groups g and k have been exposed to treatment the same number of periods up to $t'-1$ and up to period $t-1$, all those terms cancel out and one can identify the ATE of group g in t', conditional on facing history $D_{g,t}$. This would for example happen when comparing periods 2 and 4 of group g with $\mathbf{D}_{g,t} = (0,0,1,1)$ and group k with $\mathbf{D}_{k,t} = (0,1,0,1)$.

What if the number of periods of exposure to treatment between group g and k differs only in $t'-1$, but not in $t-1$, i.e. $\sum_{\ell=1}^{t-1} D_{g,\ell} = \sum_{\ell=1}^{t-1} D_{k,\ell}$, but $\sum_{\ell=1}^{t'-1} D_{g,\ell} \neq \sum_{\ell=1}^{t'-1} D_{k,\ell}$? Now the reverse difference-in-differences would identify the following object:

$$
\begin{aligned}
&E[\Delta_{g,t'}(\mathbf{D}_{g,t'-1}) - [\tau_{t'}(\sum_{\ell=1}^{t'-1} D_{g,\ell}) - \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell})]|\mathbf{D}] \\
&= E[\Delta_{g,t'}(\mathbf{0}_{t'-1}) + \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{g,\ell}) - [\tau_{t'}(\sum_{\ell=1}^{t'-1} D_{g,\ell}) - \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell})]|\mathbf{D}] \\
&= E[\Delta_{g,t'}(\mathbf{0}_{t'-1}) + \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{k,\ell})]|\mathbf{D}] \\
&= E[\Delta_{g,t'}(\mathbf{D}_{k,t'-1})]|\mathbf{D}]
\end{aligned}
\tag{4.12}
$$

Thus, in this case, the reverse difference-in-differences identifies the ATE of group g in period t' conditional on history $\mathbf{D}_{k,t'-1}$. For example, comparing a group g such that $\mathbf{D}_{g,t} = (1,0,0,1)$ to a group k with $\mathbf{D}_{k,t} = (0,0,1,1)$ in periods 3 and 4 would allow to identify the ATE of group g in period 3 when being exposed to treatment for the first time.

Now, what if the number of periods of exposure to treatment differs only in $t-1$, but not $t'-1$, i.e. $\sum_{\ell=1}^{t-1} D_{g,\ell} \neq \sum_{\ell=1}^{t-1} D_{k,\ell}$, but $\sum_{\ell=1}^{t'-1} D_{g,\ell} = \sum_{\ell=1}^{t'-1} D_{k,\ell}$? In this case, the reverse difference-in-differences would identify the following object:

$$
E[\Delta_{g,t'}(\mathbf{0}_{t'-1}) + \tau_{t'}(\sum_{\ell=1}^{t'-1} D_{g,\ell}) + [\tau_{t}(\sum_{\ell=1}^{t-1} D_{g,\ell}) - \tau_{t}(\sum_{\ell=1}^{t-1} D_{k,\ell})]|\mathbf{D}]
\tag{4.13}
$$

Thus, in this case, the reverse difference-in-differences is a biased estimator of the ATE of group g in period t' conditional on history $\mathbf{D}_{g,t'-1}$. In particular, this is the case of staggered difference-in-differences, when groups are not allowed to leave treatment.

In a similar fashion, one can show that the leaver difference-in-differences comparing periods t and t', with $t > t'$ identifies the ATE of the group switching treatment status in period t under its own history if both $\sum_{\ell=1}^{t-1} D_{g,\ell} = \sum_{\ell=1}^{t-1} D_{k,\ell}$ and $\sum_{\ell=1}^{t'-1} D_{g,\ell} = \sum_{\ell=1}^{t'-1} D_{k,\ell}$. It would identify the ATE of the group switching treatment status in period t under the history of the comparison group if $\sum_{\ell=1}^{t'-1} D_{g,\ell} = \sum_{\ell=1}^{t'-1} D_{k,\ell}$ while $\sum_{\ell=1}^{t-1} D_{g,\ell} \neq \sum_{\ell=1}^{t-1} D_{k,\ell}$.

We now know under which assumptions the reverse and leaver difference-in-differences are unbiased estimators of the ATE of group g in period t, under the treatment history of the comparison group (which may be the same as the one of group g). When only such valid comparisons exist - or that there exist no reverse or leaver difference-in-differences -

the TWFE estimator is thus robust to heterogeneity: it only weighs positively unbiased estimators of their corresponding ATE.

4.A.3 Proof of Theorem 5

We want to prove that DID_M^A is an unbiased estimator of δ^S. Let us write the expectation of DID_M^A:

$$
\begin{aligned}
E\big[DID_M^A\big] =& \sum_{t=1}^{T} E\Bigg[\frac{1}{N_S}\sum_{g:D_{g,t}=1,D_{g,t-1}=0}\big[N_{g,t}\mathbb{1}\{\exists k\neq g, D_{k,t}=D_{k,t-1}=0\}E\big[DID_{+,g,t}|\mathbf{D}\big] \\
&+ N_{g,t}\mathbb{1}\{\nexists k\neq g, D_{k,t}=D_{k,t-1}=0\}E\big[DID^R_{-,g,t+1}|\mathbf{D}\big]\big] \\
&+ \frac{1}{N_S}\sum_{g:D_{g,t}=0,D_{g,t-1}=1}\big[N_{g,t}\mathbb{1}\{\exists k\neq g, D_{k,t}=D_{k,t-1}=1\}E\big[DID_{-,g,t}|\mathbf{D}\big] \\
&+ N_{g,t}\mathbb{1}\{\nexists k\neq g, D_{k,t}=D_{k,t-1}=1\}E\big[DID^R_{+,g,t+1}|\mathbf{D}\big]\big]\Bigg]
\end{aligned}
\tag{4.14}
$$

Let us look separately at each conditional expectation:

$$
\begin{aligned}
E\left(DID_{+,g,t}|\mathbf{D}\right) =& E\left(\mathbb{1}\{D_{g,t}=1,D_{g,t-1}=0\}\left[(Y_{g,t}-Y_{g,t-1})-\sum_{k:D_{k,t}=D_{k,t-1}=0}\frac{N_{k,t}}{N_{0,0,t}}(Y_{k,t}-Y_{k,t-1})\right]|\mathbf{D}\right) \\
=& \mathbb{1}\{D_{g,t}=1,D_{g,t-1}=0\}\left[E\left(Y_{g,t}-Y_{g,t-1}|\mathbf{D}\right)-\sum_{k:D_{k,t}=D_{k,t-1}=0}\frac{N_{k,t}}{N_{0,0,t}}E\left(Y_{k,t}-Y_{k,t-1}|\mathbf{D}\right)\right]
\end{aligned}
$$

For every g that $D_{g,t-1}=0$ and $D_{g,t}=1$, we have:

$$
E\left(Y_{g,t}-Y_{g,t-1}|\mathbf{D}\right)=E\left(\Delta_{g,t}|\mathbf{D}\right)+E\left(Y_{g,t}(0)-Y_{g,t-1}(0)|\mathbf{D}\right) \tag{4.15}
$$

Following de Chaisemartin and d'Haultfoeuille (2020a), under Assumptions 4, 5 and 11, there exists a real number $\psi_{0,t}$ such that for all g,

$$\begin{aligned} E\left(Y_{g,t}(0) - Y_{g,t-1}(0)|\mathbf{D}\right) &= E\left(Y_{g,t}(0) - Y_{g,t-1}(0)|\mathbf{D}_g\right) \\ &= E\left(Y_{g,t}(0) - Y_{g,t-1}(0)\right) \\ &= \psi_{0,t} \end{aligned} \tag{4.16}$$

where $\mathbf{D}_g$ is the vector collecting treatment status of group g over time. Then, we have:

$$\begin{aligned} &E\left(DID_{+,g,t}|\mathbf{D}\right) \\ =&\mathbb{1}\{D_{g,t}=1, D_{g,t-1}=0\}\Big[E(\Delta_{g,t}|\mathbf{D}) + E(Y_{g,t}(0) - Y_{g,t-1}(0)|\mathbf{D}) \\ &- \sum_{k:D_{k,t}=D_{k,t-1}=0} \frac{N_{k,t}}{N_{0,0,t}} E(Y_{k,t}(0) - Y_{k,t-1}(0)|\mathbf{D})\Big] \\ =&\mathbb{1}\{D_{g,t}=1, D_{g,t-1}=0\}\left[E\left(\Delta_{g,t}|\mathbf{D}\right) + \psi_{o,t}(1 - \sum_{k:D_{k,t}=D_{k,t-1}=0} \frac{N_{k,t}}{N_{0,0,t}})\right] \\ =&\mathbb{1}\{D_{g,t}=1, D_{g,t-1}=0\}\left[E\left(\Delta_{g,t}|\mathbf{D}\right) + \psi_{o,t}(1 - \frac{1}{N_{0,0,t}} \underbrace{\sum_{k:D_{k,t}=D_{k,t-1}=0} N_{k,t})}_{=N_{0,0,,t}}\right] \\ =&\mathbb{1}\{D_{g,t}=1, D_{g,t-1}=0\}E\left[\Delta_{g,t}|\mathbf{D}\right] \end{aligned} \tag{4.17}$$

where the first equality follows from (4.15), the second equality from (4.16) and the third one uses the definition of $N_{0,0,t}$.

A similar reasoning yields:

$$E\left(DID_{-,g,t}|\mathbf{D}\right) = \mathbb{1}\{D_{g,t}=0, D_{g,t-1}=1\}E\left[\Delta_{g,t}|\mathbf{D}\right] \tag{4.18}$$

$$E\left(DID^R_{+,g,t}|\mathbf{D}\right) = \mathbb{1}\{D_{g,t}=1, D_{g,t-1}=0\}E\left[\Delta_{g,t-1}|\mathbf{D}\right] \tag{4.19}$$

$$E\left(DID^R_{-,g,t}|\mathbf{D}\right) = \mathbb{1}\{D_{g,t}=0, D_{g,t-1}=1\}E\left[\Delta_{g,t-1}|\mathbf{D}\right] \tag{4.20}$$

Plugging (4.17), (4.18), (4.19) and (4.20) in (4.14), we have:

$$\begin{aligned}
E\left[DID_M^A\right] =& E\left[\sum_{t=2}^{T}\frac{1}{N_S}\left[\sum_{g:D_{g,t}=1,D_{g,t-1}=0}\left[N_{g,t}\mathbb{1}\{\exists k\neq g, D_{k,t}=D_{k,t-1}=0\}E\left[\Delta_{g,t}|\mathbf{D}\right]\right.\right.\right.\\
&\left.\left.+N_{g,t}\mathbb{1}\{\nexists k\neq g, D_{k,t}=D_{k,t-1}=0\}\mathbb{1}\{D_{g,t+1}=0, D_{g,t}=1\}E\left[\Delta_{g,t}|\mathbf{D}\right]\right]\right]\\
&+\frac{1}{N_S}\left[\sum_{g:D_{g,t}=0,D_{g,t-1}=1}\left[N_{g,t}\mathbb{1}\{\exists k\neq g, D_{k,t}=D_{k,t-1}=1\}E\left[\Delta_{g,t}|\mathbf{D}\right]\right.\right.\\
&\left.\left.\left.+N_{g,t}\mathbb{1}\{\nexists k\neq g, D_{k,t}=D_{k,t-1}=1\}\mathbb{1}\{D_{g,t+1}=1, D_{g,t}=0\}E\left[\Delta_{g,t}|\mathbf{D}\right]\right]\right]\right]
\end{aligned} \tag{4.21}$$

Under Assumption 12, we have that if there is a group g such that $D_{g,t}=1$ and $D_{g,t-1}=0$ then there either exists at least one comparison group which is untreated, or g is such that $D_{g,t+1}=0$ and there exists a group k which is untreated. Then, this implies that, for a given g such that $D_{g,t}=1$ and $D_{g,t-1}=0$:

$$\mathbb{1}\{\exists k\neq g, D_{k,t}=D_{k,t-1}=0\}+\mathbb{1}\{\nexists k\neq g, D_{k,t}=D_{k,t-1}=0\}\mathbb{1}\{D_{g,t+1}=0, D_{g,t}=1\}=1$$

Similarly, for g such that $D_{g,t}=0$ and $D_{g,t-1}=1$:

$$\mathbb{1}\{\exists k\neq g, D_{k,t}=D_{k,t-1}=1\}+\mathbb{1}\{\nexists k\neq g, D_{k,t}=D_{k,t-1}=1\}\mathbb{1}\{D_{g,t+1}=1, D_{g,t}=0\}=1$$

Thus, Equation (4.21) writes:

$$\begin{aligned}
E\left[DID_M^A\right] &= \sum_{t=2}^{T}E\left[E\left[\frac{1}{N_S}\left(\sum_{g:D_{g,t}=1,D_{g,t-1}=0}N_{g,t}\Delta_{g,t}+\sum_{g:D_{g,t}=0,D_{g,t-1}=1}N_{g,t}\Delta_{g,t}\right)|\mathbf{D}\right]\right]\\
&=\delta^S
\end{aligned} \tag{4.22}$$

References

Abdulkadiroğlu, Atila and Tayfun Sönmez, "School choice: A mechanism design approach," *American economic review*, 2003, *93* (3), 729–747.

_ **, Nikhil Agarwal, and Parag A Pathak**, "The welfare effects of coordinated assignment: Evidence from the New York City high school match," *American Economic Review*, 2017, *107* (12), 3635–3689.

_ **, Parag A Pathak, Jonathan Schellenberg, and Christopher R Walters**, "Do parents value school effectiveness?," *American Economic Review*, 2020, *110* (5), 1502–1539.

Agarwal, Nikhil and Paulo J Somaini, "Demand Analysis under Latent Choice Constraints," *NBER Working Paper*, 2022.

_ **and Paulo Somaini**, "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism," *Econometrica*, 2018, *86* (2), 391–444.

_ **, Itai Ashlagi, Michael A. Rees, Paulo Somaini, and Daniel Waldinger**, "Equilibrium Allocations Under Alternative Waitlist Designs: Evidence From Deceased Donor Kidneys," *Econometrica*, 2021, *89* (1), 37–76.

Akerlof, George, "Procrastination and Obedience," *American Economic Review*, 1991, *81* (2), 1–19.

Altonji, Joseph, Peter Arcidiacono, and Arnaud Maurel, "The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects," *Handbook of the Economics of Education*, 2016, *5*, 305–96.

Arcidiacono, Peter and Robert A Miller, "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

Austin, Benjamin A, Edward L Glaeser, and Lawrence H Summers, "Jobs for the Heartland: Place-based policies in 21st century America," *NBER Working Paper*, 2018.

Azevedo, Eduardo M and Jacob D Leshno, "A supply and demand framework for two-sided matching markets," *Journal of Political Economy*, 2016, *124* (5), 1235–1268.

Bó, Ignacio and Rustamdjan Hakimov, "The Iterative Deferred Acceptance Mechanism," *Working Paper*, 2016.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event study designs: Robust and efficient estimation," *arXiv preprint arXiv:2108.12419*, 2022.

Bound, John, Jeffrey Groen, Gabor Kezdi, and Sarah Turner, "Trade in university training: cross-state variation in the production and stock of college-educated labor," *Journal of Econometrics*, 2004, *121* (1-2), 143–173.

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the incidence and efficiency of a prominent place based policy," *American Economic Review*, 2013, *103* (2), 897–947.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna, "Difference-in-differences with a continuous treatment," *arXiv preprint arXiv:2107.02637*, 2021.

Calsamiglia, Caterina, Chao Fu, and Maia Güell, "Structural estimation of a model of school choices: The boston mechanism versus its alternatives," *Journal of Political Economy*, 2020, *128* (2), 642–680.

Card, David, Jesse Rothstein, and Moises Yi, "Location, location, location," Technical Report, National Bureau of Economic Research 2023.

Chabé-Ferret, Sylvain and Anca Voia, "Are Grassland Conservation Programs a Cost-Effective Way to Fight Climate Change? Evidence from France," Technical Report 2021.

Chen, Li and Juan Sebastián Pereyra, "Time-constrained school choice," *Working Paper*, 2019.

Chetty, Raj and Nathaniel Hendren, "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects," *The Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

_ , John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan, "Mobility report cards: The role of colleges in intergenerational mobility," Technical Report, National Bureau of Economic Research 2018.

_ , _ , Nathaniel Hendren, Maggie R Jones, and Sonya R Porter, "The opportunity atlas: Mapping the childhood roots of social mobility," Technical Report, National Bureau of Economic Research 2018.

_ , Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where is the land of opportunity? The geography of intergenerational mobility in the United States," *The Quarterly Journal of Economics*, 2014, *129* (4), 1553–1623.

Chyn, Eric and Lawrence F Katz, "Neighborhoods matter: Assessing the evidence for place effects," *Journal of Economic Perspectives*, 2021, *35* (4), 197–222.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon, "Spatial wage disparities: Sorting matters!," *Journal of urban economics*, 2008, *63* (2), 723–742.

_ , _ , _ , and Sébastien Roux, "Sorting and local wage and skill distributions in France," *Regional Science and Urban Economics*, 2012, *42* (6), 913–930.

Cour des Comptes, "Admissions post-bac et accès á l'enseignment supèrieur. Un dispositif contesté á réformer," *Rapport public*, 2017.

Dauth, Wolfgang, Sebastian Findeisen, Enrico Moretti, and Jens Suedekum, "Matching in cities," *Journal of the European Economic Association*, 2022, *20* (4), 1478–1521.

de Chaisemartin, Clément and Xavier d'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–96.

— **and** — , "Two-way fixed effects regressions with several treatments," *arXiv preprint arXiv:2012.10077*, 2020.

— **and Xavier D'Haultfoeuille**, "Difference-in-differences estimators of intertemporal treatment effects," Technical Report, National Bureau of Economic Research 2022.

— **and** — , "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey," *National Bureau of Economic Research*, 2022.

— **and Xavier d'Haultfoeuille**, "Fuzzy differences-in-differences," *The Review of Economic Studies*, 2018, *85* (2), 999–1028.

Diamond, Rebecca, "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000," *American Economic Review*, 2016, *106* (3), 479–524.

— **and Cecile Gaubert**, "Spatial sorting and inequality," *Annual Review of Economics*, 2022, *14*, 795–819.

Dillon, Eleanor Wiske and Jeffrey Andrew Smith, "Determinants of the match between student ability and college quality," *Journal of Labor Economics*, 2017, *35* (1), 45–66.

Fabre, Anaïs, "Robustness of Two-Way Fixed Effects Estimators to Heterogeneous Treatment Effects," *TSE Working Paper*, 2023.

Fack, Gabrielle, Julien Grenet, and Yinghua He, "Beyond truth-telling: Preference estimation with centralized school choice and college admissions," *American Economic Review*, 2019, *109* (4), 1486–1529.

Fu, Chao, Junjie Guo, Adam J Smith, and Alan Sorensen, "Students' heterogeneous preferences and the uneven spatial distribution of colleges," *Journal of Monetary Economics*, 2022, *129*, 49–64.

Gale, David and Lloyd S Shapley, "College admissions and the stability of marriage," *The American Mathematical Monthly*, 1962, *69* (1), 9–15.

Gibbons, Stephen and Anna Vignoles, "Geography, choice and participation in higher education in England," *Regional science and urban economics*, 2012, *42* (1-2), 98–113.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gourieroux, Christian and Alain Monfort, *Simulation-based econometric methods*, Oxford university press, 1996.

Grenet, Julien, Yinghua He, and Dorothea Kübler, "Decentralizing centralized matching markets: Implications from early offers in university admissions," *WZB Discussion Paper*, 2019.

He, Yinghua, Shruti Sinha, and Xiaoting Sun, "Identification and Estimation in Many-to-one Two-sided Matching without Transfers," *arXiv preprint arXiv:2104.02009*, 2023.

Heckman, James J, "Heterogeneity and state dependence," in "Studies in labor markets," University of Chicago Press, 1981, pp. 91–140.

Hillman, Nicholas W, "Geography of college opportunity: The case of education deserts," *American Educational Research Journal*, 2016, *53* (4), 987–1021.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Howard, Greg, Russell Weinstein et al., *" Workhorses of Opportunity": Regional Universities Increase Local Social Mobility*, IZA-Institute of Labor Economics, 2022.

Hoxby and Sarah Turner, "What high-achieving low-income students know about college," *American Economic Review*, 2015, *105* (5), 514–17.

Hoxby, C and C Avery, "The missing "One-Offs": The hidden supply of high-achieving, low-income students. Washington, DC: Brookings Institution," 2013.

Hsiao, Allan, "Educational Investment in Spatial Equilibrium: Evidence from Indonesia," *Working Paper*, 2023.

Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow, "The allocation of talent and us economic growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Humphries, John Eric, Juanna Joensen, and Gregory Veramendi, "Complementarities in High School and College Investments," *Working Paper*, 2023.

Ishimaru, S, "Geographic Mobility of Youth and Spatial Gaps in Local College and Labor Market Opportunities," Technical Report, Working paper 2020.

Kapor, Adam, Christopher A. Neilson, and Seth D. Zimmerman, "Heterogeneous Beliefs and School Choice Mechanisms," *American Economic Review*, 2020.

_ , Mohit Karnani, and Christopher Neilson, "Aftermarket frictions and the cost of off-platform options in centralized assignment mechanisms," *Industrial Relations Working Paper Series (645)*, 2020.

Katz, Lawrence F and Kevin M Murphy, "Changes in relative wages, 1963–1987: supply and demand factors," *The quarterly journal of economics*, 1992, *107* (1), 35–78.

Kelchtermans, Stijn and Frank Verboven, "Program duplication in higher education is not necessarily bad," *Journal of Public Economics*, 2010, *94* (5-6), 397–409.

Kennan, John, "Spatial variation in higher education financing and the supply of college graduates," Technical Report, National Bureau of Economic Research 2015.

_ and James R Walker, "The effect of expected income on individual migration decisions," *Econometrica*, 2011, *79* (1), 211–251.

Kim, Kimin and Myoung jae Lee, "Difference in differences in reverse," *Empirical Economics*, 2019, *57* (3), 705–725.

Kline, Patrick and Enrico Moretti, "Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority," *The Quarterly journal of economics*, 2014, *129* (1), 275–331.

Ku, Hyejin, Uta Schönberg, and Ragnhild C Schreiner, "Do place-based tax incentives create jobs?," *Journal of Public Economics*, 2020, *191*, 104105.

Larroucau, Thomas and Ignacio Ríos, "Dynamic College Admissions and the Determinants of Students' College Retention," *Working Paper*, 2021.

Larroucau, Tomás and Ignacio Rios, "Do "Short-List" Students Report Truthfully? Strategic Behavior in the Chilean College Admissions Problem," Technical Report 2018.

Luflade, Margaux, “The value of information in centralized school choice systems,” *Working Paper*, 2018.

— , “The value of information in centralized school choice systems,” Technical Report 2019.

Magnac, Thierry and David Thesmar, “Identifying dynamic discrete decision processes,” *Econometrica*, 2002, *70* (2), 801–816.

— **and** — , “Identifying dynamic discrete decision processes,” *Econometrica*, 2002, *70* (2), 801–816.

Merle, Pierre, “La Catégorie socio-professionnelle des parents dans les fiches administratives des élèves. Quelles limites? Quels usages?,” *Socio-logos. Revue de l’association française de sociologie*, 2013, *8.*

Nimier-David, Elio, “Local human capital and firm creation: Evidence from the massification of higher education in France,” Technical Report, Working Paper.[1537] 2022.

Olden, Andreas and Jarle Møen, “The triple difference estimator,” *The Econometrics Journal*, 2022, *25* (3), 531–553.

Patnaik, Arpita, Matthew J. Wiswall, and Basit Zafar, “College Majors,” Technical Report, National Bureau of Economic Research 2020.

Rho, Minyoung, “Student Rank Order List Choice Model: A Dynamic Discrete Choice,” *Dissertation, Tepper School of Business, Carnegie Mellon University*, 2019.

Robins, James, “A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect,” *Mathematical modelling*, 1986, *7* (9-12), 1393–1512.

Rossi, Pauline and Paola Villar, “Private health investments under competing risks: evidence from malaria control in Senegal,” *Journal of Health Economics*, 2020, *73*, 102330.

Roth, Alvin and Uriel Rothblum, “Truncation strategies in matching markets —In search of advice for participants,” *Econometrica*, 1999, *67* (1), 21–43.

Roth, Alvin E and Xiaolin Xing, "Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists," *Journal of political Economy*, 1997, *105* (2), 284–329.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–22.

_ , Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023.

Silva, JMC Santos and Silvana Tenreyro, "The log of gravity," *The Review of Economics and statistics*, 2006, *88* (4), 641–658.

Strezhnev, Anton, "Semiparametric weighting estimators for multi-period difference-in-differences designs," Technical Report 2018.

Train, Kenneth, *Discrete choice methods with simulation*, Cambridge; New York: Cambridge University Press, 2009.

Train, Kenneth E, *Discrete choice methods with simulation*, Cambridge university press, 2009.

Troyan, Peter and Thayer Morrill, "Obvious manipulations," *Journal of Economic Theory*, 2020, *185.*

Waldinger, Daniel, "Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation," *American Economic Review*, 2021, *111* (8), 2660–2696.

Zerecero, Miguel, "The Birthplace Premium," Technical Report, mimeo 2021.