

The Choice Among Risky Prospects in Matching Markets: Evidence from Real-time Medical Match Data

Veli M. Andirin*

Yale University

October 13, 2026

[Updated frequently. Please click here for the latest version.](#)

Abstract

Agents (e.g. students, doctors) in most matching markets submit lists where they rank their options. Lists are risky prospects that differ in outcomes (e.g. a school assignment) and outcome probabilities. Standard models assume that agents would make choices to maximize their expected utility if they know how lists map to probabilities. However, agents often lack this knowledge. I study choices under risk in matching markets in a unique setting where agents have detailed information on assignment chances under alternative lists. A match simulation platform for Turkey’s compulsory medical service match tells doctors placement probabilities to each hospital on their list in real time. Using snapshots of lists and assignment probabilities from this platform and an original survey about doctors’ preferences, I build and estimate a structural model that allows departures from expected utility. I find that applicants treat probabilities as proposed by prospect theory rather than maximize expected utility: They overweight nearly certain outcomes relative to merely probable ones. An analyst who assumes expected utility and ignores probability weighting would overstate the welfare gains from information provision by more than fifty percent. Moreover, taking probability weighting into account can reverse the welfare ranking of alternative assignment mechanisms. Finally, heterogeneity in probability weighting has distributional consequences. Doctors who depart more from expected utility benefit less from information and more from strategy-proof alternatives.

*I am grateful to Joseph Altonji, Leah Boustan, Charles Hodgson, Adam Kapor, Cormac O’Dea, and Jesse Shapiro for their invaluable guidance. I thank Nicholas Barberis, Steven Berry, Barbara Biasi, Janet Currie, Philip Haile, Alena Kang-Landsberg, Yuichi Kitamura, Bentley MacLeod, Victoria Marone, Mushfiq Mobarak, Chris Neilson, Larry Samuelson, Paulo Somaini, Mai Wo, and Seth Zimmerman for helpful feedback. I thank Ahmet Ay, the owner of Proders Company, for data and continued support. The project was financially supported by the Cowles Foundation, and received approval from the IRB at Yale. All errors are my own. E-mail: velimurat.andirin@yale.edu.

1 Introduction

Centralized assignment mechanisms match people (e.g., students, doctors) to institutions (e.g., schools, residency programs) throughout the world, and in most of these settings applicants need to submit a list where they rank their options. Each list is a lottery that offers particular outcomes (e.g., admission to a school) and probabilities with which these outcomes can occur, usually conditional on the lists of other applicants.

It is common for applicants in matching markets to lack knowledge of how alternative lists translate into admission chances.¹ This lack of knowledge can arise because of a misunderstanding of the mechanism², or simply because of holding incorrect beliefs about admission chances to a particular option.³ The literature typically assumes that, absent these barriers, applicants would choose the list that maximizes expected utility.

However, a lack of information about admission chances under alternative lists might not be the only barrier to making “optimal choices” in an expected utility sense. There are alternative theories of how people decide under risk. For instance, in their seminal prospect theory paper, Kahneman and Tversky (1979) emphasize that people overweight the outcomes that are attained with certainty relative to merely probable ones. If applicants in matching markets behave according to an alternative theory of choice under risk rather than standard expected utility maximization, their choices need not maximize expected utility even if they know their admission chances.

In this paper I study how applicants make choices under risk in matching markets using a unique setting where they have access to detailed, real-time information on their admission chances. In the days leading up to Turkey’s compulsory medical service match, nearly all doctors enroll on a match simulation platform. They submit their tentative lists to this platform, and the platform feeds these lists into simulations that run every few seconds to tell doctors their placement chances to each hospital on their list in real time. As a result, when a doctor or one of her competitors changes a list, the probabilities she faces move within seconds.

Using data from this platform on doctors’ lists and the assignment probabilities combined with an original survey of their preferences, I build and estimate a model of choice that allows departures from expected utility maximization. I show that applicants weight probabilities non-linearly, and these weights are characterized by the inverse S-shape proposed in the prospect theory litera-

¹Kapor, Neilson and Zimmerman (2020), Arteaga et al. (2022), and Larroucau et al. (2025) show that inaccurate beliefs about admission chances are widespread in different contexts.

²For instance, under the widely-used Boston mechanism, priorities depend on the order in which options are ranked. As a result, a student who ranks a school first has priority over every student who ranks it second. Applicants who do not understand such strategic incentives may act naively and simply rank schools from their most preferred to least preferred (see, e.g., Kapor, Neilson and Zimmerman 2020; He 2015; Calsamiglia, Fu and Güell 2020).

³See, e.g., Arteaga et al. (2022), who also show that these inaccurate beliefs can matter even in “strategy-proof” mechanisms as overoptimism about assignment chances can lead applicants to cut their lists short because of search costs.

ture: Applicants overweight nearly certain outcomes (those with probabilities close to zero or one) relative to the merely probable ones. I argue that taking probability weighting into account matters for answers to policy-relevant questions: A model that ignores departures from expected utility overstates the welfare gains from information about probabilities by more than fifty percent, and taking probability weighting into account can flip the welfare ranking of alternative mechanisms. Finally, I extend the model to allow for heterogeneity in probability weighting. I show that there are important distributional consequences of this heterogeneity: Doctors who distort probabilities more benefit less from information provision and more from strategy-proof alternatives.

I start by describing the institutional context of the compulsory medical service in Turkey. Doctors in Turkey are required to work in underserved areas of the country after medical school for around a year. A centralized mechanism assigns them to hospitals, using Immediate Acceptance with random priorities in each round: If, at any given round of the match, the demand for a hospital exceeds its available seats, all applicants to that hospital have equal chances of getting in. These random priorities make every doctor's outcome depend on what others do, so knowing the probabilities under alternative lists matters even more here than in a setting where other factors such as exam scores are taken into account.

What makes this setting exceptional is that a simulation platform delivers this high-stakes information in fine detail and at high frequency. My main data are minute-by-minute snapshots of doctors' rank-order lists on this platform and of the assignment probabilities attached to them. I obtained these data from the simulation platform for the March and September 2025 rounds of the match. I supplement these data with original surveys about the true preferences, strategies, and future plans of doctors, conducted during the same rounds of the match. My data also include the final placement outcomes of the doctors. Comparing these final placements to the platform data, I show that the predictions on the platform track the real outcomes closely. Nearly every doctor predicted to be assigned to a hospital with certainty does end up at that hospital. Moreover, the platform-predicted match probabilities and real match shares overlap nearly one-to-one. I also gather data from official sources on the outcomes of the previous rounds of the match, the characteristics of cities in Turkey such as their development levels and reliance on the compulsory service program, and the residency exam performance of different medical schools.

Using these data, I first present descriptive evidence that further explains the key facts about this context and motivates my modeling choices. I start by showing the importance of the match both for cities in Turkey and for the doctors: Less developed areas of the country rely heavily on this program as a source of doctor supply, and survey responses by doctors show that the outcome of this match can have a major impact on their careers. The importance of the match for doctors translates into high activity on the platform: During any daytime hour, nearly 40% of doctors are active on the platform. Through event studies, I show that doctors respond almost immediately to

changes in probabilities. I also conducted a survey that solicited from doctors their stated favorite hospital. I show that nearly all survey respondents rank their ideal hospital as their top choice if they are certain to be placed there. Many doctors nonetheless drop their ideal hospital when they could reach it with a high probability (e.g., 70%) and when listing safety hospitals below it would remove the remaining non-assignment risk.

I then introduce an empirical model that combines doctors' survey responses about their ideal hospitals with their rank-order list choices on the platform. The part of the model for list-level choices allows doctors to depart from expected utility maximization and to weight probabilities non-linearly. I estimate this model using maximum simulated likelihood. For the probability weighting function, I use Bernstein polynomials as the sieve basis, which ensures that the results are not driven by the standard functional forms from the prospect theory literature.

The identification of the probability weighting function relies on two key sources of variation. First, the survey on which hospitals doctors would choose if they could do so freely provides measurements of hospital preferences independent of assignment probabilities, which allows isolating the effect of probabilities on list-level choices. Second, changes in probabilities are informative about the slope of the probability weighting function. Under probability weighting, unlike expected utility, the change in the utility of a rank-order list due to a change in the assignment probability to a hospital depends on initial probabilities. For example, under the usual inverse-S shape of prospect theory, a given change in probability moves list utility less in the "merely probable" middle range than it does near zero or near one, so doctors should respond less to it.

The estimates reveal a probability weighting function with the inverse-S shape that prospect theory predicts: An overweighting of small probabilities, a middle region where there is insensitivity to increases in probabilities, and a steep rise towards certainty. Allowing for probability weighting also changes the estimated preferences over hospital characteristics in an intuitive direction. Many hospitals in less developed areas offer certain placement. A model that ignores probability weighting therefore reads the choices of such hospitals as evidence that doctors do not mind less developed locations much. Once the premium on certainty is modeled, the estimated distaste for less developed locations gets larger.

I show that the estimated model with probability weighting fits the data much better than a model that simply assumes expected utility maximization. The probability weighting model matches the real share of doctors who choose rank-order lists that offer them guaranteed placement at a single hospital almost one-to-one, whereas the expected utility model struggles to explain why many doctors choose such lists. The model with probability weighting also outperforms the expected utility model in explaining why doctors drop their ideal hospital even when that hospital offers fairly good placement prospects.

Using the estimates from this model, I first evaluate the welfare impact of information provision

about assignment chances. It is *ex ante* unclear whether providing information about admission probabilities can increase welfare in an expected utility sense when doctors behave as if nearly certain outcomes carry an extra premium. The key challenge for evaluating the impact of information provision is the absence of data on what the rank-order lists would be in the counterfactual scenario where the doctors did not have access to information on placement probabilities. I use doctors' initial choices on the platform to approximate their counterfactual lists, since the changes to their lists are at least partially driven by responses to probability changes caused by the actions of other users. A survey on what doctors' counterfactual top two choices would be if they did not have access to the platform corroborates the suitability of this approach.

I find that information provision about placement probabilities is still welfare-improving on average even under probability weighting. However, ignoring probability weighting would overstate these gains by more than fifty percent. The gap between the two models arises because doctors switch to less desirable hospitals over time to buy certainty. Probability weighting lets the model separate more strongly how desirable a hospital is from how likely it is, so it penalizes those switches more heavily.

I then examine whether accounting for probability weighting affects comparisons between alternative assignment mechanisms. I ask how this full-information version of the Boston mechanism compares to a strategy-proof alternative when probability weighting is taken into account versus when it is ignored. Since hospitals have no preferences over applicants, I use random serial dictatorship (RSD) as the strategy-proof benchmark. Since there are more than 2,000 hospitals that applicants can rank, RSD with unlimited list length can be impractical in reality, so I compare the two mechanisms at various list-length constraints on RSD.

I show that there is a range of list lengths where the probability weighting and expected utility models disagree on whether the Boston mechanism or RSD outperforms the other in terms of the average expected welfare. Ignoring probability weighting overstates the gains from a switch to RSD: a doctor with a low priority draw under RSD lands at an undesirable hospital, and probability weighting penalizes those outcomes more. These results therefore show that taking probability distortions into account can reverse the welfare ranking of alternative mechanisms.

Finally, I extend the baseline model to allow for unobserved heterogeneity in probability weighting. To this end, I use a parametric form for the probability weighting function taken from Tversky and Kahneman (1992) that allows me to capture the degree of departures from expected utility maximization via a single parameter which I allow to be person-specific. After estimating the distribution of this parameter in the population, for each doctor, I calculate the posterior mean of a measure of distortion that captures how much each doctor departs from expected utility maximization. I show that doctors who distort probabilities more benefit less from information and more from a strategy-proof alternative. It is often said that manipulable mechanisms are difficult. These

results explain for whom they are so.

This paper connects several strands of the literature. It first shows that probability weighting, which is a key element of the prospect theory of Kahneman and Tversky (1979), is relevant in matching markets. To my knowledge, no existing empirical paper studies probability weighting in markets where the primary outcome is not monetary.⁴ As Barberis (2013) notes, most existing applications of prospect theory are in finance, insurance, and gambling markets. A closely related paper in the context of insurance choices is Barseghyan et al. (2013), who estimate a structural model featuring probability weighting using micro data on deductible choices and find that probability weighting plays a significant role in explaining behavior. Andrikogiannopoulou and Papakonstantinou (2020, 2021) show evidence of probability weighting in betting markets using data from individual bettors. In addition to estimating a structural model with probability weighting in a setting with high-stakes and non-monetary outcomes, I rely on a context where the agents are given accurate predictions of probabilities every few seconds. By contrast, existing work needs to impose assumptions such as rational beliefs or the assumption that players read probabilities off the odds to solve the problem of jointly identifying preferences and beliefs.⁵

Another strand of the literature that my paper builds on studies the role of information about assignment chances in matching markets. Kapor, Neilson and Zimmerman (2020) study the role of beliefs about admission chances under the old New Haven mechanism, which closely resembles the Boston mechanism, and find that school choice participants often hold inaccurate beliefs. Larroucau et al. (2025) show that both overoptimism and underoptimism about admission chances are common in Chilean college admissions, and sending warnings about low placement probabilities can correct these mistakes. Dur, Hammond and Morrill (2018) study school choice in the Wake County Public School System, North Carolina, where the information on the number of applicants ranking a school as their first option is displayed on a website, and they show that the applicants who visit the website more frequently have better placement outcomes. Arteaga et al. (2022) study the effects of sending warnings to students whose risk of non-placement exceeds a threshold in the contexts of school choice in Chile and New Haven where Deferred Acceptance is used and show that these warnings reduce non-placement risk. Together, these papers show that inaccurate beliefs about admission chances are common, and correcting these mistakes can be helpful. What

⁴Rees-Jones and Shorrer (2023) review the nascent literature that brings ideas from behavioral economics to the field of education market design. Recent work from this literature includes Dreyfuss, Heffetz and Rabin (2022), Meisner and von Wangenheim (2023), and Dreyfuss et al. (2025), who study the role of expectations-based reference dependence in explaining dominated choices in strategy-proof mechanisms; as well as Idoux (2023) and Wang, Wang and Ye (2025), who use models of directed cognition (Gabaix et al. 2006) to explain such choices.

⁵Other previous work that examines the evidence for probability weighting with field data includes Cicchetti and Dubin (1994), who study insurance decisions for phone line trouble, Jullien and Salanié (2000), Chiappori et al. (2019), and Gandhi and Serrano-Padial (2014), who study horse races, and Kliger and Levy (2009) who study financial markets. More broadly, this work is also related to the structural behavioral economics literature reviewed in DellaVigna (2018).

I show in this paper is that even when nearly all applicants have access to very detailed, real-time information, behavioral frictions can still prevent the use of such information to make choices that maximize expected welfare, and the variation in the extent of such behavioral frictions can determine who benefits more from information.

My work also connects to the broader empirical literature that compares the welfare impacts of alternative assignment mechanisms. Particularly related papers are He (2015), Agarwal and Somaini (2018), Calsamiglia, Fu and Güell (2020), and Kapor, Neilson and Zimmerman (2020), who study welfare outcomes under manipulable mechanisms versus strategy-proof ones. In this paper, I use Random Serial Dictatorship as the strategy-proof alternative to the Boston mechanism since hospitals do not have priorities over applicants. I show that expected utility and probability weighting models disagree on the extent of welfare gains from a switch to a strategy-proof alternative. Moreover, for a range of list-length restrictions on RSD, the comparison in terms of which mechanism outperforms the other flips when probability weighting is taken into account. Finally, my work sheds light on the distributional consequences of alternative mechanisms by showing that people who depart more from expected utility maximization benefit more from strategy-proof alternatives.⁶

The rest of the paper proceeds as follows: I present the details about the compulsory service in Turkey, the assignment mechanism, and the data in Section 2. In Section 3, I present descriptive evidence on the activity of doctors on the platform using snapshots of lists and assignment probabilities as well as their survey responses, and provide initial evidence on the value of information provision using the average non-placement risk and placement chances to high-ranked options as measures. I present the structural model with probability weighting in Section 4. I present the model estimates and discuss its fit in Section 5. I revisit the value of information in Section 6 using the estimates from this model. I compare the Boston mechanism with full information to Random Serial Dictatorship in Section 7. Section 8 estimates a model with heterogeneity in probability weighting and discusses its implications for welfare. Section 9 concludes.

⁶My work also relates to the literature on matching in medical markets. Medical matching markets have been a central motivation for the study of matching mechanisms (see Roth and Sotomayor 1992). However, the empirical literature on the subject is relatively limited. Agarwal (2015) focuses on identifying preferences on both sides of the market in the context of medical residency matching in the US, when only data on realized outcomes are observed. Friedrich et al. (2024) study the role of interdependent values in admissions to Denmark’s medical programs. Unlike these two papers, in the context I study, hospitals do not have preferences over applicants. My work also contributes to the literature on doctors’ location choices (Falcettoni 2018; Costa, Nunes and Sanches 2024), though the focus in this paper is on the role of applicant behavior under information provision, using real-time preference and match probability data.

2 Institutional Background and Data

I start this section by describing the compulsory service for doctors in Turkey. I then explain the details of the mechanism that assigns doctors to their workplaces. A discussion of the data from the match simulation platform and other sources follows. Finally, I present evidence on the overlap between platform predicted probabilities and the real match outcomes.

2.1 Overview of the Compulsory Service and the Mechanism

2.1.1 *The Compulsory Medical Service*

Doctors in Turkey are required to work in a program known as the compulsory medical service when they finish medical school, and they cannot work as medical doctors in the country unless they complete this requirement. The program aims to reduce doctor shortages in underserved areas, and has been in place under various legal frames since the foundation of Turkey in 1923.⁷

The duties of a doctor during this program vary, and depend on the healthcare institution the doctor is assigned to. Doctors assigned to hospitals commonly work in emergency rooms. Doctors can be matched to other kinds of healthcare institutions as well. For instance, they can be assigned to military clinics where their primary responsibilities are providing basic healthcare services to military staff and coordinating referrals for more advanced cases. For male doctors, assignments to military clinics can also require them to work near conflict zones for temporary periods. Appendix Section A.1 gives further details on potential tasks and workplaces during this service.

The official duration of the service depends on the development level of the district where the healthcare institution is located, and can range between 300 and 600 days. Doctors assigned to more developed areas need to serve longer. It is possible to leave the compulsory service program earlier: There is a centralized residency placement exam in Turkey that takes place in the Fall and Spring of each year, and the doctors who take the centralized residency placement exam and enroll in a residency program can leave the compulsory service to start their training. However, doctors will again be required to work in underserved areas when they finish their residency training. Moreover, it is common for doctors to keep preparing for this exam during their compulsory service work, so taking the exam earlier can jeopardize their placement chances at competitive residency programs. Doctors who are dissatisfied with their placement outcome can also consider resigning, but they will have to re-enter the match and complete their service requirement to be able to work as medical doctors, and they will have to wait at least six months before being able to re-enter. These options to enroll in a residency program and to resign give doctors some control over their service duration. The survey evidence in Section 3.1 will describe how doctors' expected length

⁷Similar programs existed even before the foundation of the Republic (İlikan-Rasimoğlu 2013).

of service varies as a function of their placement outcome.

Medical graduates are matched to their compulsory service hospitals through a centralized match that occurs nearly every two months. The rounds that take place around September are the most popular ones as they follow summer graduations. The next subsection describes the official assignment mechanism that the government uses for this match.

2.1.2 *The Mechanism*

The assignment mechanism follows a procedure that can be described as Immediate Acceptance (i.e. the Boston mechanism) with random priorities in each round. Doctors submit a rank-order list that can include up to ten hospitals. The mechanism then proceeds as follows:

- Step 1: The algorithm tries to match doctors to their first choice. Each hospital h can accept up to q_h^1 doctors among candidates who rank it as their first choice. If the number of candidates ranking it as their top choice exceeds q_h^1 , everyone has an equal chance of being admitted. Unassigned doctors and remaining seats move to Step 2.
- Step $k > 1$: Doctors unassigned in step $k - 1$ apply to their k th choice. Hospital h accepts applicants up to its remaining q_h^k seats, and rejects the rest.
- After Step 10, there is a scramble round where previously unassigned doctors and remaining seats are matched randomly.

The available seats are typically announced a few days before the match. After this announcement, doctors submit their rank-order lists on a government website before the deadline. A unique feature of this setting is the existence of a private match simulation firm that has nearly all medical graduates enrolled in each match. I describe data obtained from this platform and other data sources in the next section.

2.2 **Data**

2.2.1 *Simulation Platform*

A private match simulation platform that was established in 2021 has been simulating the compulsory service match since then.⁸ The Ministry of Health typically announces the available hospitals and quotas a few days before the official match. After this announcement, nearly all doctors entering the match submit their tentative lists to the simulation platform. The platform takes these lists as its inputs, and runs Monte Carlo simulations in the background. It tells doctors

⁸Upon its establishment, the platform quickly replaced a competitor that provided simpler information such as the number of doctors listing an institution rather than providing probabilities.

their placement chances to all hospitals on their tentative lists. The simulations run every few seconds, and whenever the doctors (or their competitors) update their lists, they will learn about their new placement chances almost instantly. The platform also provides search tools that help doctors find hospitals with high placement chances for specific positions on their lists. Appendix Figure 1 shows screenshots from the platform.

The data I have in this paper come from the rounds of the match that took place in March 2025 and in September 2025. There were 1078 doctors in the match that took place in March 2025 who were matched to 529 hospitals. The simulations on the platform started on February 27 shortly after the announcement of the available seats, and ended on March 3, a few hours before the official deadline for submitting the list. The match that took place in September was much larger in size: 13056 doctors were matched to 2290 hospitals. The simulations started on August 28, again after the seats were announced, and ended on September 2, hours before the official deadline. Due to its large size, in the main analysis, I focus mostly on the match that took place in September.

I obtained four types of data from the platform. The first one is the real-time rank-order list and assignment probability data. These data include, for each user, their rank-order list on the platform and their assignment probabilities to each option on their list (as well as their probability of entering the scramble round) for every minute throughout the simulation process.

In addition to these minute-by-minute rank-order list and assignment probability data, I also administered original surveys on the platform during both rounds of the match. The surveys included questions about the preferences of the users (such as what their top two choices would be if they could choose freely from all hospitals, what their top two choices would be if they did not have access to the platform, the factors they consider to be the most important when making their choices), and their future plans in relation to outcomes of the match (how long they plan to serve in the compulsory service program, and when they expect to start a residency program if they are placed to their top choice versus if they are placed through the scramble).

The third piece of data, which I have for the round of the match that took place in September 2025 is the activity log of users on the platform. In addition to what can be obtained already from the preference data (e.g. whether users make a change to their lists between certain minutes), these data allow me to have additional information about the activity of the users, such as whether they loaded the page of all available hospitals (but not which hospital they might have clicked on afterwards) or whether they were checking their assignment probabilities.

Finally, I obtained information on the characteristics of hospitals (e.g. their type and location), demographics of the users, and importantly, the final placement outcome of the users. The last one includes which hospital they were placed into, and whether that hospital was in their rank-order list or whether they were placed through the scramble round. In addition, there is another platform owned by the same firm for simulating the residency match, and as described in Andirin (2026),

most doctors who enroll in a residency program use that website before the actual residency match. I obtain the residency exam score for the compulsory service candidates who have an account on this second platform as well.

2.2.2 *Other Sources*

I also use publicly available data on match outcomes published by the Ministry of Health. The first kind of such data are the video recordings of the March and September 2025 match rounds. The videos display the names of doctors, the hospitals they are matched to, and the round in which each match took place. These data cannot be directly merged into mine since I do not have the names of doctors, so I mainly use these recordings to get a count of the doctors matched to each hospital in each round.

In addition to these video recordings, the Ministry of Health publishes the names of doctors, the hospitals they were matched to, and whether they were placed through the scramble round for all rounds of the match. In Section 3.1, I use these data to build initial proxies for the desirability of different locations and hospital types by looking at whether certain locations and hospital types are more likely to fill their seats through the scramble round. I also use these data to track the share of different types of hospitals over time.

Turkey’s Council of Higher Education publishes the average scores for the clinical part of the residency entrance exam. I obtain these scores and use them to build a measure of the quality of medical schools, which is later incorporated into preference heterogeneity in the model presented in Section 4.

I use data from Turkey’s Statistical Institute on the number of doctors in each city. I combine this information with the number of doctors appointed by the compulsory service program to measure the share of doctors appointed by this program in a city. I also get data from the Ministry of Industry and Technology to characterize the development levels of different cities across the country (Acar et al. 2019). In Section 3.1, I combine data on the share of doctors appointed by the compulsory service and data on development levels to show the relationship between the two. In Section 4, I use development levels of locations as a factor influencing doctors’ choices of hospitals.

2.2.3 *Overlap Between The Simulation Platform and Official Results*

The platform predictions about match probabilities need to be accurate to be useful for guiding doctors’ real choices. Two conditions are necessary to ensure the accuracy. The first one requires a large share of doctors who are entering the match to use this platform. If a large share of match participants are outsiders to the platform, then the choices of those on the platform would have

little value in predicting the real outcomes. The second condition requires users of the platform to be largely truthful in their stated preferences. If the platform users make completely different choices on the official match website, again, there would be little value to using the platform. In this subsection I show that both conditions are satisfied, and as a result, the platform-predicted match probabilities track real match shares almost one-to-one.

The first condition was satisfied for both rounds of the match for which I have data. Among all doctors who entered the match, 96.7% formed a rank-order list on the platform in the match that took place in March, and 98% of them formed a list in the September match.

To show that doctors' choices on the platform and their official lists are largely aligned, I rely on the final placement outcomes showing which hospitals the doctors were eventually placed into as well as the number of doctors appointed to a hospital in a given round. Although I cannot see what rank-order lists the doctors eventually submitted to the official website, knowing these final outcomes helps me to show that the users largely stick with their choices on the platform when submitting their official lists.

I first look at the accuracy of "certain placement predictions". In March 2025, the platform predicted that 38.9% of the doctors would be placed to a given hospital with 100% chance. 97.5% of these doctors were indeed placed into platform-predicted hospitals. In September 2025, the share of doctors for whom the platform predicted a certain placement stood at 25.7%, and these predictions were accurate for 94.8% of these doctors.

I then look at the overlap between the average probability of being assigned through the scramble round predicted by the platform, and the actual share of doctors who were assigned through this round in reality. In March, the average scramble round probability predicted by the platform was 7.1%, slightly lower than the actual share of doctors who were assigned through this round (7.3%). In September, the predicted share was 4.8%, slightly higher than the actual share, which was 4.2%. It will become clear in Section 3.7 that these discrepancies are small, especially compared to the change between the average scramble round probability at the beginning versus the end of the simulation process.

The small differences in platform-predicted and actual outcomes could be caused either by users who deviate from the final lists they had on the platform by submitting a different rank-order list to the official website at the end, or due to "outsiders" who do not have lists on the platform but submit their choices to the government website before the deadline. While it is not possible to get a full count of deviations from the lists on the platform, I can still detect a subset of those using the final outcomes. In March, 1.4% of users were placed in a hospital that they have not listed on the platform, and they were not assigned through the scramble round, suggesting that their official lists included hospitals different from the final lists on the platform. This share stood at 2.4% in September. Adding new hospitals is not the only type of deviation that can occur in practice. For

instance, users can also alter the ordering of the options between their final lists on the platform and the official lists. I estimate the share of such users to be at least 1.8% in March and 2.5% in September.⁹

These numbers suggest that users largely stick with their final choices on the platform, and any deviations are likely to be small. If these deviations are indeed negligible, then the platform-predicted probabilities should reflect true placement chances closely, which Figure 1 shows indeed to be true. In this figure, for the large majority of doctors who do not have one of the detectable deviations described above, I group their match probabilities to hospitals on their lists into equal-sized bins, and show the average probability of match across bins, and what share of these matches took place in reality. The slopes and standard errors depicted in both panels are from a regression of an indicator for whether a doctor-listed hospital match took place on the predicted probability of the match. The figure clearly shows that platform-predicted probabilities at the end of simulations track real match probabilities almost one-to-one.

Overall, the results from this section show that the platform provides accurate predictions of match probabilities which are reliable for guiding doctors' choices. This is made possible by, and plausibly also the reason for, the fact that the doctors largely stick with their final choices on the platform when submitting their official rank-order lists.

3 Descriptive Evidence

I start this section by presenting descriptive evidence on what the compulsory service match means for the cities that rely on it and for doctors. I then show that the doctors spend the entire simulation window actively, and most of their activity is driven by information about probabilities. One of the survey questions asked doctors their true ideal hospitals. I show that doctors indeed rank these hospitals as their top choice when they are guaranteed placement, but they start choosing other hospitals when the placement at their ideal hospitals is not guaranteed even when these ideal hospitals can be attained at relatively high chances (e.g. 70%). I use doctors' initial choices on the platform as a proxy for the counterfactual where the platform is absent. I corroborate the suitability of this approach by relying on a survey about counterfactual top two choices. This exercise shows that the platform lowers average non-placement probability and increases placement at high-ranked options, but these gains are accompanied by doctors switching to less desirable options over time.

⁹Since I know the number of doctors placed into a hospital in a given round, I can identify a subset of doctors who made an ordering change. For instance, if a doctor was placed to the hospital she ranked as her third choice on the platform, but that hospital received no doctors in round three in reality, the doctor must have shifted the place of the hospitals in her list.

3.1 A High-Stakes Match for Cities and Doctors

The compulsory medical service program exists because many cities in less developed areas of Turkey need this program as a source of doctor supply. Panel A of Figure 2 shows, for each city in Turkey, the share of doctors appointed by the compulsory service program among all doctors serving in the city, averaged over 2018-2023. Panel B of the same figure shows the development levels for each city in Turkey based on the development index of the Ministry of Industry and Technology. The two plots make it clear why this program is critical for the residents of the less developed cities in the country.

Doctors also care a lot about the outcomes of the match. Locations differ substantially in terms of how desirable they are to the doctors participating in the match. As a proxy for the demand for a location, I take the share of doctors assigned to a city by the scramble round of the match, as a fraction of all doctors assigned through the compulsory service program between 2018 and 2023. Panel C of Figure 2 shows that certain cities, especially those in the East of the country, receive a high share of their compulsory service doctors through the scramble round. That pattern also explains why cities in the East need the compulsory service program the most as shown in Panel A. Most of their doctors come from the compulsory service, because doctors rarely want to work in these areas when they are free to choose. The types of medical institutions also affect doctors' choices. As discussed in Section 2.1.1, the duties of doctors can differ substantially based on the institution they are assigned to. It is plausible to think that certain options such as military clinics can be less desirable, especially when they are located near the border regions. Panel D of Figure 2 shows that the share of seats filled by the scramble round goes even higher when attention is restricted to military clinics, and this is especially true for the East and Southeast of the country.

Surveys conducted on the platform during the March and September 2025 rounds of the match show that doctors' plans about their careers depend heavily on whether they are assigned to an institution they prefer versus to an institution they would rather avoid. In the March survey I asked doctors how long they expect to serve in the compulsory service program if they are assigned to their top choice versus if they are assigned through the scramble round. The plot on the left in Panel A of Figure 3 shows the distribution of responses to the first question. There is considerable variation in how long doctors expect to serve, but most anticipate serving approximately one year or more. The plot on the right in the same panel of Figure 3 shows the distribution of answers to the second question. On average, doctors expect to serve much shorter periods if they are assigned through the scramble round, and a sizeable fraction state that they would resign immediately. As explained in Section 2.1.1, a doctor who resigns is still required to complete her compulsory service but she has to wait at least six months before being able to re-enter the match. Combined with this survey evidence, this shift means that being assigned to an option that the doctor does not prefer can result in significant disruptions to her career, and potentially a significant period without income.

The impact on doctors' careers is not limited to disruptions due to the potential of resignation. Doctors' plans about when and which residency program to start can also be affected by their match outcomes. As explained in Section 2.1.1, doctors can leave the compulsory service program early if they start their residency training, but they need to serve again at the end of the residency. This gives doctors an option to quit a place they do not like. However, it is a costly way out, because it is common for doctors to prepare for the residency exam during the compulsory service and take it several times, and starting a residency program sooner can mean entering into a less competitive specialty.

In the September 2025 survey, I asked doctors when they intended to start a residency program. The match in September 2025 was right after the centralized residency exam, so doctors graduating in the summer could, in principle, start the residency program in that placement cycle. Alternatively, they could wait for future exam and placement cycles. The plot on the left in Panel B of Figure 3 shows when doctors expect to start a residency program if they are placed to their top choice. Nearly half of the doctors who took the exam prefer waiting for future placement cycles. The plot on the right in the same panel of Figure 3 shows the expected start times if they are placed through the scramble round. More doctors state that they would start the residency program in the current cycle instead of delaying it if they are assigned through the scramble round. This means that being assigned to an option that they do not prefer can force doctors to start their residency training sooner, likely in less competitive programs that they would not choose otherwise.

To sum up, the above facts show that the compulsory service program is a high-stakes match both for the residents of underserved areas and for the doctors. As I will show in the next section, the importance of this match for doctors' careers translates into high and continuous activity on the platform.

3.2 Continuous Activity and Updating

Because the outcome of the compulsory service match can have an important impact on doctors' careers, most doctors spend a significant amount of time on the platform throughout the simulation window. Panel A of Figure 4 shows, for every hour of the simulation, the share of doctors who are active on the platform based on the log files of activity. At any given daytime hour, a large share of doctors are active, and nearly all doctors are online towards the end of the simulation process. Panel B of the same figure shows the distribution of the number of active hours among all doctors, and makes it clear that the continuous activity in Panel A is not driven by a small group of outliers. The median number of active hours is 39 across all doctors, suggesting that an average doctor checks the platform nearly once every two daytime hours during this simulation process that lasts around four and a half days.

An important fraction of this activity on the platform is dedicated to experimenting with alter-

native lists. Panel C of Figure 4 shows, for every hour of the simulation, the share of doctors who make at least one change to their lists. Panel D of the same figure shows that there is a long right-tail to the distribution of the number of lists tried, but nevertheless, the median doctor experiments with many alternative configurations of hospitals.

Although the doctors spend the entire simulation window actively and tend to experiment with alternatives until the end, it is worth noting that some of these experiments are less substantial than others, and the more substantial changes take place earlier. In Appendix Figure 2, I order lists by the amount of time they were held during the user's active hours on the platform, where the highest duration rank refers to the list the user kept for the longest, and I show the distribution of active hours for lists at each duration rank. The figure shows that relatively few lists were held for several hours, suggesting that the doctors modify certain lists very quickly after learning about the placement chances that they offer. To show that later changes to rank-order lists on the platform have smaller impacts on placement probabilities compared to earlier changes, for each user, I compare their initial and final lists during every hour on the platform, and find the maximum placement probability to the positions that changed between the start and end of the hour. Appendix Figure 3 shows the median of the maximum placement probability affected by changes, and makes it clear that later changes concern positions with lower placement chances.

3.3 Information About Probabilities Drives Changes

Section 3.2 showed that the doctors continuously update their lists on the platform. This section will show that most of this activity is driven by the information about placement chances rather than learning more about the alternatives or their characteristics.

In the survey conducted during the September 2025 round of the match, I asked doctors if they made a change to their initial top-ranked hospital on the platform, and if so, what the reason for this change was. Panel A of Figure 5 shows that a large majority of users stick with their initial choices, and among those who make a modification, the biggest reason is the change in assignment probabilities.

Panels B and C of the same figure provide evidence showing that the changes in later hours are also driven by information about probabilities. Panel B shows the responses to a survey question where I asked doctors a day before the deadline if they were still considering making changes to their lists, or if they regarded their lists as finalized and were not intending to change them unless the probabilities changed. A large majority stated that they were not intending to make changes unless the probabilities changed. After the end of the simulations on the platform, I asked doctors whether they made a change to their lists in the final hours, and if so, what the reason for this change was. As can be seen in Panel C, a large share of the doctors stated probability changes as the reason.

Together, these three plots suggest that most of the changes to the lists on the platform are caused by probabilities, and there is limited scope for learning about the hospitals or their characteristics in explaining these changes. This pattern is natural in a setting where all the available hospital seats are announced just a few days before the deadline and there is not enough time to learn about the fine details of alternatives. When I asked the doctors about the main factors driving their hospital choices on the platform, an overwhelming majority stated distance as can be seen in Appendix Figure 4, and almost all doctors mention easily observable characteristics that depend on the hospital’s type and location, which also suggests that their choices would be unlikely to be affected by information about options they discover later. Overall, the evidence in this section describes a choice environment where the doctors are clear about their preferences, and changes to the list-level choices are driven by assignment prospects.

3.4 Responses to Probability Changes are Fast

The preceding section showed the importance of the information about probabilities in explaining changes to rank-order lists. This section will demonstrate that doctors react to changes in probabilities quickly.

As described in Section 2.2, the data I have consist of minute-by-minute snapshots of users’ preferences and placement probabilities to each option on their rank-order lists. As a result, I can distinguish between probability changes caused by actions taken by the user herself versus changes caused by other users.

I construct an hourly measure of the intensity of “external” probability changes by summing across the changes in non-placement probability into the user’s listed choices (and hence entering the scramble round) between each snapshot caused by other users. More formally, I construct the variable:

$$\Delta P(Scramble)_{it}^{ex} = \sum_{m \in \mathcal{M}_t} \Delta P(Scramble)_{im} ExtChange_{im}$$

where i indexes users, t refers to hours, m indexes minutes, $\mathcal{M}_t$ refers to the set of minutes in hour t , $\Delta P(Scramble)_{im}$ is the change in the probability of not being placed to an option in the user’s rank-order list (and hence entering the scramble round) between minutes m and $m - 1$, and $ExtChange_{im}$ is an indicator equal to one if the user’s own rank-order list stayed the same between m and $m - 1$.

I then estimate event-study equations of the following form:

$$\Delta Choice_{it} = \alpha_i + \gamma_t + \sum_{k=-4}^5 \beta_k \Delta P(Scramble)_{i,t-k}^{ex} + \varepsilon_{it} \quad (1)$$

where $\Delta Choice_{it}$ refers to whether the user made any modifications to her rank-order list during the hour t , α_i is a doctor fixed-effect, and γ_t refers to hour fixed-effects. Figure 6 displays the results from estimates of equation 1. The figure shows that the users respond immediately to shifts in their probability of going to the scramble round triggered by the actions of other users.

The dependent variable in equation 1 pools across all changes, including reorderings and substitutions of the options on the existing rank-order lists with newer ones. In Appendix Figure 5, I instead use the change in list length between the ends of hours t and $t - 1$. The figure shows that users make their lists lengthier following increases in the scramble round probability driven by the actions of other users, which can be explained by the addition of more safety options. In Appendix Figure 6, I replace the dependent variable with an indicator for whether the user engaged in any search activity in a given hour. The figure shows that the users engage in more search in response to increases in scramble round probability caused by the actions of other users.

These results show that it does not take very long before the users update their lists in response to probability changes. The findings also suggest that factors such as inertia are unlikely to be important in this setting.

3.5 Doctors Trade Their Favorite Hospital for Placement Certainty

The preceding two sections made it clear that the doctors care a lot about their assignment prospects. This section will show what kind of probabilities they search for. Combining survey responses and the lists on the platform, I find that doctors seek options that give them certain placement chances. Whether this is consistent with expected utility or is explained better by a model that puts a premium on absolute certainty will be answered when the model with probability weighting is introduced later.

In the September 2025 round of the match, I asked users which two hospitals they would rank as their top two choices if they could choose freely without being worried about placement prospects. I call the first hospital that a doctor names her ideal hospital. Panel A of Figure 7 shows the share of survey respondents whose stated ideal choices are the same as their final top-ranked choices on the platform, and the share of respondents whose ideal hospital and final top-ranked choice differ. For the majority of users, their stated ideal hospital appears as their top-ranked choice on final rank-order lists, but for a large share around 40%, the final top-ranked hospital is different from the stated ideal hospital.

Panel B of the same figure shows the share of users who stick with their ideal hospital as their top-ranked hospital on the platform as a function of the probability of being placed in the ideal hospital if it is ranked first. One can see that as the placement probability approaches certainty, nearly all doctors (96.0%) stick with their stated ideal choice.¹⁰ However, a large share of doctors drop

¹⁰The monotonic relation in this plot and almost all users choosing their ideal hospitals when they are guaranteed

their ideal choices even if they are reachable at fairly large probabilities. In other words, a substantial share of doctors reject having final rank-order lists where their ideal hospital is attainable at a high probability (e.g. 70%) followed by safety hospitals.

Some doctors start with their ideal hospital as their top-ranked choice on the platform, and switch to other alternatives later during the simulations. It is useful to ask what their placement probabilities to the top-ranked option look like around the time of the switch to get a sense of the magnitude of the gain they experience and to understand whether these changes are driven by shocks.

Figure 8 shows results from event-study equations of the form:

$$TopChoiceProb_{it} = \phi_i + \tau_t + \sum_{\substack{k=-4 \\ k \neq -1}}^5 \delta_k \Delta z_{i,t-k} + \delta_6 z_{i,t-6} + \delta_{-5} (1 - z_{i,t+4}) + \xi_{it} \quad (2)$$

where $\Delta z_{i,t-k}$ is an indicator equal to 1 for user-hours where a user i who started initially with her ideal hospital as her top choice switched to a different hospital exactly k periods before time t , and 0 otherwise. The variable $z_{i,t-6}$ is an indicator equal to 1 for user-hours where user i has already made a change to her top choice at least six periods before. The term $(1 - z_{i,t+4})$ is an indicator for user-hours where user i has not yet made a change and will not do so for at least four more periods. These two endpoint indicators are included to fully saturate the equation, following Freyaldenhoven et al. (2024). ϕ_i is a doctor fixed-effect, and τ_t refers to time fixed-effects. Plotted in Figure 8 are the coefficients δ_k .

The figure shows that the users who start off with their ideal hospitals as their top choice and later switch to another choice had an average placement probability to their ideal hospital equal to 65% at the time of the switch, and by switching, they increase their probability of placement at the top choice by around 20 percentage points. However, it is also worth emphasizing that these users had, on average, relatively high probabilities of being placed to the top option on their lists to begin with. The figure also shows that not all changes are explained by probability shifts: The placement probabilities to the top-ranked option are stable on average before the switch, suggesting that these doctors found options with higher placement chances absent a shock to the placement probability to their ideal hospitals. Later in the paper, I will compare the performance of the expected utility and probability weighting models in explaining the switches from ideal hospitals to other alternatives.

11

placement there also provide evidence of largely truthful responses in the surveys.

¹¹It is also useful to ask whether one can learn about violations of first-order stochastic dominance from the cases where doctors drop their ideal hospitals. For 16% of the users, at the time they dropped their ideal hospital from their first choice for a new alternative, that alternative was a sure second choice. In other words, if they had kept their ideal hospitals as their first choice and ranked the hospital they later switched to as their second choice, they would have had guaranteed placement at that hospital in the second round if they had not been placed in the first

Panels A to D of Appendix Figure 7 show the distribution of cumulative placement probabilities to different options in rank-order lists at different snapshots. The patterns in the figure are consistent with the findings from the previous two figures: Over time, the cumulative probabilities of placement in top choices shift upward, suggesting that doctors gradually switch to top choices that offer better placement prospects. As expected due to the Boston mechanism, most doctors have significant chances of being placed to one of their top choices throughout the simulation process since the mechanism gives higher priority to applicants who rank an option high compared to those who rank it lower.

3.6 How Does the Platform Affect Choices?

The preceding sections established that information about probabilities drives choices on the platform, and that doctors gradually switch to choices that give them certain placement chances. One might think that doctors' initial choices were too optimistic, and if they had to choose without the platform, they would pick alternatives with higher placement chances on their own. In other words, the platform could be just giving them a chance to experiment with their ideal hospitals, but if the platform did not exist, their counterfactual choices would be similar to what they eventually choose on the platform.

The key challenge to answering whether and how the platform affects choices is the lack of data on what the counterfactual choices of each user would be absent the platform. To tackle this challenge, in the September 2025 round of the match, I asked doctors which two hospitals they would rank as their top two choices if they did not have access to the platform. This question obviously has its own caveats: In a setting with 2290 hospitals where the platform starts running immediately after the seats are announced, the doctors might not know the exact hospitals they would choose in that counterfactual scenario even if they make their best effort in answering the survey question. However, these survey responses can still give a sense of what placement chances the doctors would face if they were making choices without the platform. Another caveat of these responses is that they are only defined for the subset of users who responded to the survey, and only for the top two choices, unlike the initial choices on the platform which are defined for all users by definition. In the rest of this section I show how the stated counterfactual choices, initial choices on the platform, and final choices relate to each other.

Out of the 10.6% of doctors who responded to the survey question on their counterfactual top two choices, 37% indicated a top-ranked hospital that was different from their final choices. For these doctors, if they stuck with their counterfactual top-ranked hospital as their final top-ranked hospital, then their placement chance to these hospitals would be 65% on average, as can be seen

round. This shows that first-order stochastic dominance violations are present for a non-negligible share of users.

in the first bar of the plot on the left in Panel A of Figure 9. On the other hand, their placement chance to their final top-ranked hospital on average was around 87% as can be seen in the second bar of the same plot. These numbers suggest that, for doctors whose stated counterfactual and final top-ranked hospitals differ, the final top-ranked hospitals offer a greater placement chance compared to stated counterfactual top-ranked hospitals.

The corresponding numbers for doctors whose initial and final top-ranked hospitals differ on the platform are very similar. 41% of the doctors on the platform had a final top-ranked hospital that was different from their initial top-ranked one. For these doctors, if they stuck with their initial top-ranked hospital, their placement chances to those hospitals would be 69% on average as can be seen in the first bar of the plot on the right in Panel A of Figure 9. In contrast, their placement chance to their final top-ranked hospital was around 85%, as can be seen in the second bar of the same plot. This similarity suggests the initial and counterfactual top-ranked hospitals are very similar in terms of the placement chances they offer.

Looking at the counterfactual, initial, and final second-ranked hospitals can also be helpful to get a sense of how the platform affects the “safety” options of the doctors. A second-ranked hospital is not a good safety option if it is predicted to be filled in the first round of the match. For 88% of survey respondents, the stated and final second-ranked hospitals differ. Among these, 86% of the counterfactual second-ranked hospitals were bad safety options because they were predicted to be filled in the first round, as can be seen in the first bar of the plot on the left in Panel B of Figure 9. Only 4% of the final second-ranked hospitals, on the other hand, were bad safety options. The platform therefore helps doctors find far safer second choices.

A striking similarity emerges again when one looks at the corresponding numbers for the doctors whose initial and final second-ranked hospitals on the platform differ. For 88% of the doctors, their initial and final second-ranked hospitals differ. Out of these initial second-ranked hospitals, 84% of them were bad safety options, compared with just 5% of the final second-ranked hospitals, as can be seen in the plot on the right in Panel B of Figure 9.

Overall, these numbers suggest that the platform leads doctors to pick first choices with higher placement chances, and find valid safety options. They also show that using initial choices as a proxy for the counterfactual choices can be reasonable, as the initial choices on the platform and the stated counterfactual choices exhibit very similar placement chances.

3.7 Aggregate Outcomes Without the Platform

The previous section showed that the platform helps users to find hospitals with higher placement chances for the top-ranked position on their lists, and also steers them away from unrealistic second choices. This section asks what the outcomes would look like if the platform did not exist, i.e., if nobody had access to it.

The challenge in this section is the same as in the previous one: Counterfactual lists in the absence of the platform are unknown. Figure 10 shows that the average scramble round probability keeps declining throughout the process, but to fully disentangle the effects of the platform from what could have happened even without it, one needs to know the counterfactual lists that the doctors would submit if the platform did not exist.

Motivated by the similarity between the initial and stated counterfactual choices shown in the previous section, to evaluate the counterfactual where the platform did not exist, I use doctors' initial choices on the platform to construct alternative rank-order lists as this gives me well-defined lists for all users unlike the counterfactual survey which is limited to survey respondents. That is, for each user i , I construct an alternative rank-order list, where the position k on that user's rank-order list is fixed to the initial alternative that was tried for that position. To give an example, if a user initially ranked Hospital A as the second alternative on her list, and later switched to Hospital B, I fix the second spot on her counterfactual list to Hospital A. I do this for all users, and run the match repeatedly to calculate the assignment probabilities under these alternative lists.¹²

Although one cannot draw welfare conclusions by just looking at how assignment probabilities change between the counterfactual and final lists, it is possible to evaluate certain outcomes of interest such as the average probability of being assigned through the scramble round or the average probability of assignment to one of the top two choices even without a model. Table 1 shows how these two outcomes change under information provision relative to a scenario where all users were stuck with their initial choices.

There is a large increase in the average probability of going to the scramble round when the match is simulated using the counterfactual lists rather than the final ones, from 4.8% to 8%. However, the gains in terms of avoiding the scramble round are concentrated among a minority of the users: for around 10% of users the probability of going to the scramble round would rise by 19.2 percentage points or more in the counterfactual.

The next two rows focus on the change in the probability of being assigned to top two choices using the alternative versus actual lists. The platform leads to large increases in the probability of being assigned to one of top two choices on average. These gains are less concentrated than the reductions in the probability of going to the scramble round. For around 30% of the users,

¹²One might think that it might be reasonable to use time-series to overcome this challenge, by comparing the placement outcomes before and after the launching of the platform. There are two limitations to this approach. The first one is that there was an alternative platform before, which provided simpler information such as the number of doctors ranking an alternative as their top choice, which makes a clean before-after comparison difficult. The second limitation is the wide variation in the numbers and characteristics of available hospitals in each match. Appendix Figure 8 shows, in Panel A, the share of doctors who are assigned through the scramble round in each round of the match. Panel B shows the share of seats affiliated with the military, as a proxy for the fraction of less desirable seats in a match. Panel C shows that the placement outcomes, as captured by the share of applicants going to the scramble round, are clearly correlated with what is offered.

these gains are very large, corresponding to an increase of at least 21.6 percentage points in the probability of top two placement.

An advantage of using initial lists to create counterfactual lists is that they allow me to define complete lists for all users. Using the responses to the survey question on counterfactual choices, on the other hand, has the caveat that it is only defined for the top two choices of a subset of doctors who responded to the survey. To make use of these responses, in an extension to the above analysis, I use the stated counterfactual top two choices in place of the initial top two choices for survey respondents, while keeping the other positions on their lists and the lists of other doctors fixed to the initial choices. The increase in the probability of being assigned to one of the top two choices for survey respondents is 15.8 percentage points, similar to what they would experience when their counterfactual choices are assumed to be their initial choices (13.2 percentage points).

These results show that the platform leads to improvements in certain outcomes at an aggregate level: The average probability of going to the scramble round goes down, and the average placement probability to one of the top two choices increases. However, these improvements do not come for free. Figure 11 shows that users gradually switch to options that are further away from them and located in lower-quality cities.¹³ These switches to less desirable options make the welfare gains from information an open question. If these declines in the desirability of the chosen options come from an overweighting of certainty rather than maximizing expected utility, the platform's effects on welfare in an expected utility sense will be ambiguous. The following sections will answer whether information provision is welfare improving even if agents are allowed to depart from expected utility maximization.

4 Rank-Order List Formation with Probability Weighting

The descriptive analysis in Section 3 revealed that doctors prefer hospitals that give them certainty even when their favorite hospitals could be reached at relatively high chances. This behavior can be consistent with expected utility maximization, but it can also emerge from an alternative model where doctors behave as if they maximize an objective function that puts an extra weight on certainty. If doctors indeed depart from expected utility maximization, then information provision can drive them towards different objectives, and they might do better under a mechanism that does not require them to work with probabilities. Thus, understanding the objectives of agents in matching markets is important for evaluating the welfare impacts of information provision and of alternative assignment mechanisms.

¹³Changes in probability-weighted average characteristics over time could partially result from changes in placement chances. To isolate the impact of users' own choices from that of probability changes that occur over time, I use final placement probabilities to individual hospitals when calculating probability-weighted average characteristics. The results are nearly identical if I use probabilities at the time a user holds a particular list.

The standard model of rank-order list formation in the school (or hospital) choice literature asserts that the agents choose the rank-order list R among the set of feasible lists $\mathcal{R}$ to maximize

$$\sum_{h \in R} p_h^R u_{ih}$$

where p_h^R refers to assignment probability to option h (e.g. a hospital or a school) if the agent submits list R , and u_{ih} is the agent i 's utility from the option h . In words, the standard assumption is to assert that the candidates are expected utility maximizers.

In this section I present a model that allows doctors to weight probabilities non-linearly when they form their rank-order lists. Specifically, I allow doctors to maximize another weighted sum of utilities, where the weights are $w(p)$, without asserting that $w(p) = p$:

$$\sum_{h \in R} w(p_h^R) u_{ih}$$

The model nests expected utility maximization, as it allows doctors to weight probabilities linearly as well. I outline the full empirical model below, and discuss its identification and estimation. In later sections, I use this model to evaluate welfare gains from information provision and from a switch to a strategy-proof alternative, and extend the model to allow for heterogeneity in probability weighting.

4.1 Model

I model the utility that doctor i gets from a hospital h at period t as follows:

$$u_{iht} = \beta_i^x X_h + \beta_i^d \text{distance}_{ih} + \varepsilon_{iht}$$

Here, X_h includes hospital characteristics (the type and the geographical region of the hospital) and the development percentile of the city where the hospital is located.¹⁴ distance_{ih} denotes the distance between the city of the hospital and the city where doctor i graduated from medical school measured in hundreds of kilometers. I assume ε_{iht} is distributed i.i.d. Type-I extreme value across individuals i , hospitals h , and time periods t .¹⁵ The outside option corresponds to not being placed in one of the options on the rank-order list. Its utility is denoted u_{i0} , and I normalize the deterministic component of it to zero.

Preferences for hospital characteristics depend on the gender of the doctor and the quality of her medical school, as measured by the average residency exam score of its graduates. Specifically,

¹⁴The development percentile is scaled to be between zero and one, with one corresponding to the least developed city.

¹⁵These time-variant hospital-level shocks allow for unobservable changes in tastes over time, but rule out permanent unobserved preferences for hospitals.

for a given characteristic k , I assume that:

$$\beta_i^k = \beta_{g(i)}^k + \beta_q^k q(i)$$

Here, $g(i)$ refers to the gender of the doctor, and $q(i)$ is the average residency exam score of the medical school the doctor graduated from, standardized to have zero mean and unit variance across doctors. (Appendix Figure 9 shows the distribution of school quality across Turkey. Average school quality has a correlation of -0.69 with the development tiers shown in Panel B of Figure 2.)

I model the utility from a rank-order list submitted to the platform as:

$$U_{it}^R = \sum_{h \in R} w(p_{h,t}^R) u_{iht} + \varepsilon_{it}^R$$

$p_{h,t}^R$ is the probability of being assigned to hospital h at time t if the rank-order list R is submitted, and $w(p)$ is the probability weighting function. If $w(p) = p$, we are back in the expected utility framework.¹⁶

In addition to the probability distortions, I also introduce a list-level shock ε_{it}^R that is distributed Type-I extreme value with variance $\kappa_{list}^2 \frac{\pi^2}{6}$ i.i.d. across individuals, lists, and time. There are several reasons to introduce a list-level shock. As Kapor (2025) argues, there might be heuristics that doctors follow when forming their rank-order lists, so their list-level choices can deviate from what hospital-level utilities and probability weights can predict alone. One might expect that deviations from lists that optimize a weighted average of hospital-level utilities might be even larger in this context: Although the platform-predicted outcomes closely match real outcomes as I argued in Section 2.2.3, there is a small share of users whose real outcomes cannot be rationalized by the lists on the platform. As a result, it would be naive to expect all lists on the platform to reflect optimal choices by users.

For the probability distortion function, I adopt a sieve-based approach (see Chen 2007) and use Bernstein polynomials as the basis with the normalizations $w(0) = 0$ and $w(1) = 1$. This approach ensures that the standard functional forms of the prospect theory literature do not drive the results. I outline the details in Appendix Section A.2.2. In Section 8, I adopt a common parametric form for the probability weighting curve to incorporate heterogeneity.

¹⁶The probability weights in this formulation do not depend on the true or stated ranking of the hospitals by individuals. This is in line with Bernheim and Sprenger (2020) who show that probability weights do not depend on the ranking of the alternatives. Note that this formulation allows for violations of first-order stochastic dominance. As I show in Section 3.5, such violations are indeed present in the data.

4.2 Estimation and Identification

As discussed in Section 3.5, I asked survey respondents which hospitals are their true top two favorite hospitals. Let T_s denote the time at which doctor i answered the survey. I assume that her stated most preferred hospital $h_{i,1}$ and second most preferred hospital $h_{i,2}$ satisfy:

$$h_{i,1} = \arg \max_{h \in \mathcal{H}} u_{i,h,T_s}$$

$$h_{i,2} = \arg \max_{h \in \mathcal{H} \setminus \{h_{i,1}\}} u_{i,h,T_s}$$

where $\mathcal{H}$ denotes the set of all hospitals.

I also asked the users a day before the deadline whether they were still working on their lists or they were planning to stick with them unless the probabilities changed. If they chose the latter response, I assume that their lists at the time of the answer, which I call the pre-final lists, satisfy:

$$R_{i,T-1} = \arg \max_{R'} \sum_{h \in R'} w(p_{h,T-1}^{R'}) u_{ih,T-1} + \varepsilon_{i,T-1}^{R'}$$

To calculate the assignment probabilities $p_{h,T-1}^{R'}$ under alternative lists R' for user i , I set her list to be R' , fix all other users' rank-order lists to what they had on the platform at the exact time of i 's survey answer, and simulate the match. Finally, I assume that the final rank-order lists on the platform satisfy:

$$R_{i,T} = \arg \max_{R'} \sum_{h \in R'} w(p_{h,T}^{R'}) u_{ih,T} + \varepsilon_{i,T}^{R'}$$

I then construct the likelihood of the observed sequence of choices on the platform. Let $\alpha = (\beta, \kappa_{list}, \theta)$ denote the parameters of the model, where θ refers to the coefficients on the Bernstein basis polynomials used for the probability weighting function. Let D_i collect the characteristics of doctor i . Because the hospital-level shocks ε_{iht} are independent over time, the likelihood for the observed sequence of choices of doctor i if she responded to all survey questions can be written as:

$$P(h_{i,1}, h_{i,2}, R_{i,T-1}, R_{i,T} \mid \alpha, D_i) = P(h_{i,1}, h_{i,2} \mid \alpha, D_i) \cdot P(R_{i,T-1} \mid \alpha, D_i) \cdot P(R_{i,T} \mid \alpha, D_i) \quad (3)$$

The likelihoods for users without parts of the survey data can be constructed similarly. The list-level choice probabilities do not have closed-form expressions due to the presence of hospital-level shocks, and they need to be simulated. Appendix Section A.2 gives explicit expressions for each factor and details on the simulated maximum-likelihood estimation.

In practice, all of the parameters of the model are identified jointly. Here I briefly discuss the key sources of identification. Both the stated ideal choices, and the rank-order list choices are informative for pinning down the β parameters in the utility function. Repeated observations on list-level choices as well as the discrepancies between survey-indicated preferences and list-level choices allow pinning down the scale parameter κ_{list} .

The presence of preference measurements outside the list-level choices is helpful for pinning down the probability weighting function. In addition, differential responses to changes in probabilities based on the initial probabilities before change are helpful for identifying the probability weighting function: Fixing the user, the change in utility from an increase in the assignment probability to hospital h , keeping assignment chances to other hospitals constant, is approximately $w'(p_h)\Delta p_h u_h$, which is useful for pinning down the slope of $w(p)$ (and the function itself when the normalizations that $w(0) = 0$ and $w(1) = 1$ are used). Hence, changes in probabilities between the pre-final and final rank-order lists provide identifying variation. Similar to comparing the reactions to changes in placement probabilities when the initial placement chances differ, cross-sectional comparisons across individuals are also used in identifying the probability weighting function. This is because of the fact that one gets observationally equivalent individual-hospital pairs with different assignment probabilities after the utility function is parametrized.

5 Results

5.1 Coefficient Estimates

Figure 12 shows the estimated probability weighting curve. The estimated curve reveals that the doctors depart from expected utility maximization and that they weight probabilities non-linearly. Moreover, it shows a striking similarity to the inverse S-shaped pattern commonly found in the prospect theory literature: An overweighting of small probabilities followed by a nearly flat region, and an eventual steep rise as the probability approaches one.¹⁷

Appendix Table 1 presents the estimates of the model's other parameters. In the last three columns, I present the results from a version that assumes expected utility, where I constrain $w(p)$ to equal p . Under both the probability weighting and expected utility models, the coefficients on the hospital characteristics (i.e. β parameters) have the expected signs and reasonable magnitudes. Doctors derive disutility from being further away from the city where their medical school is located, they prefer working in more developed cities, and military clinics are the least preferred hospital types. Females have stronger preferences for working in more developed cities. The

¹⁷For reference, Appendix Figure 10 reproduces the estimated curve from Tversky and Kahneman (1992). Findings with curves that look very similar to the one in Tversky and Kahneman (1992) are ubiquitous in the prospect theory literature (see Barberis 2013 for a review of this literature).

relative distaste for military clinics is higher for males, consistent with the fact that males might be required to work near conflict zones outside the country when they are assigned to military clinics, lowering the desirability of these options.

However, there are some important and intuitive differences in the estimated β parameters when one accounts for departures from expected utility maximization: Under probability weighting, doctors have a stronger preference for working in more developed cities. This change in magnitude is sensible, because without probability weighting, the attractiveness of the hospitals in less developed areas is attributed to the hospitals rather than to the certain placement chances that they offer. It is also worth noting that the value of the inside options goes down under probability weighting. Without probability weighting, the model reads avoidance of even small probabilities of being assigned through the scramble round as a strong preference for being assigned to one of the ranked options. With probability weighting, some of that preference is attributed to the overweighting of small probabilities.

The results show that doctors weight probabilities when forming their rank-order lists, and this weighting reflects the core features discussed in the prospect theory literature. Moreover, accounting for probability weighting has an intuitive impact on the coefficients for hospital characteristics. In the next two sections, I compare probability weighting and expected utility models in terms of their performance in explaining the facts in the data. After evaluating the model fit, I discuss the welfare impacts of real-time information provision, and compare this “full-information” Boston mechanism to the strategy-proof random serial dictatorship. For both exercises, I discuss how predictions change when probability weighting is ignored.

5.2 Expected Utility, Probability Weighting, and Changes to the Top Choice

In Section 3.5, we have seen that the survey respondents who started with their ideal hospital as their top choice and later switched to alternative hospitals had, on average, a probability of placement around 65% to their ideal choice at the time they switched, which appears relatively large. In this section, I compare the performance of the expected utility and probability weighting models in explaining these changes.

When one compares the deterministic utility from rank-order lists before and after the switch, the expected utility model finds that only 49.8% of the users are better off with their new lists. This number rises to 69.7% with the probability weighting model, meaning that the latter does a better job in terms of rationalizing the behavior of users who drop their true top choices. This is intuitive, since a jump of 20 percentage points from 65% means a much larger increase under the estimated probability weighting curves depicted in Figure 12 than would be the case under expected utility.

Figure 13 shows the changes in the deterministic component of list-level utility, measured in kilometers, around the time of the switch through the lenses of both models. Panel A is an event-

study plot that replaces the dependent variable in Equation (2) with the deterministic component of utility based on the expected utility model, and Panel B is an event-study plot that does the same using the model with probability weighting. On average, the jump in utility is much clearer under probability weighting.

5.3 Placement Chances Under Real and Model-Predicted Choices

We learned from the analysis in Section 3.2 that the users of the platform gradually switch to first choices that give them certain placement chances. Similarly, the results in Section 3.5 showed that the users are likely to drop their ideal hospitals from their first choices unless they can be placed there with certainty. Can this preference for certain placement chances be rationalized under the expected utility model?

Figure 14 answers this question. The first bar shows the share of users who are predicted to be placed to their top choices with certainty in the actual data. The second bar shows the corresponding share predicted by the expected utility model. To calculate this share, for each user-time pair¹⁸, I calculate the probability that the user chooses a list that offers certain placement chances at a hospital by using the model-predicted list choice probabilities, and then average across user-time pairs. The third bar does the same for the probability weighting model. As the figure shows, the expected utility model has difficulty in explaining the strong preference for certainty, whereas the probability weighting model closely matches the real data.

More generally, one can ask which model is better at explaining the distribution of probabilities in chosen lists in the real data. To do so, in Appendix Figure 11, I compare the empirical CDFs of the chosen probabilities obtained from the real data and from predicted choice probabilities using the two models. I do this separately for the first, second, and third choices. To construct the empirical CDFs implied by the models, I calculate:

$$\widehat{F}_M^{(k)}(x) = \frac{1}{N} \sum_{(i,t)} \sum_{R \in \mathcal{R}_{it}} \Pr_M(R \mid \mathcal{R}_{it}) \mathbf{1}\{p_{k,t}^R \leq x\} \quad (4)$$

where $\Pr_M(R \mid \mathcal{R}_{it})$ denotes the model-implied probability that user i chooses list R from her choice set $\mathcal{R}_{it}$ under model M , $p_{k,t}^R$ denotes the placement probability at the k th position of list R at time t , and N is the total number of (user, time) choices. The empirical CDF from the actual data takes the same form, with the inner sum replaced by $\mathbf{1}\{p_{k,t}^{R_{i,t}} \leq x\}$. The figure shows that the probability weighting model again does a better job than the expected utility model in capturing the distribution of probabilities at different ranks.

¹⁸A user can have more than one chosen list in the estimation data as outlined in Section 4.2.

5.4 Estimated $w(p)$ and Changes to the Lists

In the event-study specification in equation 1, I showed that the users react to changes in their probability of being assigned through the scramble round caused by the actions of other users. Based on the estimated probability weighting curves in Figure 12, the impact of such a change should depend on the initial probability. For instance, an increase of one percentage point from 5% to 6% would seem much larger than a change of the same magnitude from 50% to 51%. In order to test this intuition using the estimated probability weighting curve, I augment equation 1 as:

$$\begin{aligned}\Delta Choice_{it} = & \alpha_i + \gamma_t + \sum_{k=-4}^5 \beta_k \Delta P(Scramble)_{i,t-k}^{ex} \\ & + \sum_{k=-4}^5 \theta_k \Delta P(Scramble)_{i,t-k}^{ex} \cdot w'(P(Scramble_{t-k-1})) + \varepsilon_{it}\end{aligned}$$

where $w'(P)$ represents the slope of the estimated probability weighting curve at P , and $Scramble_t$ refers to the probability of going to the scramble round averaged over minutes in hour t . Intuitively, changes in the probability of going to the scramble round would seem larger if the slope of the probability weighting function before the change was high. Appendix Figure 12 presents the coefficient estimates of θ from the equation above, which confirms this intuition.

6 The Value of Real-Time Information

In this section, I quantify the welfare impact of information provision about assignment probabilities. To approximate the counterfactual scenario where the applicants need to decide without getting information on probabilities from the platform, as discussed in Section 3.6, I use their initial rank-order lists, and simulate the match outcomes with these lists. I first use changes in expected utility as the key metric for evaluating welfare. Although the doctors depart from expected utility maximization, a policy-maker plausibly cares about doctors' expected placement outcomes, which speak to longer-term welfare, rather than about the decision weights doctors apply during a short choice process. However, I also provide welfare results based on the estimated probability weighting curve.

For each doctor, I compute the difference between the expected utility at the end of the simulations and expected utility in the counterfactual case where all doctors use their initial choices on the platform as described in Section 3.7:

$$\Delta E[u_i] \equiv E^A[u_i] - E^C[u_i] = \sum_{h \in R_i^A} p_{ih}^A E[u_{ih}] - \sum_{h \in R_i^C} p_{ih}^C E[u_{ih}]$$

where R_i^A refers to the actual rank-order list of doctor i at the end of the simulation process, and p_{ih}^A refers to doctor i 's placement probability to hospital h when all doctors use these actual lists. R_i^C and p_{ih}^C refer to the lists and placement probabilities in the scenario where doctors use counterfactual lists. The two right-most expectations are with respect to the distribution of idiosyncratic shocks.

I calculate the change in welfare using the parameter estimates from both the expected utility and probability weighting models. Figure 15 shows the results. Under expected utility, the increase in welfare from information provision is on average equivalent to being 68 kilometers closer to one's medical school. When the welfare is evaluated using the parameters from the probability weighting model, the average welfare gain falls to 44 kilometers, so ignoring probability weighting overstates welfare gains from information by more than fifty percent. As noted in Section 3.7, doctors gradually switch to less desirable options to improve their placement chances. The coefficient estimates presented in Section 5 showed that the probability weighting model penalizes lower-quality options more than the expected utility model, which explains the decline in welfare gains when probability weighting is taken into account.

In Appendix Table 2 I present welfare gains at different deciles of the distribution. The table shows that the changes in welfare are centered around zero: Half of the users gain from information provision, but the other half loses, under both the expected utility and the probability weighting models. This spread also highlights the value of using the model to evaluate the welfare consequences of information provision. When attention was restricted to the change in the scramble round probability as a proxy for welfare in Section 3.7, very few users appeared to be affected by the platform, and the effects were concentrated in top and bottom deciles. The model-based results in this section reveal much broader welfare consequences. The heterogeneity in probability weighting that will be introduced in Section 8 will be helpful for characterizing who gains versus loses from information provision.

As mentioned at the beginning of the section, one can also evaluate welfare based on the probability weights that the doctors use when forming their lists. For this purpose, I define the change in $w(\cdot)$ -weighted utility:

$$\sum_{h \in R_i^A} w(p_{ih}^A) E[u_{ih}] - \sum_{h \in R_i^C} w(p_{ih}^C) E[u_{ih}]$$

Panel C in Appendix Table 2 shows the distribution of changes according to this alternative measure. The average increase in welfare is larger, and the distribution is more dispersed. Gaining

certainty (or losing it due to the competition from now-informed competitors) is priced more under this metric, making the “perceived” stakes larger.

7 Boston with real-time information versus a strategy-proof mechanism

The previous section showed that real-time information provision improves the outcomes of the Boston mechanism compared to the case where the mechanism is run with less information on the side of the candidates. The next question concerns how this “upgraded” version of the Boston mechanism compares to a strategy-proof alternative.

As a point of comparison, I use random serial dictatorship. This mechanism is widely used for random allocation problems, where the objects do not have preferences over applicants, which is also the case in the setting of the compulsory medical service in Turkey.

In random serial dictatorship, the candidates (here, doctors) are randomly assigned priorities. First, the doctor with the highest priority picks her favorite hospital. Then, the doctor with the second highest priority chooses among what remains, and the mechanism proceeds in this manner until all doctors are assigned. As a result, the mechanism is strategy-proof, as no candidate can benefit from misrepresenting her true preferences. In a setting where the agents distort probabilities, the benefits of using a strategy-proof mechanism can be even larger, which further increases the value of comparisons to such mechanisms. To measure the benefits from switching to a strategy-proof alternative, I calculate the difference between expected utilities under RSD and the final lists on the platform:

$$\Delta E[u_i] \equiv E^{RSD}[u_i] - E^A[u_i] = \sum_{O \in \mathcal{O}} p^O E[u_{ih}^O] - \sum_{h \in R_i^A} p_{ih}^A E[u_{ih}] \quad (5)$$

where $\mathcal{O}$ is the set of potential orderings under RSD, and p^O is the probability with which a particular ordering occurs, which is equal to $\frac{1}{|\mathcal{O}|}$. $E[u_{ih}^O]$ refers to the expected utility the candidate receives from the hospital she is assigned to under the ordering O .

$E^{RSD}[u_i]$ is an expectation over all potential orderings of 13056 doctors, which is impossible to calculate directly. To ease the problem, I randomly order applicants 250 times and take the average across these orderings.

Although the Random Serial Dictatorship is strategy-proof when there are no restrictions on the list length, such restrictions are in place in most matching markets. As a result, it is also important to look at how more realistic versions of this mechanism compare to the Boston mechanism. To do so, I simulate the outcomes under truncated versions of the RSD as well. Under truncated RSD with list length k , if the listed k options are exhausted by the time it is an applicant’s turn, then the applicant gets assigned randomly. Following the existing literature (see, e.g., Kapor, Neilson and

Zimmerman 2020) I assume that the agents naively behave truthfully under the truncated version of this otherwise strategy-proof mechanism. Although strategy-proofness is not guaranteed after truncation, Pathak and Sönmez (2013) suggest that truncated RSD is still less manipulable than the Boston mechanism.

Figure 16 shows the welfare comparisons between the “full-information” Boston mechanism used in the compulsory service match and the Random Serial Dictatorship, when the list length under RSD is unrestricted (the dashed lines on top) and when a maximum of k hospitals can be listed, with values of k varying on the x-axis. As we can see in the figure, in all cases, the probability weighting model predicts the welfare gains from a switch to RSD to be smaller than the expected utility model predicts. The reason is similar to that for the decline in gains from information when probability weighting is accounted for: with probability weighting, assignment to less desirable hospitals lowers utility more, and because the unlucky draws of priority under RSD lead to assignment to such hospitals, the overall utility from RSD goes down.

Probability weighting and expected utility models agree on which mechanism outperforms the other when the list length under RSD is at the extremes: When one assumes that the doctors rank all hospitals without any restrictions, RSD is far better than the length-constrained Boston mechanism employed in the compulsory service match. Similarly, when the doctors can only rank a small number of hospitals under RSD (and they are assumed to be behaving naively even under such a restriction), the Boston mechanism clearly outperforms the RSD. More interesting comparisons emerge for a range of list lengths in the middle. For these list lengths, the expected utility model finds RSD to outperform the Boston mechanism, whereas the probability weighting model says otherwise. This result shows that an analyst can reach different conclusions about which mechanism outperforms the other for a range of realistic list lengths depending on whether probability weighting is taken into account or not.

In Appendix Table 3 I show the percentiles of welfare gains from a switch to an unconstrained version of RSD. Panel A shows the results from a model that assumes expected utility, and Panel B shows the results from a model that allows for probability weighting. Under both models, there is a wide dispersion of gains. Although the majority gains from an untruncated RSD, there is a substantial share of doctors who would actually lose from it. Section 8 will explore the role of heterogeneity in probability weighting in explaining the distributional consequences of a mechanism change.

When comparing two random assignment mechanisms, it is also useful to ask how volatile the resulting utilities are. To do that, I compare standard deviations of the utilities using a procedure similar to the one in equation 5. Appendix Table 4 shows that, under both the expected utility and probability weighting models, the full-information Boston mechanism clearly outperforms the Random Serial Dictatorship in terms of offering less volatile placement outcomes. As we have

seen in Section 3.5, information helps doctors to find safety options with high placement chances, which allows them to secure themselves against the worst outcomes when they are unlucky.

8 Extension: Heterogeneity In Probability Weighting and Its Implications

In Section 6 we have seen that there are winners and losers of information provision, and Section 7 showed that the same is true for a switch to alternative assignment mechanisms. This section explores the role of heterogeneity in probability weighting in explaining who these winners and losers are.

It is natural to expect the degree of probability weighting to vary across individuals. Indeed, the previous literature has documented substantial heterogeneity in probability distortions (see, for instance, Barseghyan et al. 2013; Bruhin, Fehr-Duda and Epper 2010). Heterogeneity in probability weighting can be important for evaluating the distributional consequences of information provision and a switch to a strategy-proof mechanism. Agents who distort probabilities more significantly can be expected to have lower expected utility gains when they are provided with information about placement chances, as their objective functions will depart more from expected utility. They might also benefit more from a switch to a strategy-proof mechanism: Since their probability distortions will no longer matter, they will be on an equal footing with other agents.

To learn about heterogeneity in probability weighting, I modify the model in Section 4. I assume that the utility that a user gets from a rank-order list is:

$$U_{it}^R = \sum_{h \in R} w_i(p_{h,t}^R) u_{ih} + \varepsilon_{it}^R$$

where w_i is:

$$w_i(p) = \frac{p^{\gamma_i}}{(p^{\gamma_i} + (1-p)^{\gamma_i})^{1/\gamma_i}} \quad (6)$$

which is the probability weighting function used in Tversky and Kahneman (1992). It captures probability weighting via a single parameter $\gamma_i > 0$, where values of $\gamma < 1$ produce the inverse S-shaped curve found in this paper and in the previous literature. Smaller values of γ indicate less sensitivity to probability increases in the middle of the probability range, and hence, higher degrees of departure from expected utility.

I further assume that:

$$\ln \gamma_i \sim N(\mu_{\gamma_i}, \sigma_{\gamma_i}^2)$$

implying that γ_i is log-normally distributed, ensuring that it is always positive. An attractive

feature of these functional form assumptions is the capability to summarize heterogeneity in probability weighting with a single parameter. I assume that μ_{γ_i} depends on the gender and the school quality of the applicant as follows:

$$\mu_{\gamma_i} = \mu_{\gamma_{g(i)}} + \mu_{\gamma_q} q(i)$$

where $g(i)$ denotes the gender of doctor i and $q(i)$ is the standardized mean residency exam score of her medical school. I assume that the doctors of the same gender share a common variance term, i.e. $\sigma_{\gamma_i}^2 = \sigma_{\gamma_{g(i)}}^2$.

In order to estimate the heterogeneity in the probability weighting function, I use more periods than the one in the baseline estimation in Section 4.2. For each user, I define a set of stable lists as the lists that the user kept for more than two of her active hours, to lower the possibility of using merely exploratory lists, in line with the list durations depicted in Appendix Figure 2. Similar to the baseline model, I use simulated maximum likelihood, but with an additional layer of integration. Appendix Section A.2.4 gives the details.

Panels A and B of Figure 17 show the estimated probability weighting curves at the median values of γ_i for males and females, respectively. Appendix Table 5 shows the parameter estimates and the median values of γ_i for men and women. Men distort probabilities slightly more than women, and graduating from a higher-quality medical school also implies slightly higher levels of distortions. However, observables explain little of the variation in distortion relative to the standard deviation of γ_i .

For a given value of γ_i , define the following measure of departures from expected utility that I call absolute distortion:

$$AD_i(\gamma_i) = \int_0^1 |w_i(p; \gamma_i) - p| dp$$

Following Revelt and Train (2001), I compute each doctor's posterior distribution of γ_i from her observed choices and the estimated population distribution, and then take the posterior mean of AD_i for each doctor i . The details of this process are explained in Appendix Section A.2.5.

Figure 18 shows how the welfare gains from information and from a switch to a strategy-proof alternative relate to the absolute distortion measure. Panel A shows the results for the gains from information provision. To construct this figure, I calculate expected utility differences between the actual case with the final snapshots on the platform and the counterfactual case using initial choices, similar to the exercise in Section 6. The results show that doctors who distort probabilities more heavily benefit less from information provision. This result is intuitive. The winners from information provision in an expected-utility sense are those who resemble expected utility

maximizers more, and those who do not can actually lose from competing with them.

Panel B focuses on the gains from a switch to a strategy-proof alternative. Doctors who distort probabilities more benefit more from such a switch. This pattern means that strategizing under manipulable mechanisms can still be hard for some people, including members of a highly trained profession, and these people would benefit more from mechanisms that do not require them to use probabilistic information.

9 Conclusion

Most existing economic models of applicant behavior in matching markets assume expected utility maximization, and the literature points to a lack of information about admission probabilities as a barrier to making optimal choices in an expected utility sense. In this paper I show that applicants might not maximize their expected utility even when they have access to detailed, real-time information about their admission chances. I find that the applicant behavior is better characterized by an overweighting of nearly certain outcomes relative to merely probable ones. Taking this non-linear weighting of probabilities into account matters for conclusions about the welfare impacts of information about probabilities and alternative assignment mechanisms. An analysis of the heterogeneity in probability weighting reveals that the agents who depart more from expected utility benefit less from information, and more from strategy-proof alternatives. The literature often points to the difficulty of strategizing under manipulable mechanisms. The results in this paper show for whom there is more difficulty.

References


- Acar, S., L. Bilen Kazancık, M. C. Meydan, and M. Işık (2019). *İllerin ve Bölgelerin Sosyo-ekonomik Gelişmişlik Sıralaması Araştırması SEGE - 2017*. Ankara: Kalkınma Ajansları Genel Müdürlüğü, Sanayi ve Teknoloji Bakanlığı.
- Agarwal, N. (2015). An empirical model of the medical match. *American Economic Review* 105(7), 1939–1978.
- Agarwal, N. and P. Somaini (2018). Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica* 86(2), 391–444.
- Andrikogiannopoulou, A. and F. Papakonstantinou (2020). History-dependent risk preferences: Evidence from individual choices and implications for the disposition effect. *The Review of Financial Studies* 33(8), 3674–3718.
- Andrikogiannopoulou, A. and F. Papakonstantinou (2021). Modeling preference heterogeneity within and across behavioral types: Evidence from a real-world betting market. *Swiss Finance Institute Research Paper* (13–53).
- Arteaga, F., A. J. Kapor, C. A. Neilson, and S. D. Zimmerman (2022). Smart matching platforms and heterogeneous beliefs in centralized school choice. *Quarterly Journal of Economics* 137(3), 1791–1848.
- Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives* 27(1), 173–196.
- Barseghyan, L., F. Molinari, T. O’Donoghue, and J. C. Teitelbaum (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review* 103(6), 2499–2529.
- Bernheim, B. D. and C. Sprenger (2020). On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting. *Econometrica* 88(4), 1363–1409.
- Bruhin, A., H. Fehr-Duda, and T. Epper (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica* 78(4), 1375–1412.
- Calsamiglia, C., C. Fu, and M. Güell (2020). Structural estimation of a model of school choices: The Boston mechanism versus its alternatives. *Journal of Political Economy* 128(2), 642–680.
- Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. In J. J. Heckman and E. E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6 of *Handbook of Econometrics*, pp. 5549–5632. Elsevier.

- Chiappori, P.-A., B. Salanié, F. Salanié, and A. Gandhi (2019). From aggregate betting data to individual risk preferences. *Econometrica* 87(1), 1–36.
- Cicchetti, C. J. and J. A. Dubin (1994). A microeconomic analysis of risk aversion and the decision to self-insure. *Journal of Political Economy* 102(1), 169–186.
- Costa, F., L. Nunes, and F. M. Sanches (2024). How to attract physicians to underserved areas? Policy recommendations from a structural model. *Review of Economics and Statistics* 106(1), 36–52.
- DellaVigna, S. (2018). Structural behavioral economics. Working Paper 24797, National Bureau of Economic Research.
- Dreyfuss, B., O. Glicksohn, O. Heffetz, and A. Romm (2025). Deferred acceptance with news utility. *Management Science*.
- Dreyfuss, B., O. Heffetz, and M. Rabin (2022, November). Expectations-based loss aversion may help explain seemingly dominated choices in strategy-proof mechanisms. *American Economic Journal: Microeconomics* 14(4), 515–55.
- Dur, U., R. G. Hammond, and T. Morrill (2018). Identifying the harm of manipulable school-choice mechanisms. *American Economic Journal: Economic Policy* 10(1), 187–213.
- Falcattoni, E. (2018). The determinants of physicians’ location choice: Understanding the rural shortage. Available at SSRN 3493178.
- Freyaldenhoven, S., C. Hansen, J. Perez Perez, and J. M. Shapiro (2024). Visualization, identification, and estimation in the linear panel event-study design. In *Advances in Economics and Econometrics: Twelfth World Congress*. Cambridge University Press.
- Friedrich, B., M. B. Hackmann, A. Kapor, S. J. Moroni, and A. B. Nandrup (2024). Interdependent values in matching markets: Evidence from medical school programs in Denmark. Technical report, National Bureau of Economic Research.
- Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review* 96(4), 1043–1068.
- Gandhi, A. and R. Serrano-Padial (2014). Does belief heterogeneity explain asset prices: The case of the longshot bias. *Review of Economic Studies* 82(1), 156–186.
- He, Y. (2015). Gaming the Boston school choice mechanism in Beijing. Technical report, Toulouse School of Economics.
- Idoux, C. (2023). Integrating New York City schools: The role of admission criteria and family preferences. Working paper, MIT Blueprint Labs.

- İlikan-Rasimoğlu, C. G. (2013). Taşra'yı iyileştirmek: 19. Yüzyıl Osmanlı imparatorluğu'nda Memleket Hekimleri [Treating the Countryside: Provincial Physicians in the 19th Century Ottoman Empire]. *Lokman Hekim Journal* 3(1), 1–6.
- Jullien, B. and B. Salanié (2000). Estimating preferences under risk: The case of racetrack bettors. *Journal of Political Economy* 108(3), 503–530.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Kapor, A. (2025). Transparency and percent plans. *Econometrica* 93(6), 2123–2157.
- Kapor, A. J., C. A. Neilson, and S. D. Zimmerman (2020). Heterogeneous beliefs and school choice mechanisms. *American Economic Review* 110(5), 1274–1315.
- Kliger, D. and O. Levy (2009). Theories of choice under risk: Insights from financial markets. *Journal of Economic Behavior and Organization* 71(2), 330–346.
- Larroucau, T. and I. Rios (2018). Do “short-list” students report truthfully? Strategic behavior in the Chilean college admissions problem.
- Larroucau, T., I. Rios, A. Fabre, C. Neilson, et al. (2025). College application mistakes and the design of information policies at scale. Technical report, National Bureau of Economic Research.
- Meisner, V. and J. von Wangenheim (2023). Loss aversion in strategy-proof school-choice mechanisms. *Journal of Economic Theory* 207, 105588.
- Pathak, P. A. and T. Sönmez (2013). School admissions reform in chicago and england: Comparing mechanisms by their vulnerability to manipulation. *American Economic Review* 103(1), 80–106.
- Rees-Jones, A. and R. Shorrer (2023). Behavioral economics in education market design: A forward-looking review. *Journal of Political Economy Microeconomics* 1(3), 557–613.
- Revelt, D. and K. Train (2001). Customer-specific taste parameters and mixed logit: Households' choice of electricity supplier. Unpublished.
- Roth, A. E. and M. Sotomayor (1992). Two-sided matching. In *Handbook of Game Theory with Economic Applications*, Volume 1, pp. 485–541. Elsevier.
- Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4), 297–323.
- Vrioni, I. (2023). Inequity in centralized college admissions with public and private universities: Evidence from Albania.


Wang, A., S. Wang, and X. Ye (2025). Option value neglect: Evidence from centralized college admission. Technical report, National Bureau of Economic Research.

Figure 1: Platform Predictions and True Placement Chances



Notes: The figure shows the relationship between platform-predicted placement probabilities and actual placement chances. For a given round of the match, I group doctor-listed hospital pairs into equal-sized bins based on platform-predicted placement probabilities. For each bin, the x-axis value shows the average platform-predicted probability, and the y-axis value shows the share of matches that actually took place. The slopes and standard errors reported are from regressions of an indicator for whether the match took place on platform-predicted match probability. The plots exclude doctors whose final lists on the platform and actual lists are determined to be different based on the criteria described in Section 2.2.3, as well as matches that are predicted to have less than one percent probability.


Figure 2: Reliance on the compulsory service program and its scramble round across cities




Notes: The figure shows the reliance of cities in Turkey on the compulsory service program and its scramble round. Panel A shows the share of doctors appointed through the compulsory service program among all doctors serving in a city, averaged over 2018-2023. Panel B shows the development levels of cities in Turkey according to a 2019 report by the Ministry of Industry and Technology. Higher numbers indicate lower levels of development. Panel C shows the share of doctors appointed through the scramble round among all doctors assigned to a city by the compulsory service program over the same years. Panel D repeats the same structure, but considers only the seats associated with the military.

Figure 3: Anticipated Length of Service and Residency Start

Panel A: Expected Length of Service




Panel B: Expected Residency Start



Notes: The figure shows the distribution of responses to questions about anticipated length of service during the compulsory service and the anticipated start date to residency programs under different scenarios. In Panel A, the plot on the left shows the share of survey respondents with a given anticipated length of service if they are assigned to their top choice, and the plot on the right shows the shares when the respondents are assigned through the scramble round instead. In Panel B, the plot on the left shows the share of respondents with a given anticipated start date to residency programs if they are assigned to their top choice, and the plot on the right shows the shares when the respondents are assigned through the scramble round instead.


Figure 4: Distribution of Activity Over Time and Across Users




Notes: Panel A shows, for all hours during the simulation window, the share of users who were active on the platform based on the activity log files. Panel B shows the distribution of the number of hours with at least one activity across all users. Panel C shows, for all hours during the simulation window, the share of users who make at least one change to their lists. Panel D shows the distribution of the number of unique lists across all users.

Figure 5: Survey Evidence on the Reasons for List Changes


Panel A: Reason for Change to First Choice



Panel B: List Finalization Status




Panel C: Reason for Last-Minute Change




Notes: Panel A shows the distribution of reasons for changing the option ranked first initially on the platform. Panel B shows the distribution of responses to a question about whether users considered their lists finalized a day before the deadline. Panel C shows the distribution of reasons for last-minute changes among users who made changes in the final hours.

Figure 6: Responses to External Changes in Scramble Probability




Notes: The figure shows coefficient estimates from an event-study specification. The dependent variable is an indicator for making a change to the rank-order list. The main independent variables of interest are leads and lags of cumulative changes to the probability of entering the scramble round not caused by the user's own changes across snapshots in a given hour. The specifications include doctor and hour fixed effects. Standard errors are clustered by doctor. The numbers inside parentheses on the y-axes refer to the mean of the dependent variable.

Figure 7: Ideal and Final Top Choices



Notes: Panel A shows the share of survey respondents whose top choice on the platform at the deadline matches their stated ideal choice in the survey. Panel B shows the share of users retaining their ideal choice as a function of the bins of placement probabilities to their ideal choice if they had ranked it as their top choice on the platform.


Figure 8: Placement Probability Around First Switch from Ideal Hospital




Notes: The figure shows coefficient estimates from an event-study specification. The dependent variable is the probability of placement to the top choice on the rank-order list. The main independent variables of interest are leads and lags of an indicator for whether a user i who started with her ideal hospital as her top choice made a switch for the first time in time t . The specifications include doctor and hour fixed effects, and endpoint indicators following (Freyaldenhoven et al. 2024). Standard errors are clustered by doctor. The number inside parentheses on the y-axis refers to the mean of the dependent variable at event-time -1 for doctors who made a switch from their top choice.

Figure 9: Placement Chances to Initial, Counterfactual and Final Top Two Choices

Panel A: Top-Ranked Hospitals




Panel B: Second-Ranked Hospitals



Notes: The plot on the left in Panel A shows the placement probability to counterfactual versus final top-ranked hospitals for survey respondents whose counterfactual and final top-ranked hospitals differ. The plot on the right in Panel A shows the placement probability to initial versus final top-ranked hospitals for doctors whose initial and final top-ranked hospitals on the platform differ. The plot on the left in Panel B shows the share of second-ranked choices that are predicted to be filled in the first round for the counterfactual and final second-ranked hospitals for doctors whose counterfactual and final second-ranked hospitals differ. The plot on the right in Panel B shows the corresponding shares for the initial and final second-ranked hospitals for doctors whose initial and final second-ranked hospitals differ.


Figure 10: Share of Doctors without a Choice and Average Scramble Probability over Time




Notes: The blue line shows the average probability of entering the scramble round across users. The orange line shows the share of match participants who are enrolled on the platform but do not have a rank-order list. The green line shows their difference. All measurements are taken at the end of a given hour. The first day is excluded to improve the scaling.

Figure 11: Average Probability-Weighted Characteristics Over Time

Panel A: Distance




Panel B: Development Rank




Notes: Panel A shows the probability-weighted average distance to hospitals on a doctor's list, averaged over all doctors with a non-empty list at the end of a given hour. Panel B shows the probability-weighted average development rank of all cities on a doctor's list, averaged over all doctors similarly. The cities with a higher development rank are less developed. Final per-round placement probabilities of hospitals are used to calculate probability-weighted averages.

Figure 12: Estimated Probability Weighting Curve




Notes: The figure shows the estimated probability weighting curve using the coefficients on the Bernstein basis polynomials presented in Appendix Table 1. The dotted line is the 45-degree line. The shaded area is a pointwise 95 percent confidence band.

Figure 13: List Utility Around First Switch from Ideal Hospital




Notes: The figure shows coefficient estimates from an event-study specification. The dependent variable is the deterministic utility from a user's rank-order list converted to kilometers, using the estimates from the expected utility model (Panel A) or from the probability-weighting model (Panel B). The main independent variables of interest are leads and lags of an indicator for whether a user i who started with her ideal hospital as her top choice made a switch for the first time in time t . The specifications include doctor and hour fixed effects, and endpoint indicators following (Freyaldenhoven et al. 2024). Standard errors are clustered by doctor.

Figure 14: Share of Users Choosing Certainty for Top Choice




Notes: The figure shows the share of users choosing options that offer certain placement chances for their top choice both in the real data and under the choice probabilities predicted by alternative models. To calculate the share of users choosing top choices with certain placement chances under a given model, for each user-time pair, I calculate the probability that the top choice at the chosen list has certain placement chances by using model-predicted list choice probabilities, and then average across user-time pairs.

Figure 15: Average Welfare Gains From Information




Notes: The figure shows the average welfare gains from information provision, by calculating the average of the differences in expected utility between the final predictions on the platform and the counterfactual scenario where all doctors use their initial lists. The welfare is evaluated using both the parameter estimates from a model that assumes expected utility and a model that allows for probability weighting.

Figure 16: Comparing Mechanisms Under Alternative List Lengths



Notes: The figure shows the welfare improvement from a switch to Random Serial Dictatorship as the maximum list length allowed under RSD varies. The dashed line “Full-ranking RSD (EU)” shows the average welfare improvement when doctors are allowed to rank all options and the welfare is evaluated using the expected utility model. “Full-ranking RSD (PW)” shows the average welfare improvement under the probability weighting model. Triangles correspond to average welfare gains under the expected utility model and the maximum list lengths shown on the x-axis. Circles correspond to the welfare gains under the probability weighting model. A third dashed line marks a welfare gain of zero. The shaded region shows the list length range where the expected utility and probability weighting models disagree about which mechanism outperforms the other. Its ends are the list lengths at which the average welfare gain crosses zero under each model, interpolated linearly between the simulated list lengths.

Figure 17: Heterogeneity in Probability Weighting




Notes: Panels A and B show the estimated probability weighting function at the median value of γ_i from the estimated population distribution, for males and females, respectively. The solid red curve is the probability weighting function $w(p) = p^\gamma / (p^\gamma + (1-p)^\gamma)^{1/\gamma}$, and the dotted gray line represents the expected utility case ($w(p) = p$). The shaded area is a pointwise 95 percent confidence band.

Figure 18: Welfare Gains and Heterogeneity in Probability Distortions

Panel A: Gains from Information Provision

Panel B: Gains from Strategy-Proof Mechanism



Notes: Each point represents the mean welfare gain for one of ten equal-width bins of the absolute distortion measure. The absolute distortion measure is the posterior mean of each individual's absolute distortion. The dashed line is the least-squares fit to the individual observations, and the annotation reports its slope and standard error. Panel A shows the relationship between welfare gains from information provision and probability distortions. Panel B shows the relationship between gains from switching to a strategy-proof mechanism (Random Serial Dictatorship) and probability distortions. Welfare gains are expressed in kilometers of distance to the medical school the individual graduated from.

Table 1: Comparison of Actual and Counterfactual Placements

Panel A: Using Alternative Lists Based on Initial Choices								
Average scramble probability								
	Final				Counterfactual			
	0.048				0.080			
Percentiles of reduction in scramble round probability								
10	20	30	40	50	60	70	80	90
-0.036	-0.004	0.000	0.000	0.000	0.000	0.000	0.014	0.192
Average top two placement chance								
	Final				Counterfactual			
	0.906				0.787			
Percentiles of increase in top two placement chance								
10	20	30	40	50	60	70	80	90
-0.088	-0.042	-0.006	0.000	0.002	0.094	0.216	0.338	0.480
Panel B: Using Stated Counterfactual Top Two Choices								
Average top two placement chance								
	Final				Counterfactual			
	0.935				0.776			
Percentiles of increase in top two placement chance								
10	20	30	40	50	60	70	80	90
0.000	0.000	0.049	0.094	0.131	0.173	0.226	0.296	0.378

Notes: The table shows placement chances under the final rank-order lists compared to placement chances under counterfactual scenarios. In Panel A, I rerun simulations using alternative rank-order lists. These alternative lists are constructed by taking the initial choices the doctors made on the platform for each alternative position on the list. Average scramble probability and average top two placement chance using the actual and these counterfactual rank-order lists are displayed, as well as the percentiles of the decrease in the probability of being assigned through the scramble round and the increase in top two placement chance in the actual relative to the counterfactual scenario. In Panel B, I rerun simulations using the stated top two alternatives for doctors who responded to a survey question on what their top two choices would be if they did not use the platform. For each doctor who responded, I fix their top two choices to these counterfactual alternatives, while keeping their remaining choices and those of other doctors same as in the final lists. I do this separately for each survey respondent. For respondents whose stated counterfactual choices were the same as their final choices, their counterfactual placement chances equal their actual placement chances. Average top two option placement probabilities with the actual and counterfactual lists, as well as percentiles of increase in top two placement chance are displayed.

A Appendix

A.1 Tasks and Workplaces during the Compulsory Service

Doctors can choose among several workplace types to complete their compulsory medical service. These workplace types are important for determining the tasks and wages of doctors. Assignments to state hospitals, research hospitals (often affiliated with universities), and city hospitals (hospitals recently built by the private sector and operated by the state) usually imply working in emergency rooms. Emergency-room work has some advantages such as gaining experience with a diverse set of patients and higher wages, but it also brings longer work hours and less time to prepare for the medical specialty exam.

Doctors can also work in emergency stations or coordination centers. The former often refers to working in ambulances alongside other medical care staff, and the latter refers to coordinating emergency and patient transfer cases. Since emergency stations and coordination centers in the same district are usually managed from the same center, transitions between them are possible. Similar to emergency rooms, the types of cases that the doctors can face in emergency stations vary from severe ones such as cardiac arrests to less severe ones.

In addition to patient-facing workplaces, there are also alternatives that offer more administrative tasks, such as public health centers and district health headquarters. Doctors working in these places coordinate healthcare services such as vaccination and routine screening programs. They can still be called to work in emergency rooms when there are staffing needs, but they mostly work within a fixed schedule, and for lower wages. The advantage of these positions is mainly having more time to prepare for the medical residency exam.

Finally, there are positions affiliated with the Ministry of Defense. Doctors appointed to these positions work in military clinics. Similar to other positions, there are both advantages and disadvantages associated with working in military clinics, and these vary greatly based on the location. Doctors appointed to these positions undergo medical and security checks, and usually have to wait around a month before starting to work, which can be an advantage for residency exam preparation. They also have more regular hours of work compared to doctors working in emergency clinics. However, positions affiliated with the military can involve serving outside the country in areas close to conflict zones. Even though such appointments are temporary, they are still undesirable for many doctors.

A.2 Details of Estimation

This section discusses the details of the models in Sections 4 and 8. Throughout, I omit the dependence on doctor observables D_i when writing the likelihood terms. I write the likelihoods for

doctors who responded to all survey questions, as the likelihoods for other users can be constructed similarly.

A.2.1 The Likelihood

In this section I provide further details on the estimation of the model discussed in Section 4.2.

Recall that the utility from a hospital at time t is:

$$u_{iht} = \beta_i^x X_h + \beta_i^d \text{distance}_{ih} + \varepsilon_{iht}$$

with ε_{iht} i.i.d. Type-I extreme value across (i, h, t) . Let $\varepsilon_{i,t} \equiv \{\varepsilon_{iht}\}_{h \in \mathcal{H}}$ denote the vector of hospital-level shocks for doctor i at time t . As in the main text, let α collect the parameters of the model.

Following the distributional assumptions on hospital-level shocks, the likelihood of survey responses to ideal choices takes the following rank-ordered logit form:

$$P(h_{i,1}, h_{i,2} \mid \alpha) = \frac{\exp(\beta X_{i,h_{i,1}})}{\sum_{h \in \mathcal{H}} \exp(\beta X_{ih})} \cdot \frac{\exp(\beta X_{i,h_{i,2}})}{\sum_{h \in \mathcal{H} \setminus \{h_{i,1}\}} \exp(\beta X_{ih})}.$$

List factors. With the distributional assumption on the list-level shocks, the probability that doctor i chooses the rank-order list R^* at $t \in \{T-1, T\}$, conditional on the hospital-level shocks at t , is:

$$\Pr(R^* = R_{i,t} \mid \alpha, \varepsilon_{i,t}) = \frac{\exp\left(\frac{\sum_{h \in R^*} w(p_{h,t}^{R^*}) u_{iht}}{\kappa_{list}}\right)}{\sum_{R \in \mathcal{R}} \exp\left(\frac{\sum_{h \in R} w(p_{h,t}^R) u_{iht}}{\kappa_{list}}\right)}.$$

The hospital-level shocks need to be integrated out to get the marginal probability:

$$P(R_{i,t} \mid \alpha) = \int \Pr(R_{i,t} \mid \alpha, \varepsilon_{i,t}) f(\varepsilon_{i,t}) d\varepsilon_{i,t}, \quad t \in \{T-1, T\}.$$

This integral does not have a closed-form expression and is thus simulated. The dimension of this integral is equal to the number of available hospitals. As I explain in Section A.2.3, I sample from the choice sets, which makes the dimension manageable. Let N_d denote the number of draws from the distribution of $\varepsilon_{i,t}$. I compute:

$$\hat{P}(R_{i,t} \mid \alpha) = \frac{1}{N_d} \sum_{d=1}^{N_d} \Pr(R_{i,t} \mid \alpha, \varepsilon_{i,t}^d), \quad t \in \{T-1, T\}.$$

I set $N_d = 500$. The simulated likelihood for doctor i is then the product:

$$P(h_{i,1}, h_{i,2} \mid \alpha) \cdot \hat{P}(R_{i,T-1} \mid \alpha) \cdot \hat{P}(R_{i,T} \mid \alpha).$$

A.2.2 Probability Weighting Function

In the main analysis, I approximate the probability weighting function $w(\cdot)$ using a sieve-based approach with Bernstein polynomials as the basis. The k th Bernstein basis polynomial of degree n is a function $[0, 1] \rightarrow [0, 1]$:

$$B_{k,n}(p) = \binom{n}{k} p^k (1-p)^{n-k}, \quad k = 0, 1, \dots, n \quad (7)$$

I approximate the probability weighting function using a degree n Bernstein polynomial, implying:

$$w(p) = \sum_{k=0}^n \theta_k B_{k,n}(p)$$

I ensure that $w(p)$ is bounded below by 0 by imposing that $\theta_k \geq 0$ for all k , and that it is bounded above by 1 by imposing that $\theta_k \leq 1$ for all k . I also set $w(1) = 1$ and $w(0) = 0$ by setting $\theta_n = 1$ and $\theta_0 = 0$. In the baseline model, I set $n = 4$.

A.2.3 Choice Sets and the Assignment Probabilities for Alternative Lists

I assume that the choice set of user i at time t consists of the one-shot swaps of her list R at time t :

$$\mathcal{R}_{it} = \mathcal{S}_{R_{it}}$$

where $\mathcal{S}_{R_{it}}$ refers to the set of lists that can be obtained by replacing one of the hospitals in R_{it} by an alternative hospital that is not included in the list. This restriction is motivated by two facts. First, letting the choice set include the set of all potential lists would be computationally infeasible, given that there are 2290 hospitals, and the lists can include up to ten hospitals. In addition to the computational challenge, it would also be unrealistic to assume that the users consider all lists at any time. Many combinations of lists are unlikely to be realistic choices. In fact, Appendix Figure 13 shows that in the majority of cases where the longest-held list in a given hour differs from the one in the preceding hour, this difference is limited to one position on the list. A similar restriction

motivated by both computational feasibility and applicant behavior is found in Idoux (2023).^{19 20}

Even after restricting the choice sets to consist of one-shot swaps, the set of alternatives from which an applicant can choose remains large. In order to overcome the remaining challenge, I randomly select 30 alternative lists in addition to the chosen lists. Randomly sampling from the choice set satisfies the uniform conditioning property in McFadden (1978). The consistency of this approach for mixed logit models has been shown in Guevara and Ben-Akiva (2013).

For calculating the placement probabilities associated with an alternative list R' at time t , I replace user i 's chosen list at time t , R_{it} , with the alternative R' , and keep all other users' lists at time t the same. Then, I run the assignment mechanism 100 times, and calculate the share of runs in which the user was assigned to a given hospital on the list.

A.2.4 Estimation with Heterogeneity in Probability Weighting

The survey factor stays as described in Section A.2.1 since it does not depend on γ_i . I rewrite the conditional list-level choice probability to emphasize its dependence on γ_i as follows:

$$\Pr(R^* = R_{i,t} \mid \tilde{\alpha}, \gamma_i, \varepsilon_{i,t}) = \frac{\exp\left(\frac{\sum_{h \in R^*} w(p_{h,t}^{R^*}; \gamma_i) u_{iht}}{\kappa_{list}}\right)}{\sum_{R \in \mathcal{R}} \exp\left(\frac{\sum_{h \in R} w(p_{h,t}^R; \gamma_i) u_{iht}}{\kappa_{list}}\right)}$$

Let $\tilde{\alpha} = (\beta, \kappa_{list}, \mu_\gamma, \sigma_\gamma)$ denote the set of parameters to be estimated in this model with heterogeneity. Because the hospital-level shocks ε_{iht} are independent over time, conditional on γ_i , the list-level choices are independent. The likelihood of the observed sequence of choices can be written as:

$$P(h_{i,1}, h_{i,2} \mid \tilde{\alpha}) \cdot \int \prod_t P(R_{i,t} \mid \tilde{\alpha}, \gamma_i) f_\gamma(\gamma_i; \mu_\gamma, \sigma_\gamma) d\gamma_i$$

where the list-level choice probabilities, conditional on γ_i , are:

$$P(R_{i,t} \mid \tilde{\alpha}, \gamma_i) = \int \Pr(R_{i,t} \mid \tilde{\alpha}, \gamma_i, \varepsilon_{i,t}) f(\varepsilon_{i,t}) d\varepsilon_{i,t}.$$

¹⁹The sufficiency of choice sets constructed similarly by swapping alternatives instead of the chosen options has been shown in other contexts under models that assume expected utility, see Vrioni 2023 and Larroucau and Rios 2018 for examples. Although there is no guarantee of sufficiency under probability weighting, one-shot swaps give informative comparisons that are close to the actual choice process by applicants as discussed above.

²⁰Another reasonable restriction could be to assume that the choice set consists of lists chosen by at least one doctor, as in Kapor (2025). The downside of this approach in this setting is that the median number of seats per hospital is three. As a result, if a doctor submits exactly the same list as another doctor, the placement chances under that list would be greatly diminished.

Similar to the baseline model, I use maximum simulated likelihood to estimate this model. For simulating the outer integral over γ_i , I use 50 Halton draws.

A.2.5 Posterior Mean of the Absolute Distortion For Individual Doctors

Denote by $\Pr(y_i | \gamma_i)$ the likelihood of the observed sequence of choices conditional on γ_i :

$$\Pr(y_i | \gamma_i) = P(h_{i,1}, h_{i,2} | \tilde{\alpha}) \cdot \prod_t \int \Pr(R_{i,t} | \tilde{\alpha}, \gamma_i, \varepsilon_{i,t}) f(\varepsilon_{i,t}) d\varepsilon_{i,t} \quad (8)$$

Given the parameter estimates (which include those governing the population distribution of γ_i), and the observed choices, the posterior distribution of γ_i for each individual can be calculated using Bayes' rule:

$$h(\gamma_i | y_i) = \frac{\Pr(y_i | \gamma_i) \cdot f(\gamma_i | \hat{\mu}_\gamma, \hat{\sigma}_\gamma)}{\Pr(y_i)}$$

where $\hat{\mu}_\gamma$ and $\hat{\sigma}_\gamma$ are the estimated parameters. The population distribution $f(\gamma_i | \hat{\mu}_\gamma, \hat{\sigma}_\gamma)$ serves as the prior, and the individual likelihood from observed choices updates this prior.


Let AD_i denote the absolute distortion for doctor i at a particular γ_i value:

$$AD_i(\gamma_i) = \int_0^1 |w_i(p; \gamma_i) - p| dp$$

The numerical approximation to the posterior mean of AD_i can then be calculated as:


$$\overline{AD}_i = \frac{\sum_{s=1}^S AD_i(\gamma_i^{(s)}) \Pr(y_i | \gamma_i^{(s)})}{\sum_{s=1}^S \Pr(y_i | \gamma_i^{(s)})}$$

Appendix Figure 1: Screenshots from the Platform




Notes: The figure shows the simulation platform's screens with the text translated to English by the author. Panels A and B reproduce the platform's display for the final list of one participant in our data: the hospital names, quotas, choosing counts, and placement probabilities are those of the participant's list, drawn in the platform's layout. Panels C and D are screenshots from the platform. Panel A shows the unconditional placement probabilities to all hospitals on the participant's list as well as the probability of being assigned through the scramble round. Panel B shows the first three hospitals on the list, with conditional probabilities of being assigned to a hospital at rank k in round k if the doctor is not assigned in the preceding $k - 1$ rounds. Panel C shows, for a selected hospital, per-round conditional placement probabilities if the doctor is not assigned to any hospital in the preceding rounds. Panel D shows the filtering screen, where the doctor is searching for hospitals with at least 20% placement chances in round 4 conditional on not being assigned in the first 3 rounds.

Appendix Figure 2: Duration of Lists by Duration Rank




Notes: For each user, I sort the lists by the number of hours held during active hours, which I call the “Duration Rank”, and plot the distribution of hours held for a given duration rank, with box plots showing the median and interquartile range.

Appendix Figure 3: Maximum Placement Probability of Changed Positions over Time




Notes: The figure shows the median of the maximum placement probability held in positions that were affected by a list change. To construct this figure, for each user, I compare her list at the beginning and at the end of an hour, and find the positions that are different. For each position that is different, I calculate the maximum placement probability to that position across both lists, then take the maximum of these maxima, which I call the maximum placement probability affected by the list change. Then, I calculate the median of the maximum placement probability affected by a list change across all users, restricting attention to users who made a change to a position with non-zero placement probability.

Appendix Figure 4: Most Important Factor for the First Choice




Notes: The figure shows the distribution of responses to a question about the most important factor when choosing the first option on the rank-order list.

Appendix Figure 5: Changes in List Length in Response to External Changes in Scramble Round Probability



Notes: The figure shows coefficient estimates from an event-study specification where the dependent variable is the change in the list length between the end of an hour and the end of the preceding hour, and the main independent variables of interest are leads and lags of cumulative changes to the probability of entering the scramble round not caused by the user's own changes across snapshots in a given hour. The specifications include doctor and hour fixed effects. Standard errors are clustered by doctor. The numbers inside parentheses on the y-axes refer to the mean of the dependent variable.


Appendix Figure 6: Search in Response to External Changes in Scramble Probability




Notes: The figure shows coefficient estimates from an event-study specification. The dependent variable is an indicator for searching for alternatives in a given hour, as captured by whether the user loaded the page of all available hospitals at least once. The main independent variables of interest are leads and lags of cumulative changes to the probability of entering the scramble round not caused by the user's own changes across snapshots in a given hour. The specifications include doctor and hour fixed effects. Standard errors are clustered by doctor. The numbers inside parentheses on the y-axes refer to the mean of the dependent variable.

Appendix Figure 7: Cumulative Placement Probabilities at Different Snapshots


Panel A: Cumulative Probabilities, 24th Hour




Panel B: Cumulative Probabilities, 48th Hour



Panel C: Cumulative Probabilities, 72nd Hour




Panel D: Cumulative Probabilities, Final Hour




Notes: Panels A to D show boxplots of the cumulative placement probabilities to the options on rank-order lists, at the 24th, 48th, and 72nd hours of the simulation window and at the final snapshot. For each list position, the cumulative probability is the sum of the placement probabilities of that position and all higher-ranked positions.

Appendix Figure 8: Scramble Round Share and the Share of Military Seats Across Rounds




Notes: Panel A shows the share of users assigned through the scramble round in different rounds of the match. Panel B shows the share of seats affiliated with the military in each of these rounds of the match. Panel C shows the correlation between the share of military seats and the share of participants who are assigned through the scramble round.

Appendix Figure 9: Medical School Quality Across Turkey




Notes: The figure shows the average quality of medical schools in cities across Turkey based on the average residency exam scores. For cities with more than one medical school, the plot shows the average of those schools' scores, weighted by the number of doctors each sends into the compulsory service match.

Appendix Figure 10: The Probability Weighting Function




Notes: The figure reproduces the probability weighting function from (Tversky and Kahneman 1992), using their parameter estimates from the experimental data. p corresponds to objective probabilities, and $w(p)$ corresponds to weighted probabilities.

Appendix Figure 11: Empirical CDFs of Placement Probabilities at Each Rank Using Real and Simulated Data




Notes: The figure shows, for a given rank, the empirical CDF of the placement probabilities. “Actual” refers to the empirical CDF of the placement probabilities obtained from the final snapshot on the platform. I calculate the model-implied CDFs using equation 4, which weights, for each (user, time) choice, the indicator that the placement probability at the rank is at or below the threshold by the model-implied list choice probability, sums across lists in the user’s choice set, and averages across (user, time) choices. Panel A shows the results for options ranked as the first choice. Panel B and C repeat the same structure for the second and third choices, respectively.

Appendix Figure 12: Differential Responses to External Changes in Scramble Probability By The Slope At the Average Scramble Probability Before the Change



Notes: The figure shows coefficient estimates from an event-study specification. The dependent variable is an indicator for making a change to the rank-order list. The main independent variables of interest are leads and lags of cumulative changes to the probability of entering the scramble round not caused by the user's own changes across snapshots in a given hour, and the same leads and lags interacted with the slope of the probability weighting function evaluated at the average scramble round probability in the preceding hour. Plotted are the coefficients on the interaction terms. The specifications include doctor and hour fixed effects. Standard errors are clustered by doctor. The numbers inside parentheses on the y-axis refer to the mean of the dependent variable.

Appendix Figure 13: Differences Between Longest-Held Lists in Consecutive Hours



Notes: The figure shows the distribution of the number of positions that differ between the longest-held lists during consecutive hours. To construct this figure, for each user and hour, I find the list that the user held for the longest duration. Then, I calculate the number of positions in which these lists in consecutive hours differ. Changes to the positions where the assignment probability is less than one percent in both hours are ignored.

Appendix Table 1: Estimation Results

	Probability Weighting			Expected Utility		
	Female	Male	Res. Score	Female	Male	Res. Score
Military clinic	9.01 (0.45)	6.97 (0.44)	1.54 (0.24)	12.56 (0.55)	11.15 (0.57)	1.94 (0.35)
Emergency station	9.79 (0.38)	10.16 (0.39)	0.95 (0.19)	13.45 (0.49)	14.42 (0.53)	1.33 (0.32)
State hosp.	11.78 (0.38)	11.94 (0.39)	0.91 (0.19)	15.31 (0.50)	16.09 (0.53)	1.25 (0.32)
Research hosp.	12.52 (0.38)	12.71 (0.39)	0.84 (0.20)	15.99 (0.50)	16.77 (0.54)	1.20 (0.32)
Public health cent.	9.77 (0.38)	9.36 (0.39)	0.66 (0.20)	13.44 (0.49)	13.67 (0.53)	1.02 (0.32)
District health cent.	11.20 (0.38)	10.62 (0.39)	0.90 (0.19)	14.76 (0.50)	14.77 (0.53)	1.25 (0.32)
City hosp.	13.37 (0.39)	13.52 (0.39)	0.73 (0.20)	16.86 (0.50)	17.62 (0.54)	1.08 (0.32)
Development pct.	-1.86 (0.14)	-1.13 (0.15)	0.23 (0.11)	-1.49 (0.14)	-0.75 (0.15)	0.29 (0.11)
Distance (100km)	-0.62 (0.01)	-0.60 (0.01)	-0.03 (0.01)	-0.63 (0.01)	-0.61 (0.01)	-0.03 (0.01)
κ_{list}		1.54 (0.03)			1.61 (0.03)	
θ_1		0.38 (0.02)				
θ_2		0.29 (0.02)				
θ_3		0.33 (0.02)				

Notes: The table shows the parameter estimates from estimating the model in Section 4 with simulated maximum likelihood. Development pct. is the percentile rank of the city's development scores, which takes values between 0 and 1 and equals 1 for the least developed city. The "Res. Score" columns show the effect of a one standard deviation increase in the average school residency exam score on the estimated coefficient. θ parameters show the coefficients on the Bernstein basis polynomials of degree one to three. The expected utility columns show the estimates from an alternative model, where $w(p) = p$. Both models include controls for NUTS-1 regions of Turkey. The indicator for NUTS-1 region TRB is omitted.

Appendix Table 2: Increases in Utility from Information Provision

<i>Panel A: Increases in Expected Utility (Without Probability Weighting)</i>								
<i>Mean increase</i>								
67.58								
<i>Percentiles of increase</i>								
10	20	30	40	50	60	70	80	90
-205.38	-110.37	-62.10	-24.90	0.00	4.72	70.30	199.29	483.95
<i>Panel B: Increases in Expected Utility (With Probability Weighting)</i>								
<i>Mean increase</i>								
44.39								
<i>Percentiles of increase</i>								
10	20	30	40	50	60	70	80	90
-191.55	-102.77	-58.97	-24.57	0.00	3.78	65.51	178.45	388.52
<i>Panel C: Increases in $w()$-Weighted Utility (With Probability Weighting)</i>								
<i>Mean increase</i>								
58.20								
<i>Percentiles of increase</i>								
10	20	30	40	50	60	70	80	90
-270.64	-148.78	-64.16	0.00	13.86	81.34	155.42	265.03	430.65

Notes: The table shows changes in utility when the final rank-order lists are used instead of counterfactual lists constructed from initial choices, expressed in kilometer equivalents. Panel A uses parameter estimates from a model without probability weighting. Panel B uses parameter estimates from the model with probability weighting. Panel C shows changes in the objective function in Section 4, using estimated $w(p)$ to weight utilities.

Appendix Table 3: Differences in Expected Utility between Random Serial Dictatorship and Full-Information Boston Mechanism

<i>Panel A: Differences in Expected Utility, RSD - Boston (Without Probability Weighting)</i>								
<i>Mean difference</i>								
143.70								
<i>Percentiles of difference</i>								
10	20	30	40	50	60	70	80	90
-171.41	-93.82	-34.90	18.37	75.47	144.84	239.64	363.22	560.13

<i>Panel B: Differences in Expected Utility, RSD - Boston (With Probability Weighting)</i>								
<i>Mean difference</i>								
125.93								
<i>Percentiles of difference</i>								
10	20	30	40	50	60	70	80	90
-187.00	-111.16	-47.74	7.35	61.37	128.56	226.04	345.71	540.08

Notes: The table shows the distribution of differences in expected utility between Random Serial Dictatorship and the full-information Boston mechanism, expressed in kilometer equivalents. For the full-information Boston mechanism, the final rank-order lists at the end of the simulation process are used. For Random Serial Dictatorship, participants are randomly ordered 250 times, the mechanism is run with each ordering, and the mean utility is calculated. Panel A uses parameter estimates from a model without probability weighting. Panel B uses parameter estimates from the model with probability weighting.

Appendix Table 4: Differences in Standard Deviation of Utility between Random Serial Dictatorship and Full-Information Boston Mechanism

<i>Panel A: Differences in SD, RSD - Boston (Without Probability Weighting)</i>								
<i>Mean difference</i>								
168.48								
<i>Percentiles of difference</i>								
10	20	30	40	50	60	70	80	90
-44.96	30.44	82.69	126.45	169.58	215.58	264.30	321.65	400.66
<i>Panel B: Differences in SD, RSD - Boston (With Probability Weighting)</i>								
<i>Mean difference</i>								
186.12								
<i>Percentiles of difference</i>								
10	20	30	40	50	60	70	80	90
-24.78	46.72	97.65	139.34	181.83	228.51	277.92	335.43	415.54

Notes: The table shows the distribution of differences in the standard deviation of utility between Random Serial Dictatorship and the full-information Boston mechanism, expressed in kilometer equivalents. For the full-information Boston mechanism, the final rank-order lists at the end of the simulation process are used. For Random Serial Dictatorship, participants are randomly ordered 250 times, the mechanism is run with each ordering, and the standard deviation of utilities is calculated. Panel A uses parameter estimates from a model without probability weighting. Panel B uses parameter estimates from the model with probability weighting.

Appendix Table 5: Estimated Distribution of the Probability Weighting Parameter

	Estimate
μ_{γ_M}	-0.531 (0.015)
μ_{γ_F}	-0.503 (0.016)
μ_{γ_q}	-0.041 (0.010)
$\sigma_{\gamma_M}^2$	0.276 (0.014)
$\sigma_{\gamma_F}^2$	0.336 (0.015)
Median γ_i (Men)	0.588 (0.009)
Median γ_i (Women)	0.605 (0.010)
SD γ_i (Men)	0.380 (0.017)
SD γ_i (Women)	0.452 (0.021)

Notes: The table reports the estimated parameters of the distribution of the probability weighting parameter, where $\ln \gamma_i \sim N(\mu_{\gamma_i}, \sigma_{\gamma_i}^2)$, $\mu_{\gamma_i} = \mu_{\gamma_{g(i)}} + \mu_{\gamma_q} q(i)$ and $\sigma_{\gamma_i}^2 = \sigma_{\gamma_{g(i)}}^2$.

Additional References

- Guevara, C. A. and M. E. Ben-Akiva (2013). Sampling of alternatives in logit mixture models. *Transportation Research Part B: Methodological* 58, 185–198.
- Idoux, C. (2023). Integrating New York City schools: The role of admission criteria and family preferences. Working paper, MIT Blueprint Labs.
- McFadden, D. (1978). Modelling the choice of residential location. In A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull (Eds.), *Spatial Interaction Theory and Planning Models*, pp. 75–96. Amsterdam: North-Holland.