

PROBABILISTIC PREDICTIONS FOR FINITE RANDOM MATCHING MARKETS

NICK ARNOSTI, UNIVERSITY OF MINNESOTA

ABSTRACT. This paper introduces a class of stable matching models, which includes the traditional finite market model and the continuum model of Azevedo and Leshno (2016) as special cases. For any model in this class, stable matchings exist and form a lattice. The key innovation is the introduction of a “vacancy function” which translates expected interest at each school into the probability that the school fails to fill its seats. This makes it possible to generate agent-level probabilistic predictions for finite random matching markets.

I generate predictions using a vacancy function motivated by the assumption that interest in each school should follow a Poisson distribution. I show that this model always has a unique stable outcome, and that it accurately predicts the simulated distribution of cutoffs in markets with as few as ten schools and twenty students. This model also generates new insights about the number and quality of matches. When schools are homogeneous, it provides upper and lower bounds on students’ average rank, which match results from Ashlagi et al. (2017) but apply to more general settings. I also use this model to provide analytical expressions for the number of matches formed in a platform pricing setting considered by Marx and Schummer (2021).

1. INTRODUCTION

Ever since Gale and Shapley (1962) defined stability in two-sided matching markets, the topic has generated a great deal of interest from academics and practitioners alike: their paper has over 8,000 citations, and variants of their deferred acceptance algorithm are used to assign medical residencies in the United States and public school seats in cities across the globe. These developments prompted the award of the 2012 Nobel Prize to Alvin Roth and Lloyd Shapley “for the theory of stable allocations and the practice of market design.”

Despite this attention, the relationship between market primitives and the set of stable outcomes remains poorly understood. For example, Dur et al. (2018) demonstrated that a walk zone compromise in Boston which had been in place for a decade had almost no effect! In New York, policymakers collected preferences before finalizing tiebreaking rules, in order to simulate and compare several alternatives (Abdulkadiroglu et al., 2009). The Brookings institution highlights the difficulty of anticipating final outcomes as a key challenge facing school choice initiatives:

Even if DA [Deferred Acceptance] algorithms are relatively simple, predicting how student assignment policies will affect enrollment and outcomes is difficult... This creates challenges for policymakers to assess a priori how policy decisions will affect students and schools – and creates potential for unintended negative consequences.

(Kasman and Valant, 2019)

These challenges are not limited to policymakers in complex environments. They arise even in highly stylized matching environments considered by academic researchers, as the following example shows.

Example 1. *There are n students and n schools, each with a single seat. Each student lists $\ell < n$ uniformly random schools (in a uniformly random order). Each school assigns each student independent $U[0, 1]$ lottery numbers (higher is better).*

Related examples were first proposed by Wilson (1972); Knuth (1976); Pittel (1989), and have been studied by several subsequent papers. I pose the following questions.

- I What fraction of students will match to some school on their list?
- II Given a student’s priority at a particular school, what is her chance of being admitted?
- III How do answers change if schools use a different method for determining priority?

Each of these questions is of interest to some stakeholder. The first addresses an aggregate statistic which might be reported by the media or elected officials. The second is salient to students and their parents. The third is relevant to administrators who set prioritization policies.

Section 2 explains in detail why previous models of stable matching fail to address these questions. In summary, past work can be grouped into three categories: models of finite random markets, models with a continuum of participants on both sides, and models with a continuum of students matching to a finite number of schools.

- **Finite Number of Students and Schools.** Complex interactions between individual participants make it difficult to precisely analyze outcomes. As a result, past work imposes strong assumptions on preferences and priorities (like those in Example 1). Even so, it only provides loose bounds on the number of unassigned students.
- **Continuum of Students and Schools.** These models either assume a finite number of agent types (with a continuum of identical agents of each type) or embed types in a topological space, with agents who are “close” in this space having similar preferences and desirability to the other side. Both cases rule out idiosyncratic preference shocks, which are present in Example 1 and many econometric matching models.
- **Continuum of Students, Finitely Many Schools.** These models can accommodate very general preferences and priorities. However, their predictions are only accurate when schools have large capacities, so that demand for each school is predictable. By predicting a deterministic outcome for each student, this approach fails to capture the inherent uncertainty in markets with small capacities. For example, when the model of Azevedo and Leshno (2016) is applied to Example 1, it predicts that every student matches to her first choice, regardless of how priorities are determined. This is clearly incorrect, as a simple balls-in-bins argument establishes that with high probability, only approximately $1 - 1/e \approx 63\%$ of schools are listed as the first choice of some student.

Recent work by Arnosti (2022) answers the first two questions above by establishing limiting equations that describe outcomes in finite markets with many students and schools. That paper provides sharper results than prior work (for example, it calculates the exact fraction of students who remain unassigned in the limit), in a more general model (for example, the schools on students' lists need not be sampled uniformly). Fundamentally, however, the challenges of analyzing outcomes in finite markets remain, and force Arnosti (2022) to maintain two important assumptions from prior work. First, the schools on each student's list are sampled without replacement from a fixed distribution over schools, implying that the top schools on a student's list provide almost no information about the identity of the remaining schools. Second, school priorities come either from independent uniform lotteries or a single lottery that applies across all schools.

1.1. Summary of Contributions.

The preceding discussion illustrates the difficulty of addressing fundamental questions about even relatively simple matching environments. This paper presents a new approach which can be used to generate individual-level *probabilistic* predictions in random matching markets with arbitrary preferences, priorities, and capacities. The resulting expressions are analytically tractable and offer answers to the three questions above, as well as other insights about aggregate match outcomes. I now elaborate on these contributions.

Unifying Framework. I provide a definition of stability that depends (naturally) on school capacities and the joint distribution of preferences and priorities, as well as (much more novelly) on a “vacancy function.” The vacancy function takes as arguments the capacity of a school and the expected demand at that school, and generates a value in $[0, 1]$ which can be interpreted as the probability that the school will have a vacancy.

I show that using a deterministic vacancy function (whose output is always 0 or 1) recovers the traditional definition of stability from Gale and Shapley (1962) (Proposition 1) as well as the definition from the continuum model of Azevedo and Leshno (2016) (Proposition 2). I also show that well-known properties of stable matchings continue to hold for general vacancy functions. In particular, the set of stable matchings forms a non-empty lattice (Theorem 1), and the “Rural Hospital Theorem” holds so long as there are no ties in priority (Theorem 2). Finally, I establish that there is a unique stable matching so long as the probability of a vacancy decreases strictly as expected demand increases (Theorem 3). This uniqueness result is complementary to that of Azevedo and Leshno (2016), and helps to explain the small core observed empirically by Roth and Peranson (1999), and theoretically by Immorlica and Mahdian (2005), Kojima and Pathak (2009), and Ashlagi et al. (2017).

Probabilistic Predictions. One of the most novel and intriguing aspects of this framework is its ability to generate individual-level probabilistic predictions through the use of non-deterministic vacancy functions. By contrast, existing continuum models do not incorporate individual uncertainty: the model of Azevedo and Leshno (2016) predicts a deterministic assignment for each student, and Greinecker and Kah (2021) note that “*there is nothing*

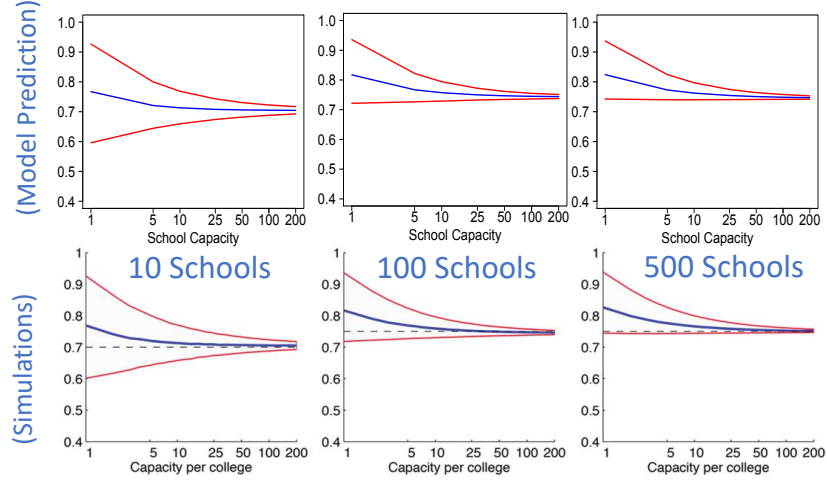


FIGURE 1. The distribution of “cutoff scores” required for admission in a class of examples considered by Azevedo and Leshno (2016). In all scenarios, there are twice as many students as seats, students submit complete lists drawn uniformly at random, and schools’ priority scores are imperfectly correlated: they consist of the average of the student’s quality (drawn uniformly on $[0, 1]$) and iid student-school terms (also drawn uniformly on $[0, 1]$). This correlation renders direct analysis of the finite random market intractable. In the first column there are ten schools, in the second there are 100, and in the third there are 500. The x -axis denotes the number of seats per school, which ranges from 1 to 200.

The bottom row displays simulation results reported by Azevedo and Leshno (2016): the blue line shows the average cutoff score, with the red lines representing the 5th and 95th percentile of the empirical distribution. Their model predicts a deterministic cutoff score, shown by the black dotted line. This prediction does not depend on the number of seats at each school, and does not capture the uncertainty in cutoffs, which is significant unless capacities are large. By contrast, my model can be used to predict the *distribution* of cutoff scores, and captures the fact that there is greater uncertainty in markets with smaller capacities. The top column shows the average, 5th percentile and 95th percentile of the predicted distribution.

random about a matching in our distributional model: the underlying matching of agents is deterministic.”

Using a vacancy function motivated by the assumption that demand at each school follows a Poisson distribution, I generate numerical predictions that closely match simulation results. Figure 1 shows that this approach accurately predicts the *distribution* of school cutoff scores in markets with as few as 20 students and 10 schools, where school priorities are positively (but imperfectly) correlated. As capacities grow, the predicted distributions concentrate around the cutoff predicted by Azevedo and Leshno (2016).

This model also generates predictions about aggregate outcomes. Figure 2 demonstrates that the model predicts the striking finding of Ashlagi et al. (2017) that students’ average rank sharply increases as the market transitions from having a slight surplus of seats to a slight shortage. Figure 6 shows that the model accurately predicts the difference between the number of matches when using two different priority rules, by comparing to simulation




FIGURE 2. Students’ average rank for their assigned school. In this market, there are 40 schools, each with a single seat. The number of students is given along the x axis, and both student preferences and school priorities are drawn iid and uniformly at random. The left panel shows simulation results from Ashlagi et al. (2017) demonstrating that (i) the difference between the school-optimal and student-optimal stable match is typically small, and (ii) in a balanced market (highlighted in gray), adding or removing one student has a dramatic effect. Our proposed model of stable matching generates a unique prediction (right), which closely matches the simulations and captures the dramatic effect of additional students in nearly-balanced markets.

results from Marx and Schummer (2021). Meanwhile, the model of Azevedo and Leshno (2016) predicts that these priority rules should result in the same number of matches.

Analytical Insights. Finally, the Poisson vacancy function can be used to derive analytical insights. Under the assumptions imposed by Arnosti (2022), the Poisson vacancy function results in equations that coincide with the limiting expressions established by his Theorem 1. Arnosti (2022) uses these expressions to derive insights about the number of unmatched students. In this paper, I show that these expressions can also be used to reproduce and generalize insights about students’ average rank from Ashlagi et al. (2017). Specifically, when there is an excess of school seats I provide an upper bound which matches that of Ashlagi et al. (2017) when each school has a single seat (Corollary 1), and generalizes this bound to cases where schools have multiple seats (Proposition 3). Proposition 4 and Corollary 2 provide analogous bounds for the case when there are more students than seats.

Section 2 provides an overview of existing papers on stable matching in large and/or continuum markets, and explains why they cannot answer the questions posed above for Example 1. Section 3 provides a heuristic analysis of Example 1, which motivates my definition of stability, as well as the Poisson vacancy function. Section 4 introduces the model. Section 5 presents results for general vacancy functions, and shows that with a deterministic vacancy function, my definition of stability coincides with those of prior work. Section 6 presents results when using the Poisson vacancy function.

2. RELATED WORK

There is a vast literature on stable matching. This section focuses exclusively on papers that use “large market” or “continuum” models to describe or approximate stable matching outcomes.

2.1. Finite Random Markets. One line of work studies outcomes in large finite markets. To maintain tractability, these papers typically impose strong assumptions. For example, Pittel (1989), Knuth (1996), Ashlagi et al. (2017), Ashlagi et al. (2019), Ashlagi and Nikzad (2019) and Kanoria et al. (2021) all assume that schools are symmetric (student preferences are iid and uniformly distributed), and all but Knuth (1996) assume that school priorities are either identical or drawn independently and uniformly at random.

Immorlica and Mahdian (2005) and Kojima and Pathak (2009) allow for slightly more general preferences, in which the schools on a student’s list are generated by repeatedly sampling from a *non*-uniform distribution over schools. However, these papers focus primarily on establishing the existence of a “nearly unique” stable outcome, and that participants have little to gain from misreporting their preferences. They do not attempt to get precise answers to questions about, for example, the number of students who match. Recent work by Arnosti (2022) adopts the preference formation model of these papers, and does manage to provide a fairly precise description of match outcomes. However, that work continues to assume that school priorities are either identical or drawn independently and uniformly at random.

To be sure, these papers provide many interesting insights. Pittel (1989) shows that there can be a significant gap between the student-proposing and school proposing outcomes. Ashlagi et al. (2017) show that students’ average rank rises dramatically once the number of students exceeds the number of school seats. Ashlagi and Nikzad (2019) demonstrate that in such cases, using a single lottery “almost stochastically dominates” using independent lotteries at each school. However, the analysis underlying these insights is not flexible enough to accommodate more realistic assumptions on preferences and priorities, making it unclear whether the insights derived using these models should be expected to hold in practice. Indeed, Ashlagi et al. (2017) demonstrate that the gap identified by Pittel (1989) vanishes if the market is not perfectly balanced, and Kanoria et al. (2021) argue that the results of Ashlagi et al. (2017) rely on sufficiently long preference lists, and may not be relevant for the assignment of public schools in New York City and elsewhere.

2.2. Markets with a Continuum of Students and Schools. A number of papers consider models with a continuum of students and schools. Azevedo and Hatfield (2018) assume a finite number of agent types on both sides, and a continuum of agents of each type. A similar model is considered by Jagadeesan and Vocke (2021). This approach is fairly tractable, and makes it possible to study many-to-many matching with contracts. However, because agents of the same type are not distinguishable, it is impossible for this type of model to capture the idiosyncratic preferences present in Example 1.

Another approach embeds agents into a topological space which incorporates features that determine each agent’s preferences and desirability. The simplest approach assumes a single “quality” dimension, with all agents on the opposite side preferring a higher-quality partner. For example, Bodoh-Creed and Hickman (2018) assume that the only feature of college which matters to students is its “quality”, while Agarwal (2015) assumes that residency programs have a common ranking of residents. Peski (2011) considers a more general setting in which men and women are embedded into \mathbb{R}^d , and all agents prefer higher values in each coordinate. He is unable to establish existence of a stable match, although he shows that if one exists, it is unique.

Greinecker and Kah (2021) consider a yet-more-general one-to-one matching model, which can accommodate matching with contracts, externalities, and transferable utility. However, one important assumption is that preferences and desirability are continuous in the underlying type space. They write that “If a woman of type w prefers a man of type m to a man of type m' then a woman with a type sufficiently similar to w prefers a man of type sufficiently similar to m to a man of type sufficiently similar to m' .” While seemingly natural, this assumption rules out idiosyncratic utility shocks which are present in Example 1 and commonly assumed by econometricians and theorists. Greinecker and Kah (2021) acknowledge that “our model cannot be used to study the asymptotic stochastic behavior of large finite marriage models in which every agent’s preferences are independently and uniformly chosen from the set of strict rankings of agents on the other side, the approach of Pittel (1989) and many subsequent papers.”

2.3. Markets with Large Schools. An alternative approach assumes a finite number of schools, each of which matches to a continuum of students. This allows very general priorities and preferences: the model of Azevedo and Leshno (2016) accommodates an arbitrary measure over student “types,” which encode both a ranking of schools and a priority score at each school. Given this type measure and the capacity of each school, they find a vector of “cutoff scores” that “clear the market”: if each school admits students whose priority is above its cutoff, and students attend their favorite school where they are admitted, expected “demand” (enrollment) is equal to “supply” (capacity) at each school with a positive cutoff. This approach turns out to be quite tractable, and Che et al. (2019) generalize it by allowing schools to have non-responsive (i.e. substitutable or complementary) choice functions. Many papers have used related models to derive various insights (Abdulkadiroglu et al., 2015; Shi, 2015; Ashlagi and Shi, 2016; Abdulkadiroglu et al., 2017; Bodoh-Creed, 2020; Shi, 2022; Allman et al., 2022).

Azevedo and Leshno (2016) establish that their model describes the limiting behavior of finite markets as the number of seats at each school increases. Intuitively, if each school can match to many students, then the law of large numbers implies that demand for each school is highly predictable. However, their model fails to capture the variability that is inherent in markets with small or modest capacities. Although their model predicts a deterministic

cutoff score for each school, Figure 1 shows simulations from their paper demonstrating that cutoff scores vary significantly when school capacities are not large.

The inability to generate stochastic predictions causes their model to make inaccurate predictions when applied to Example 1. It predicts a cutoff score of 0 at each school, implying that every student matches to his or her first choice. This prediction holds for any length of student lists, and any method for determining student priorities. This is clearly incorrect: a simple balls-in-bins analysis demonstrates that only approximately $1 - (1 - 1/n)^n \approx 63\%$ of schools are listed as some student’s first choice. Furthermore, simulations clearly indicate that the number of unassigned students depends on both the length of student lists and on school priorities.

2.4. Simulation. As a result of the limitations discussed above, existing matching models are often of limited use when trying to tackle practical problems. Parents want to know how likely their child is to be admitted to a particular school. Administrators want to predict how a proposed policy change will affect the number of students who fail to match to any school on their list. For these problems, the best tool is often simulation. Abdulkadiroglu et al. (2009) use simulations to compare different tiebreaking procedures in New York City and Boston, and de Haan et al. (2018) do the same for Amsterdam. Ashlagi and Nikzad (2019) use simulations based on data from New York to study the effect of changing priorities, and Kanoria et al. (2021) use the same data to study the effect of changing the length of student lists. While simulation is a flexible and accurate tool, it typically does not offer much *insight*. It shows what is true, but not *why* it is true. If a pattern is observed through simulation, it can be difficult to predict whether the same pattern will continue to hold in other settings.

3. MOTIVATING THE MODEL

To introduce and motivate my proposed definition of stability, I first present heuristic reasoning about Example 1. I then describe how this reasoning can be generalized beyond this example. In what follows, I say that a student is “admitted” to a school if her priority score is above that school’s cutoff, even if she chooses not to attend. I say that a student who matches to her k^{th} choice is “interested in” (or has “proposed to”) her top k choices. However, she is *not* interested in any school ranked below the one she matches to.

3.1. Analysis of Example 1. I now present a heuristic analysis of Example 1. I decompose this analysis into several steps.

3.1.1. Specify an admissions probability. By symmetry, each student is equally likely to be admitted to each school on his or her list. Let’s call this common admissions probability p .

3.1.2. Determine student outcomes from the admissions probability. A student matches to her k^{th} choice if and only if she is accepted to that school and rejected from her top $k - 1$ schools. For any student, outcomes should be approximately independent across schools:

being rejected from a top choice provides little information about the level of competition at other schools.

Thus, the fraction of students who match to their k^{th} choice should be approximately $p(1-p)^{k-1}$, and the answer to question I should be approximately

$$(1) \quad 1 - (1-p)^\ell.$$

3.1.3. *Determine expected “interest” received by each school.* Since a student who matches to her k^{th} choice sends k proposals and a student who goes unassigned sends ℓ proposals, it follows that the average number of proposals sent by students should be

$$\mu(p) = (1 - (1-p)^\ell)/p.$$

Because there are equal number of students and schools, the average number of proposals received by each school must also be $\mu(p)$, and because priorities are iid $U[0, 1]$, the expected number of proposals that a school receives from students with priority above s should be $\mu(p)(1-s)$.

3.1.4. *Determine Admissions Probabilities.* Turning to Question II, note that a student with priority s is admitted only if no higher-priority student is interested in the school. From above, the *expected* number of students with priority above s who propose to the school should be $\mu(p)(1-s)$. But what is the *distribution* of this number?

Before we have sampled student preferences, the probability that a student proposes to a particular school h will be small (bounded above by the probability that they list h , which is ℓ/n). A sum of nearly independent binary random variables each with a small mean is well-approximated by a Poisson distribution. Therefore, the number of proposals received by h from students with priority above s should approximately follow a Poisson distribution with mean $\mu(p)(1-s)$. A student with priority s will be admitted only if there are no such proposals, which should occur with probability

$$(2) \quad e^{-\mu(p)(1-s)}.$$

3.1.5. *Define a Consistency Condition.* The heuristic reasoning above starts from a hypothesized admissions probability p . But what should the value of p be? We can answer by noting that p should arise from integrating (2) over possible lottery draws:

$$p = \int_0^1 e^{-\mu(p)(1-s)} ds = (1 - e^{-\mu(p)})/\mu(p).$$

Multiplying both sides by $n\mu(p)$ yields the equation

$$(3) \quad n(1 - (1-p)^\ell) = n(1 - e^{-(1-(1-p)^\ell)/p}),$$

which can be interpreted as saying that the number of matched students (on the left) must equal the number of matched schools (on the right). It can be shown that (3) identifies a unique value of p . From there, the expressions in (1) and (2) answer questions I and II.

3.2. Generalizing this Approach.

Arnosti (2022) shows that this logic can be made rigorous, and generalized to accommodate schools with multiple seats, students who submit lists of varying lengths, and the case where school lotteries are perfectly correlated (rather than independent). However, that paper allows only those two priority structures, rendering it unable to truly answer Question III. Furthermore, it makes restrictive assumptions on student preferences. This paper shows how to generalize the reasoning above to arbitrary distributions of preferences and priorities.

In Section 3.1.1, we started by assuming an admissions probability p . In Example 1, p was the same for all schools, and represented an expectation across the student’s possible lottery draws. In general, the admissions probability should depend on both the specific school, and the student’s priority. Thus, I start from an *admissions function* which gives, for each school h and each possible priority score $p \in [0, 1]$, the probability that a student who has priority p at h will be admitted.

Section 3.1.2 used the hypothesized admissions function to predict how likely students are to match to each school on their list. This step assumed independence of outcomes across schools. In Example 1, priorities are determined by independent lotteries, but in general, priorities could be either positively correlated across schools (if based on exam scores or grades) or negatively correlated (if based on neighborhood). However, I will assume that *conditional* on a student’s priority scores, outcomes are independent across schools. This makes it possible to determine a *fractional matching* for each student from any hypothesized admissions function.

Section 3.1.3 used predicted match outcomes to determine, for each school, the expected number of proposals (“interest”) that it receives from students with priority above any given threshold. A student is interested in her k^{th} choice if and only if she was not admitted to any of her top $k - 1$ choices. I will use this observation to go from a fractional matching to a student *interest function* which specifies, for each school h and cutoff score $s \in [0, 1]$, the expected number of students with priority above s at h who will be interested in attending.

Section 3.1.4 shows how to predict a new *admissions function*: for a student with a given priority at a given school, what is her chance of getting in? For Example 1, I argued that the number of proposals received by each school should approximately follow a Poisson distribution. In general, I use a black box “vacancy function” which takes as input the capacity of the school and the expected number of higher-priority students, and generates as output a predicted probability of admission. Although many of my results hold for *any* vacancy function, in Section 6 I present predictions from using a particular vacancy function motivated by the Poisson distribution.

To summarize, the generalizable aspects of the analysis above are as follows. From any hypothesized admissions function, I generate predicted *match probabilities*, from which I determine *expected interest*, which can then be used to determine a new *admissions function*.

The final step imposes a consistency condition: this new admissions function should match the one that we started with. Any admissions function with this property is *stable*.

4. MODEL

There is a finite set of high schools \mathcal{H} . School $h \in \mathcal{H}$ has capacity $C_h \in \mathbb{N}$. Let $\mathcal{H}_0 = \mathcal{H} \cup \{\emptyset\}$ denote the set of schools along with the outside option of going unassigned, and define $C_\emptyset = \infty$. Let \mathcal{R} be the set of complete orders over \mathcal{H}_0 . There is no restriction on the number of acceptable schools for each student (i.e. the number of schools preferred to the outside option \emptyset), and the order of schools ranked below \emptyset will be irrelevant.

Students are characterized by their type $\theta = (\succ^\theta, \mathbf{p}^\theta)$, where $\succ^\theta \in \mathcal{R}$ indicates the student's preferences and $\mathbf{p}^\theta \in [0, 1]^{\mathcal{H}}$ indicates the student's priority score at school h (higher is better). Let $\Theta = \mathcal{R} \times [0, 1]^{\mathcal{H}}$ denote the space of student types. Students are distributed according to a positive finite measure η over Θ .

A fractional matching is a function M mapping each $\theta \in \Theta$ to a probability distribution on \mathcal{H}_0 . For each $h \in \mathcal{H}_0$ and $\theta \in \Theta$, the quantity $M_h(\theta)$ can be interpreted as the probability that a student of type θ is assigned to h . Hereafter, I use “matching” to mean a fractional matching, and denote the space of matchings by \mathfrak{M} .

I now define what it means for a matching to be *stable*. This definition uses two auxiliary concepts, which are based on the perspective of individual agents. What matters to each student is the set of schools that admit them. What matters to a school is the set of students who are “interested,” meaning that they would attend if admitted. In my model, these are described by

- An *admissions function* $A : [0, 1] \rightarrow [0, 1]^{\mathcal{H}_0}$.
- An *interest function* $I : [0, 1] \rightarrow \mathbb{R}_+^{\mathcal{H}_0}$.

Given $h \in \mathcal{H}$ and $p \in [0, 1]$, $A_h(p)$ can be interpreted as the probability that a student with priority p at h will be admitted, while $I_h(p)$ can be interpreted as the expected number of students whose priority at h exceeds p , and who are “interested” in h , meaning that they weakly prefer h to their assigned school. Given these interpretations, it is natural that A_h should be increasing and I_h should be decreasing. Let \mathfrak{A} denote the set of componentwise weakly increasing functions from $[0, 1]$ to $[0, 1]^{\mathcal{H}_0}$, and let \mathfrak{I} denote the set of componentwise weakly decreasing functions from $[0, 1]$ to $\mathbb{R}_+^{\mathcal{H}_0}$.

Next, I define consistency conditions that link a matching M to school interest I and student admissions decisions A . Formally, I define maps $\mathcal{I} : \mathfrak{M} \rightarrow \mathfrak{I}$, $\mathcal{A} : \mathfrak{I} \rightarrow \mathfrak{A}$ and $\mathcal{M} : \mathfrak{A} \rightarrow \mathfrak{M}$, and define a stable matching as a fixed point of the composition of these maps. This approach is illustrated in Figure 3.

Although our approach may seem unfamiliar, our definition subsumes existing ones. The function \mathcal{I} depends on the measure of student types η , and the function \mathcal{A} depends on a *vacancy function* \mathcal{V} . Section 5.1 shows that for a particular choice of η and \mathcal{V} , our definition of a stable matching coincides with the absence of blocking pairs in a finite market. Section

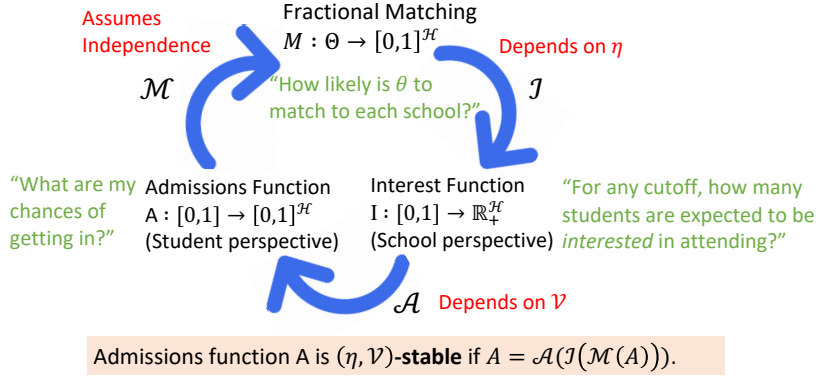


FIGURE 3. I present three ways to describe a random outcome, each of which is most relevant to different stakeholders. I present functions that link these descriptions, and define an outcome to be *stable* if it is a fixed points of the composition of these functions.

5.2 shows that when η is changed to a continuous measure, any matching that is stable according to our definition is associated with a set of market-clearing cutoffs, and vice versa.

4.1. Matching to Interest. Given any matching $M \in \mathfrak{M}$, define $\mathcal{I}(M) \in \mathfrak{I}$ to be the interest function I^M such that for each $h \in \mathcal{H}_0$ and $p \in [0, 1]$,

$$(4) \quad I_h^M(p) = \int \mathbf{1}(p_h^\theta \geq p) \left(1 - \sum_{h' \succ^\theta h} M_{h'}^\theta\right) d\eta(\theta).$$

Note that the sum in (4) gives the probability under matching M that student type θ matches to a school preferred to h , so the interpretation of (4) is that students are interested in h if they are not matched to any preferred school. The indicator ensures that the only students contributing to $I_h^M(p)$ are those with priority above p at h , allowing us to interpret $I_h^M(p)$ as the expected number of students with priority above p who are interested in h .

4.2. Interest to Admissions. The interest function I describes expected interest at each school $h \in \mathcal{H}$ and priority level $p \in [0, 1]$. From this, I will determine an admissions function $A : [0, 1] \rightarrow [0, 1]^{\mathcal{H}}$, where $A_h(p)$ is interpreted as the probability that a student with priority p at h will be admitted to h (equivalently, the probability that school h has a final cutoff below p). I define A using a *vacancy function* $\mathcal{V} : \mathbb{R}_+ \times \mathbb{N} \rightarrow [0, 1]$. Formally, let $\mathcal{A}(I)$ be the admissions function $A^I \in \mathfrak{A}$ that satisfies, for each $h \in \mathcal{H}$ and $p \in [0, 1]$,

$$(5) \quad A_h^I(p) = \mathcal{V}(I_h(p), C_h).$$

Define $A_\emptyset^I(p) = 1$ for all $p \in [0, 1]$ (students are always admitted to the outside option).

The choice of vacancy function is an important feature of the model, and one of the key innovations in this paper. The quantity $\mathcal{V}(\lambda, C)$ is interpreted as the probability that when *expected* interest is equal to λ , *realized* interest will be below C . Thus, if schools consider

students in descending order of priority, $\mathcal{V}(I_h(p), C_h)$ gives the probability that school h will still have at least one vacancy when it considers a student with priority p .

One natural choice of vacancy function is

$$(6) \quad \mathcal{V}^{det}(\lambda, C) = \mathbf{1}(\lambda < C) \quad \forall \lambda \in \mathbb{R}_+, C \in \mathbb{N},$$

In other words, realized interest is *deterministically* equal to expected interest, and there is still a vacancy if and only if expected interest is below the school's capacity. This choice produces a deterministic prediction for each student type θ and each school h . Sections 5.1 and 5.2 show that this choice of vacancy function recovers the definition of stability in a finite market from Gale and Shapley (1962), as well as that used by Azevedo and Leshno (2016).

One limitation of the deterministic vacancy function is that its predicted admissions probabilities are always zero or one. In a random matching market, there is uncertainty about where each student will be admitted. To capture this uncertainty, Section 6 uses an alternative choice of vacancy function, which assumes that when expected interest is equal to λ , realized interest follows a Poisson distribution with mean λ .

4.3. Admissions to Matching. Recall that an admissions function A describes the probability that a student of any given priority $p \in [0, 1]$ will be admitted to each school. From this, I construct an associated fractional matching $\mathcal{M}(A) = M^A$ given by

$$(7) \quad M_h^A(\theta) = A_h(p_h^\theta) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)).$$

This says that a student matches to h if and only if she is admitted to h and not to any preferred school. Note that this formula implicitly assumes independence of admissions outcomes across schools. A straightforward inductive argument implies that for any $A \in \mathfrak{A}$, $\theta \in \Theta$ and $h \in \mathcal{H}_0$,

$$(8) \quad 1 - \sum_{h' \succ^\theta h} M_{h'}^A(\theta) = \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)),$$

with both sides interpreted to be 1 if h is the first choice of θ .

4.4. Definition of Stability. The economy $\mathcal{E} = (\mathcal{H}, C, \eta)$ determines student preferences and priorities as well as school capacities. An additional key feature of the model is the vacancy function \mathcal{V} . Note that η determines the function $\mathcal{I} : \mathfrak{M} \rightarrow \mathfrak{I}$, while \mathcal{V} and C determine the function $\mathcal{A} : \mathfrak{I} \rightarrow \mathfrak{A}$. I use the following definition of stability.

Definition 1.

Matching $M \in \mathfrak{M}$ *is* $(\mathcal{E}, \mathcal{V})$ -**stable** *if* $M = \mathcal{M}(\mathcal{A}(\mathcal{I}(M)))$.

Admissions function $A \in \mathfrak{A}$ *is* $(\mathcal{E}, \mathcal{V})$ -**stable** *if* $A = \mathcal{A}(\mathcal{I}(\mathcal{M}(A)))$.

Interest function $I \in \mathfrak{I}$ *is* $(\mathcal{E}, \mathcal{V})$ -**stable** *if* $I = \mathcal{I}(\mathcal{M}(\mathcal{A}(I)))$.

Outcome $(M, I, A) \in \mathfrak{M} \times \mathfrak{I} \times \mathfrak{A}$ *is* $(\mathcal{E}, \mathcal{V})$ -**stable** *if* $M = \mathcal{M}(A)$, $I = \mathcal{I}(M)$, *and* $A = \mathcal{A}(I)$.

Definition 1 makes it clear that there is a one-to-one correspondence between stable matchings, stable admissions functions, stable interest functions, and stable outcomes. I include each of these definitions because it is sometimes most convenient to work with stable matchings, and at other times simpler to work with stable interest functions or stable outcomes.

5. RESULTS FOR GENERAL VACANCY FUNCTIONS

Our definition of stability in Definition 1 may seem strange to those familiar with more traditional definitions based on the absence of blocking pairs, or cutoffs that clear the market. It more closely resembles fixed-point characterizations of stable matchings by Adachi (2000), Fleiner (2003), and Echenique (2004). Our first results bridge this divide by showing that when using the deterministic vacancy function \mathcal{V}^{det} from (6), our definition encompasses more traditional definitions as special cases. Section 5.1 shows that in finite markets, our definition corresponds to the absence of blocking pairs. Section 5.2 shows that in continuum markets, our definition is equivalent to one based on market-clearing cutoffs.

While these results both assume that the vacancy function \mathcal{V} is as given in (6), we proceed to establish that for *any* η and \mathcal{V} , several classic results hold: the set of stable matchings is a non-empty lattice, the extreme points of this lattice can be found using the deferred acceptance algorithm, and the rural hospital theorem applies. Finally, we prove that if η has strict priorities and \mathcal{V} is strictly decreasing, there is a unique stable matching.

5.1. Finite Markets: Stability = No Blocking Pairs. Traditionally, stable matching problems involve a finite set of students $\mathcal{S} \subset \Theta$. We adopt the standard assumption that \mathcal{S} has “strict priorities”: no two students in \mathcal{S} have identical priority at any school. Given $h \in \mathcal{H}$ and $p \in [0, 1]$, define

$$(9) \quad \Theta_h(p) = \{\theta : h \succ^\theta \emptyset, p_h^\theta = p\}$$

to be the set of student types that consider school h acceptable and have priority p at school h .

Definition 2 (Strict Priorities).

*A finite subset $\mathcal{S} \subset \Theta$ has **strict priorities** if for each $h \in \mathcal{H}$ and $p \in [0, 1]$, $|\mathcal{S} \cap \Theta_h(p)| \leq 1$.*

*A positive measure η on Θ has **strict priorities** if for each $h \in \mathcal{H}$ and $p \in [0, 1]$, $\eta(\Theta_h(p)) = 0$.*

The second part of this definition is motivated by the study of *random* finite matching markets, where \mathcal{S} is generated by drawing student types iid from some measure η over Θ . In this case, the condition above ensures that \mathcal{S} has strict priorities with probability one.¹

¹The assumption that there are no ties is essential to many of our results. This is not an artifact of our definitions or proof techniques, but rather reflects fundamental challenges to defining stable matchings with indifferences.

We now give a version of the traditional definition of stability based on the absence of blocking pairs. We refer to this concept as “no blocking pairs” to distinguish it from the definition of stability in Definition 1.

Definition 3 (No Blocking Pairs). *Given any finite set $\mathcal{S} \subset \Theta$, an \mathcal{S} -matching is a function $\mu : \mathcal{S} \rightarrow \mathcal{H}_0$. An \mathcal{S} -matching μ is **feasible** if for each $h \in \mathcal{H}_0$,*

$$(10) \quad |\{\theta \in \mathcal{S} : \mu(\theta) = h\}| \leq C_h.$$

*An \mathcal{S} -matching μ **has no blocking pairs** if it is feasible, and for each $\theta' \in \mathcal{S}$ and each $h \in \mathcal{H}_0$ such that $h \succ^{\theta'} \mu(\theta')$,*

$$(11) \quad |\{\theta \in \mathcal{S} : \mu(\theta) = h, p_h^\theta > p_h^{\theta'}\}| = C_h.$$

Definition 3 states that a feasible \mathcal{S} -matching has no blocking pairs if for each student $\theta' \in \mathcal{S}$, each school that θ' prefers to its assignment is filled with higher-priority students. Note that this implies individual rationality: because $C_\emptyset = \infty$, the outside option is never filled to capacity. Therefore, if μ has no blocking pairs, then it does not assign any student to a school that she considers inferior to the outside option.

Our first result is to show that our definition of stability corresponds with the traditional definition based on the absence of blocking pairs. To state this result formally, we note that each finite set $\mathcal{S} \subset \Theta$ is naturally associated with an associated counting measure $\eta^\mathcal{S}$ over Θ , defined by

$$(12) \quad \eta^\mathcal{S}(\tilde{\Theta}) = |\tilde{\Theta} \cap \mathcal{S}| \quad \forall \tilde{\Theta} \subseteq \Theta.$$

Similarly, there is a natural correspondence between \mathcal{S} -matchings (which define an assignment only for student types in \mathcal{S}) and deterministic matchings (which define an assignment for all types in Θ). Any deterministic matching M naturally defines an \mathcal{S} -matching μ^M : for each $\theta \in \mathcal{S}$, let

$$(13) \quad \mu^M(\theta) = h \Leftrightarrow M_h(\theta) = 1.$$

Similarly, each \mathcal{S} -matching μ naturally induces a deterministic matching M^μ as follows. Define the admissions outcome A^μ by

$$(14) \quad A_h^\mu(p) = \mathbf{1}(|\{\theta \in \mathcal{S} : p_h^\theta > p, \mu(\theta) = h\}| < C_h),$$

and define $M^\mu = \mathcal{M}(A^\mu)$. In other words, (14) says that student $\theta \in \Theta$ is admitted to h if there are fewer than C_h higher-priority students from \mathcal{S} matched to h under μ , and $M^\mu(\theta)$ is the matching that results when each student type θ is assigned to its most-preferred school among those where it is admitted.

The following result says that if priorities are strict, then the functions $M \rightarrow \mu^M$ and $\mu \rightarrow M^\mu$ define a bijection between the set of $(\eta^\mathcal{S}, \mathcal{V}^{det})$ -stable matchings, and the set of \mathcal{S} -matchings with no blocking pairs. The proof of this result is deferred to Appendix A.1.

Proposition 1 (No Blocking Pairs Corresponds to a Stable Matching). *Let \mathcal{S} be a finite subset of Θ with strict priorities. If M is a $(\eta^{\mathcal{S}}, \mathcal{V}^{det})$ -stable matching, then μ^M has no blocking pairs. If μ is an \mathcal{S} -matching with no blocking pairs, then M^μ is $(\eta^{\mathcal{S}}, \mathcal{V}^{det})$ -stable, and $\mu^{M^\mu} = \mu$.*

5.2. Continuum Markets: Stability = Market-Clearing Cutoffs. Azevedo and Leshno (2016) provide a continuum model in which a market is described by a positive measure η over Θ and a stable matching is described by a vector of priority cutoffs $P \in [0, 1]^{\mathcal{H}}$. Students are admitted to school h if and only if their priority at h exceeds its cutoff P_h . They define demand for school h at cutoffs P by²

$$(15) \quad D_h(P) = \int \mathbf{1}(p_h^\theta > P_h) \prod_{h' \succ^\theta h} \mathbf{1}(p_{h'}^\theta \leq P_{h'}) d\eta(\theta).$$

That is, demand for h at cutoffs P is equal to the measure of students who are admitted to h and are not admitted to any school that they prefer to h . Each cutoff vector is naturally associated with a deterministic matching in which students attend the school that they demand. The definition of stability used by Azevedo and Leshno (2016) is that the cutoff vector should clear the market.

Definition 4. A cutoff $P \in [0, 1]^{\mathcal{H}}$ is η -**market-clearing** if $D_h(P) \leq C_h$ for all $h \in \mathcal{H}$, with equality if $P_h > 0$.

In this section, we show that in continuum markets, a cutoff vector P is η -market-clearing if and only if a corresponding interest function is $(\mathcal{E}, \mathcal{V}^{det})$ -stable. To formalize this claim, we first define a natural associate between cutoff vectors and interest functions. Each cutoff vector P is naturally associated with an interest function I^P defined for each $h \in \mathcal{H}$ and $p \in [0, 1]$ by

$$(16) \quad I_h^P(p) = \int \mathbf{1}(p_h^\theta > p) \prod_{h' \succ^\theta h} \mathbf{1}(p_{h'}^\theta \leq P_{h'}) d\eta(\theta).$$

That is, students contribute to this quantity if they have priority above p at h and do not “clear the cutoff” at any school that they prefer. Note that when $p = P_h$, we have $I_h^P(P_h) = D_h(P)$.

Conversely, from any interest function $I \in \mathfrak{I}$, we can define the associated cutoffs $\mathcal{P}(I) = \{\mathcal{P}_h(I)\}_{h \in \mathcal{H}} \in [0, 1]^{\mathcal{H}}$ by

$$(17) \quad \mathcal{P}_h(I) = \inf\{p \geq 0 : I_h(p) < C_h\}.$$

²An astute and informed reader might notice that our choice of A^P assumes that student types θ with $p_h^\theta = P_h$ are not admitted to h , whereas Azevedo and Leshno (2016) assume that they are admitted. Because η is a continuous measure with strict priorities in both cases, this distinction is of consequence only to sets of η -measure zero.

Equation (16) defines a mapping from cutoffs to interest functions, while (17) defines mapping from interest functions to cutoffs. It turns out that these mappings take stable interest functions to market-clearing cutoffs, and vice versa.³

Proposition 2 (Market-Clearing Cutoffs Correspond to Stable Interest Functions).

Let η have strict priorities. If P is η -market-clearing, then I^P is $(\mathcal{E}, \mathcal{V}^{det})$ -stable. If I is $(\mathcal{E}, \mathcal{V}^{det})$ -stable, then $\mathcal{P}(I)$ is η -market-clearing, and $I = I^{\mathcal{P}(I)}$.

We prove this result in Appendix A.2.

5.3. Existence and Lattice Structure. Having established that when $\mathcal{V} = \mathcal{V}^{det}$, our definition of stability nests existing definitions, we now prove results for general type measures η and vacancy functions \mathcal{V} . The first of these results shows that stable matchings always exist and form a lattice. To state this result, we define the following partial orders:

- $M \succeq^{\mathfrak{M}} \tilde{M}$ if for each $h \in \mathcal{H}_0$ and $\theta \in \Theta$,

$$\sum_{h' \succeq^{\theta} h} M_{h'}(\theta) \geq \sum_{h' \succeq^{\theta} h} \tilde{M}_{h'}(\theta).$$

That is, $M \succeq^{\mathfrak{M}} \tilde{M}$ if each student prefers M to \tilde{M} in the sense of first-order stochastic dominance.

- $A \succeq^{\mathfrak{A}} \tilde{A}$ if for each $h \in \mathcal{H}_0$ and $p \in [0, 1]$, $A_h(p) \geq \tilde{A}_h(p)$.

That is, $A \succeq^{\mathfrak{A}} \tilde{A}$ if admissions probabilities are uniformly higher under A .

- $I \succeq^{\mathfrak{I}} \tilde{I}$ if for each $h \in \mathcal{H}_0$ and $p \in [0, 1]$, $I_h(p) \geq \tilde{I}_h(p)$.

That is, $I \succeq^{\mathfrak{I}} \tilde{I}$ if each school receives more interest at every cutoff under I .

- $(M, A, I) \succeq (\tilde{M}, \tilde{A}, \tilde{I})$ if $M \succeq^{\mathfrak{M}} \tilde{M}$, $A \succeq^{\mathfrak{A}} \tilde{A}$, and $I \succeq^{\mathfrak{I}} \tilde{I}$.

Theorem 1 (Existence and Lattice Structure). *If the vacancy function \mathcal{V} is weakly decreasing in its first argument, then for any $(\mathcal{H}, \mathbf{C}, \eta)$, the set of $(\mathcal{E}, \mathcal{V})$ -stable outcomes is non-empty, and forms complete lattice with partial order \succeq .*

Proof of Theorem 1. Define the function $\xi : \mathfrak{M} \rightarrow \mathfrak{M}$ by

$$(18) \quad \xi(M) = \mathcal{M}(\mathcal{A}(\mathcal{I}(M))).$$

Note that

- By (4), $\tilde{M} \succeq^{\mathfrak{M}} M$ implies $\mathcal{I}(\tilde{M}) \succeq^{\mathfrak{I}} \mathcal{I}(M)$.
- By (5) and monotonicity of \mathcal{V} , $I \succeq^{\mathfrak{I}} \tilde{I}$ implies $\mathcal{A}(\tilde{I}) \succeq^{\mathfrak{A}} \mathcal{A}(I)$.

³We briefly comment on a subtlety that explains why Proposition 2 is stated in terms of the interest function I^P rather than the admissions function A^P or the matching M^P . In general, multiple market-clearing cutoffs may correspond to the same stable matching (up to a set of measure zero). For example, suppose that there is a single school h with capacity C , and that the total measure of students is $\eta(\Theta) = 2C$, with priorities uniformly distributed on $(0, 1/3) \cup (2/3, 1)$. Then any cutoff $P \in [1/3, 2/3]$ clears the market. Our definition of stability eliminates this redundancy: the unique $(\mathcal{E}, \mathcal{V}^{det})$ -stable matching corresponds to a cutoff of $2/3$ and leaves students in $[0, 2/3]$ unassigned. Thus, if $P \in [1/3, 2/3)$, P clears the market but M^P is not $(\mathcal{E}, \mathcal{V}^{det})$ -stable. By contrast, for any P , $I^P(p) = \eta(\{\Theta : p_h^\Theta > p\})$ is a stable interest function.

- By (7), $\tilde{A} \succeq^{\mathfrak{A}} A$ implies $\mathcal{M}(\tilde{A}) \succeq^{\mathfrak{M}} \mathcal{M}(A)$.

From this, we draw two conclusions. First, if (M, I, A) and $(\tilde{M}, \tilde{I}, \tilde{A})$ are stable outcomes, then $(M, I, A) \succeq (\tilde{M}, \tilde{I}, \tilde{A})$ if and only if $M \succeq^{\mathfrak{M}} \tilde{M}$. Second, the function ξ is an order preserving function, so Tarski's fixed point theorem implies that the set of fixed points of ξ (that is, the set of stable matchings) forms a complete lattice with respect to $\succeq^{\mathfrak{M}}$ (and in particular is non-empty). \square

5.4. Deferred Acceptance Algorithm. Theorem 1 establishes the existence of stable outcomes, but does not address how to find them. However, the proof suggests a natural procedure: start from a matching M and repeatedly apply the function ξ defined by $\xi(M) = \mathcal{M}(\mathcal{A}(\mathcal{I}(M)))$. If one starts from the matching \overline{M} which assigns each student to her most preferred school, then this procedure corresponds to the student-proposing deferred acceptance algorithm, and converges to the student-optimal stable matching. To see that it converges, note that $\xi(\overline{M}) \preceq^{\mathfrak{M}} \overline{M}$, from which the fact that ξ is order-preserving implies that the sequence $\{\xi^k(\overline{M})\}_{k=0}^{\infty}$ is decreasing. Therefore, it converges by completeness of \mathfrak{M} . Conversely, repeatedly applying ξ from the student-pessimal matching \underline{M} (defined by $\underline{M}_\theta(\theta) = 1$ for all θ) produces an increasing sequence of matchings that converges to the school-optimal stable matching.

Because of the correspondence between stable matchings, stable interest functions, and stable admissions functions, it is also possible to apply an analogous iterative process using the admissions function as the primitive of interest. In that case, one could start from $A_h(p) = 1$ for all h and p (resulting in convergence to the student-optimal stable matching) or $A_h(p) = 0$ for all h and p (resulting in convergence to the student-pessimal stable matching). Figure 4 shows the sequence of admissions functions that result when this iterative process is applied to an example from Azevedo and Leshno (2016).

Although convergence is guaranteed, in general it does not occur in finitely many steps. In examples that I have tried, convergence happens quickly enough that this algorithm can be applied fruitfully. The main practical challenge is computing $\mathcal{I}(M)$, which requires taking an integral over student types. Although this may be challenging for arbitrary measures η , it is tractable for many cases of interest.

5.5. Rural Hospital Theorem. I now establish a “rural hospital theorem,” which states that for any two stable matchings, each student's probability of assignment and each school's measure of matched students is identical. This is a generalization of the corresponding result for finite markets, proved by McVitie and Wilson (1970) and Roth (1986).

Theorem 2 (Rural Hospital Theorem). *If η has strict priorities and \mathcal{V} is weakly decreasing in its first argument, then the set of matched agents is identical across stable outcomes: if (M, I, A) and $(\tilde{M}, \tilde{I}, \tilde{A})$ are $(\mathcal{E}, \mathcal{V})$ -stable outcomes, then for each $h \in \mathcal{H}_0$,*

$$(19) \quad \int M_h(\theta) d\eta(\theta) = \int \tilde{M}_h(\theta) d\eta(\theta),$$




FIGURE 4. Finding a stable admissions function for the example from Azevedo and Leshno (2016), using vacancy function \mathcal{V}^{pois} defined in (22). Because schools are symmetric, $A_h = A_{h'}$ for $h, h' \in \mathcal{H}$. We iterate the map $A \rightarrow \mathcal{A}(\mathcal{I}(\mathcal{M}(A)))$. Starting from $A(p) = 1$ results in a decreasing sequence of admissions functions (in blue) converging to the student-optimal stable admissions function. Starting from $A(p) = 0$ results in an increasing sequence of admissions functions (in green) converging to the school-optimal stable admissions function. Because \mathcal{V}^{pois} is strictly decreasing, Theorem 3 guarantees that these admissions functions coincide. Notice that the admissions function is the CDF of the cutoff distribution whose 5th and 95th percentiles are shown in Figure 1.

and for each $\theta \in \Theta$ outside of a set of η -measure zero,

$$(20) \quad \sum_{h \succ^{\theta} \emptyset} M_h(\theta) = \sum_{h \succ^{\theta} \emptyset} \tilde{M}_h(\theta).$$

In contrast to the existence result in Theorem 1, Theorem 2 requires an assumption of strict priorities. This assumption is essential for the result to hold.⁴ The failure of the rural hospital theorem when there are ties in priority is not specific to my definition of stability: in finite markets with indifferences, it is known that strongly stable matchings may not exist

⁴To see that the conclusion of Theorem 2 may fail to hold if η does not have strict priorities, consider an example with two schools, A and B , each with a single seat. The measure η corresponds to a finite market with three students, x, y, z . Students x and y prefer A to B , while student z prefers B to A . Student x has priority $1/4$ at school A and $3/4$ at school B . Student y has priority $1/4$ at school A and $2/4$ at school B . Student z has priority $3/4$ at school A and $1/4$ at school B .

We claim that there are two $(\mathcal{E}, \mathcal{V}^{det})$ -stable matchings, and that y is assigned in one and unassigned in the other. In the school-optimal stable matching, x goes to B , z goes to A , and y is unassigned. In the student-optimal stable matching, x and y go to A , and z goes to B . Note that the student-optimal stable matching is infeasible (two students are assigned to A). This illustrates that our definition of stability (which was intended for markets with strict priorities) does not enforce capacity constraints in markets with ties.

(Irving, 1994), and weakly stable ones may not satisfy the rural hospital theorem (Manlove, 1999).

5.6. Uniqueness. Finally, I establish conditions under which there is a unique stable outcome.

Theorem 3 (Uniqueness). *If η has strict priorities and \mathcal{V} is strictly decreasing in its first argument, then there is a unique $(\mathcal{E}, \mathcal{V})$ -stable outcome.*

The intuition underlying this result is as follows. By Theorem 1, there are student-optimal and student-pessimal stable admissions functions A and \tilde{A} , with $A \succeq \tilde{A}$. It follows that all students will be weakly more likely to match under A . If \mathcal{V} is strictly decreasing, then $A \succ \tilde{A}$ implies that some students will be strictly more likely to match under \tilde{A} . This contradicts the rural hospital theorem, implying that we must have $A = \tilde{A}$. The complete proof is provided in Appendix A.3.

If \mathcal{V} is only weakly decreasing, it is possible that there are multiple stable matchings that lead to different outcomes for a positive η -measure of students: see Azevedo and Leshno (2016) for an example with $\mathcal{V} = \mathcal{V}^{det}$. However, their Theorem 1 shows that even in this case, there is typically a unique stable matching: this holds if η has full support, or for a generic set of school capacities.

6. RESULTS FOR THE POISSON VACANCY FUNCTION

Theorems 1, 2, 3 present results for general vacancy functions, while Propositions 1 and 2 show that prior definitions of stability correspond to the use of a deterministic vacancy function. The “secret sauce” of my alternative approach is to use a different vacancy function which produces probabilistic predictions.

My goal is to predict outcomes in a random market with n students whose types are sampled iid from a probability measure $\tilde{\eta}$ over Θ . For example, given a student who knows her own type but not the types of others, I wish to predict the student’s interim probability of being assigned to each school on her list. Which vacancy function should be used? This section proposes vacancy functions motivated by the Binomial and Poisson distributions, and shows that the resulting expressions can be used to reproduce and generalize insights from Ashlagi et al. (2017).

To motivate the Binomial and Poisson distributions, consider first how we could use the model of Azevedo and Leshno (2016) to generate a prediction. If we define the measure η by $\eta(S) = n\tilde{\eta}(S)$ and use a deterministic vacancy function, we get their predicted cutoffs for each school. These cutoffs partition the student type space. Let $\Theta_h(P)$ denote the set of types that are assigned to h when cutoffs are P . That is, $\Theta_h(P)$ consists of students with priority above P_h at h , and priority below $P_{h'}$ at every h' that they prefer to h .

Note that when student types are sampled iid, the number of students with a type in $\Theta_h(P)$ will follow a Binomial distribution with parameters n and $\tilde{\eta}(\Theta_h(P))$. If expected demand for

h is $\lambda = n\tilde{\eta}(\Theta_h(P))$ and capacity is C , the probability that demand will be less than capacity (that is, the probability that h will have a vacancy) is

$$(21) \quad \mathcal{V}^{bin}(\lambda, C) = \sum_{k=0}^{C-1} \binom{n}{k} \left(\frac{\lambda}{n}\right)^k \left(1 - \frac{\lambda}{n}\right)^{n-k}.$$

The uncertainty in realized demand will translate to uncertainty in h 's cutoff score. Although Azevedo and Leshno (2016) generate binary predictions, the true admissions probabilities should vary continuously: a student whose priority is a bit above the predicted cutoff score is not assured a spot, and one with priority below the predicted cutoff still has a chance of admission.

This fact could be captured by using the binomial vacancy function in (21) rather than the deterministic vacancy function in (6). If the number of students n is large relative to the demand at each school λ , then the Binomial distribution is well-approximated by a Poisson distribution, inspiring the following slightly simpler vacancy function:

$$(22) \quad \mathcal{V}^{pois}(\lambda, C) = \sum_{k=0}^{C-1} \frac{e^{-\lambda} \lambda^k}{k!}.$$

The remainder of this section presents predictions when using the vacancy function \mathcal{V}^{pois} .

In the limit as the number of students and the capacity of each school grow, the predictions of this new model converge to those of Azevedo and Leshno (2016).⁵ This is because Poisson random variables with large means are highly concentrated. The advantage of my model is not in this limit, but rather lies in its ability to generate probabilistic predictions for markets where capacities are modest and cutoffs are uncertain.

I will use this model to generate predictions about two quantities of interest: the number of matched students and the average rank of matched students (both defined formally below). I compare these predictions to existing analytical and simulation results for finite markets. To predict outcomes from random finite markets with n students whose types are drawn iid from a probability distribution $\tilde{\eta}$ on Θ , I define the measure η by $\eta(\tilde{\Theta}) = n\tilde{\eta}(\tilde{\Theta})$ for all $\tilde{\Theta} \subseteq \Theta$, and study $(\mathcal{E}, \mathcal{V}^{pois})$ -stable matchings.

Section 6.1 provides results for students' average rank, using the results of Ashlagi et al. (2017) as the primary comparison. Section 6.2 provides results for the number of matches, and compares against findings from Marx and Schummer (2021). In both cases, the model with a Poisson vacancy function accurately predicts simulation results for finite markets of

⁵To be more precise, given a market $\mathcal{E} = (\mathcal{H}, \mathbf{C}, \eta)$, define a sequence of markets $\mathcal{E}^m = (\mathcal{H}, \mathbf{C}^m, \eta^m)$, with $C_h^m = mC_h$ for all $h \in \mathcal{H}$ and $\eta^m(\tilde{\Theta}) = m \cdot \eta(\tilde{\Theta})$ for all measurable $\tilde{\Theta} \subseteq \Theta$. That is, we simply scale up the number of students and number of seats per school by a factor of m . This is the same limiting regime as considered by Azevedo and Leshno (2016). If η has full support, then Azevedo and Leshno (2016) show that there is a unique $(\mathcal{E}, \mathcal{V}^{det})$ -stable matching M^* . In this case, the unique $(\mathcal{E}, \mathcal{V}^{pois})$ -stable matching of \mathcal{E}^m converges to M^* as m grows. Intuitively, this is because Poisson random variables with large means are highly concentrated. More formally, for any $C \in \mathbb{N}$ and any $\lambda \in \mathbb{R}_+ \setminus \{C\}$, $\mathcal{V}^{pois}(m\lambda, mC) \rightarrow \mathcal{V}^{det}(\lambda, C)$ as $m \rightarrow \infty$.

moderate size. In addition, this model provides new analytical expressions and insights, described in more detail below.

6.1. Average Rank. I now present results on students' average rank, and compare to simulations and numerical bounds from Ashlagi et al. (2017). Denote θ 's rank of $h \in \mathcal{H}$ by

$$R_h(\theta) = |\{h' \in \mathcal{H}_0 : h' \succeq^\theta h\}|$$

and define⁶

$$(23) \quad \text{AverageRank}(M) = \frac{\int \sum_{h \in \mathcal{H}} M_h(\theta) R_h(\theta) d\eta(\theta)}{\int \sum_{h \in \mathcal{H}} M_h(\theta) d\eta(\theta)}.$$

Ashlagi et al. (2017) simulate outcomes in finite markets with 40 schools, each with a single seat, and a varying number of students. Each student ranks all 40 schools in a uniformly random order, and priorities are iid uniform across schools. To generate predictions for this market, I take η to be a uniform measure over complete preference profiles with $\eta(\Theta)$ equal to the total number of students, and use the Poisson vacancy function \mathcal{V}^{pois} .

There are several reasons that the model's predictions might not match the simulation results. First, Theorem 3 implies that the model generates a unique prediction, whereas finite markets may have multiple stable matchings. Thus, the "prediction error" must at least be comparable to the variation across stable matchings. Second, the assumption of independent outcomes across schools introduces error: when the number of students is below 40, every student in the finite market must match, whereas our model predicts that each student has a positive (albeit small) probability of going unassigned.

Despite these concerns, the model's predictions are excellent. Figure 2 displays simulations from Ashlagi et al. (2017) alongside model predictions. The curves do not merely appear qualitatively similar, they also match quantitatively. To emphasize this point, I reproduce their Table 1 in Figure 5. In their simulations, the average rank of students differs by at most 2% between student-optimal and school-optimal stable matches. My model's predictions always lie in this interval.

In addition to providing excellent quantitative predictions, my model can be used to derive analytical insights. I illustrate the point by providing analytical bounds on students' average rank. Figure 2 makes it clear that the behavior of the market is very different depending on whether the number of students is less or greater than the number of seats. Accordingly, the

⁶Note that (23) counts the average rank among matched students. I make this choice to align with the convention adopted by Ashlagi et al. (2017). One could instead work with the average rank across *all* students, counting a student who lists ℓ schools and goes unassigned as receiving their $(\ell + 1)^{st}$ choice. In general, these definitions are incomparable (either can be larger than the other). However, Proposition 3 holds for both definitions. To see this, note that

$$\frac{\int \sum_{h \in \mathcal{H}_0} M_h(\theta) R_h(\theta) d\eta(\theta)}{\eta(\Theta)} = \frac{\sum_{h \in \mathcal{H}} I_h(0) + I_\emptyset(0)}{\sum_{h \in \mathcal{H}} \int M_h(\theta) d\eta(\theta) + I_\emptyset(0)} \leq \frac{\sum_{h \in \mathcal{H}} I_h(0)}{\sum_{h \in \mathcal{H}} \int M_h(\theta) d\eta(\theta)}.$$

$ \mathcal{S} = \mathcal{H} + 10$	100	200	500	1000	2000	5000
Student-Optimal	29.5	53.6	115.8	203.8	364.5	793.1
School-Optimal	30.1	54.7	118.0	207.5	370.8	804.7
AKL Bound	25.3	45.7	98.2	175.2	314.6	690.5
Model Prediction	29.6	53.9	115.8	205.5	366.1	793.4

FIGURE 5. A reproduction of the first column of Table 1 from Ashlagi et al. (2017), alongside my model predictions. Each column corresponds to a different market size, and there are always 10 more students than schools. The first two rows show the (simulated) average rank under the student optimal and school optimal stable matches, which differ by at most 2%. The third row shows lower bounds from Ashlagi et al. (2017), which underestimate the average rank by approximately 15%. The final row shows my model's predictions, which always lie between the simulation results for the extremal stable matchings.

analysis will consider these cases separately. Define ρ to be the ratio of students to schools:

$$(24) \quad \rho = \eta(\Theta)/|\mathcal{H}|,$$

and let C denote the (common) capacity at each school. I analyze the case with more seats ($\rho < C$) in Section 6.1.1, and the case with more students ($\rho > C$) in Section 6.1.2.

I focus on two special cases considered by Ashlagi et al. (2017): independent lotteries across schools (“IID markets”) and a single lottery used by all schools (“RSD markets”).

Definition 5. *The measure η describes a **symmetric IID market** if (i) the restriction of \succ^θ to \mathcal{H} is uniformly distributed, and (ii) for each $\succ \in \mathcal{R}$, the conditional distribution of \mathbf{p}^θ given $\succ^\theta = \succ$ is uniform on $[0, 1]^\mathcal{H}$.*

*The measure η describes a **symmetric RSD market** if (i) the restriction of \succ^θ to \mathcal{H} is uniformly distributed, and (ii) for each $\succ \in \mathcal{R}$, the conditional distribution of \mathbf{p}^θ given $\succ^\theta = \succ$ is uniform on $\{\theta : p_h^\theta = p_{h'}^\theta \text{ for all } h, h' \in \mathcal{H}\}$.*

6.1.1. *More Seats than Students.* Define

$$(25) \quad \text{Enrollment}(\lambda, C) = \int_0^\lambda \mathcal{V}^{\text{pois}}(x, C) dx,$$

and note that the definition of $\mathcal{V}^{\text{pois}}$ in (22) implies that for any $C \in \mathbb{N}$,

$$(26) \quad \int_0^\infty \mathcal{V}^{\text{pois}}(\lambda, C) = C.$$

For any $\rho < C$, define $\Lambda(\rho, C)$ by the equation

$$\text{Enrollment}(\Lambda(\rho, C), C) = \rho.$$

Proposition 3. *Fix $C \in \mathbb{N}$, and let $C_h = C$ for all $h \in \mathcal{H}$. Let η^{IID} describe a symmetric iid market, and let M^{IID} be the unique $(\eta^{\text{IID}}, \mathcal{V}^{\text{pois}})$ -stable matching guaranteed by Theorem*

3. If there are more seats than students ($\rho < C$), then

$$\text{AverageRank}(M^{IID}) \leq \Lambda(\rho, C)/\rho.$$

To clarify the relationship with results from Ashlagi et al. (2017), I parameterize ρ and apply Proposition 3 to the special case where $C = 1$.

Corollary 1. *In a symmetric iid market where schools have a single seat and $\rho = \frac{n}{n+k} < 1$,*

$$\text{AverageRank}(M^{IID}) \leq \frac{n+k}{n} \log \left(\frac{n+k}{k} \right).$$

This upper bound exactly matches that from Theorem 2 of Ashlagi et al. (2017). Because we work with different models, neither result directly implies the other. However, the bound in Proposition 3 provides insight beyond the cases considered by Ashlagi et al. (2017). First, it allows for arbitrary school capacity, and improves as school capacity grows: fixing the ratio of students to seats $\rho/C = 0.97$, the bound on average rank is approximately 3.6 when $C = 1$, 2.0 when $C = 3$, 1.4 when $C = 10$, and tends to 1 as $C \rightarrow \infty$. Second, my bound does not assume that students submit complete (or long) lists: the distribution of list lengths can be arbitrary.⁷ Instead, fixing school capacity C , the bound depends only on the ratio of students to schools ρ .

6.1.2. *More Students than Seats.* I now turn to the case where students outnumber seats.

Proposition 4. *Suppose that $C_h = C_{h'}$ for all $h, h' \in \mathcal{H}$ and that there are more students than seats ($\rho > C$). Let η^{IID} describe a symmetric IID market in which all students list ℓ schools, and let M^{IID} be the unique $(\eta^{IID}, \mathcal{V}^{pois})$ -stable matching. Then*

$$\text{AverageRank}(M^{IID}) \geq \ell \left(1 - \frac{\rho}{C} - \frac{1}{\log(1 - C/\rho)} \right).$$

Let η^{RSD} describe a symmetric RSD market in which all students list ℓ schools, and let M^{RSD} be the unique $(\eta^{RSD}, \mathcal{V}^{pois})$ -stable matching. Then

$$\text{AverageRank}(M^{RSD}) \leq 1 + \log(\ell).$$

To facilitate a comparison with bounds from Ashlagi et al. (2017), I parameterize ρ and consider the special case where $C = 1$.

Corollary 2. *In a symmetric iid market where $C = 1$, $\rho = \frac{n+k}{n}$ and all students list n schools,*

$$\text{AverageRank}(M^{IID}) \geq \frac{n}{\log \left(\frac{n+k}{k} \right)} - k.$$

⁷While it is well known that increasing a student's list makes outcomes worse for all other students, this does not imply that extending a list increases the average rank, because average rank is calculated only for matched students. If some students submit short lists, extending their list may cause these students to match in place of (or in addition to) others with longer lists, thereby decreasing the average rank.

Proposition 4 and Corollary 2 imply the now well-known result that RSD priorities result in a much better average rank than IID priorities when students outnumber seats. In addition, they give quantitative bounds on the performance of each approach. The lower bound in Corollary 2 is very similar to the bound of $\frac{n}{1 + \frac{n+k}{n} \log(\frac{n+k}{k})}$ from Theorem 2 of Ashlagi et al. (2017) (in fact, my bound is tighter for $k \leq n$).⁸

More importantly, Proposition 4 clarifies the effect of list length, school capacity, and market size. In the model of Ashlagi et al. (2017), n simultaneously represents the number of school seats, number of schools, and the length of student lists. It is not clear how their bound changes if students list only a subset of the market, or if schools have multiple seats. By contrast, Proposition 4 establishes that students' average rank under RSD is logarithmic in the *list length* (rather than the market size). With IID priorities, students' average rank is linear in the list length, with a constant that depends on the ratio of students to seats ρ/C . This implies that with IID priorities, large capacities don't result in meaningfully better outcomes: given a fixed ratio of students to seats ρ/C , the lower bound does not depend on whether schools are small or large.⁹

6.2. Number of Matches. Another metric of interest is the number of matches. Few theoretical papers study this quantity, despite its salience in practice. In fact, many papers assume that students submit complete lists, or at least lists that are long enough that the short side of the market matches fully.

One recent exception is Marx and Schummer (2021). They consider the problem facing a matching platform that helps to match men and women, and charges both sides for each match. A man and woman can only be matched if they are both willing to pay the fee. The platform thus faces a tradeoff: if its prices are too high, there will be few mutually acceptable pairs, and few matches will form. The goal of the platform is to choose prices to maximize its revenue.

The main technical challenge they confront is analyzing the number of matches that form at given prices. They consider two matching algorithms. In the first, both sides rank acceptable partners, and a stable match is selected. Let $V^{IID}(W, M, q)$ be the expected number of matches when using this procedure with W women, M men, and probability of mutual compatibility $1 - q$. They argue that it is more tractable to analyze an alternative procedure in which men declare all acceptable partners, and then women are placed in a random order and sequentially allowed to choose their favorite unmatched mutually acceptable man. This is in essence a woman-selecting random serial dictatorship. Marx and Schummer (2021) show

⁸Letting $\beta = k/n$, algebra reveals that the bound in Corollary 2 is tighter than the bound from Theorem 2 of Ashlagi et al. (2017) so long as $(1 + \beta)\beta \log^2(1 + \beta) \leq 1$, which holds for $\beta \leq 1$.

⁹This is in contrast to RSD. In this case, fixing the ratio of students to seats ρ/C , the average rank is decreasing in C , and converges to one as C grows (although this fact is not reflected in the bound in Proposition 4).

that the expected number of matches that form in this case is

$$(27) \quad V^{RSD}(W, M, q) = \sum_{j=1}^{\min(M, W)} \frac{\prod_{i=0}^{j-1} (1 - q^{M-i}) \prod_{i=0}^{j-1} (1 - q^{W-i})}{1 - q^j}.$$

Figure 6 includes two plots from their paper: one uses (27) to plot V^{RSD} , and the second compares V^{RSD} and V^{IID} using simulation. Based on the second plot, Marx and Schummer (2021) argue that $V^{RSD}(W, M, q)$ is a reasonable approximation of $V^{IID}(W, M, q)$.

My model can be used to generate accurate and tractable approximations to both quantities. Let W be the mass of students and M be the number of schools, each with capacity $C = 1$. The cases studied by Marx and Schummer (2021) correspond to IID and RSD markets, as defined in Definition 5. These cases have also recently been studied by Arnosti (2022), whose expressions coincide with special cases of the model presented in this paper.

Under the assumptions of Marx and Schummer (2021), the length of student lists follows a binomial distribution, which satisfies the convexity condition in Theorem 3 of Arnosti (2022). This implies that the number of matched students in the unique $(\eta^{RSD}, \mathcal{V}^{pois})$ stable matching is lower than in the corresponding $(\eta^{IID}, \mathcal{V}^{pois})$ stable matching. In other words, the procedure that Marx and Schummer (2021) use to approximate the size of a stable matching should offer a lower bound.

The work of Arnosti (2022) can also be used to provide tractable closed-form approximations to the number of matches under each procedure. In Marx and Schummer (2021), the number of mutually compatible partners for each woman follows a binomial distribution with parameters M and $(1 - q)$. If we approximate this by a Poisson distribution with mean $M(1 - q)$ then rearranging the expressions in Proposition 2 in Arnosti (2022) yields the following approximations when $W \leq M$:¹⁰

$$(28) \quad \hat{V}^{RSD}(W, M, q) = W - \frac{\log(1 + e^{-(M-W)(1-q)} - e^{-M(1-q)})}{1 - q}.$$

$$(29) \quad (1 - q)\hat{V}^{IID}(W, M, q) = \log \left(1 - \frac{\hat{V}^{IID}(W, M, q)}{W} \right) \log \left(1 - \frac{\hat{V}^{IID}(W, M, q)}{M} \right).$$

Note that the expression in (28) is simpler than that in (27), and more amenable to optimization. Furthermore, whereas exactly calculating $V^{IID}(W, M, q)$ is intractable, (29) gives a concise closed-form expression relating the probability of incompatibility q and the match rate \hat{V} . Figure 6 shows predictions from using (28) and (29) alongside the original graphs from Marx and Schummer (2021). Despite the error in approximating the binomial distribution with a Poisson, the graphs closely match.

¹⁰In the model of Marx and Schummer (2021), the choice $W \leq M$ is without loss of generality, and cleans up the expression in (28) by allowing the use of W in place of $\min(W, M)$ and M in place of $\max(W, M)$.




FIGURE 6. Studying the number of matches that form, with 50 participants on one side and varying the number of participants on the other side. First panel: exact expressions from Marx and Schummer (2021) for V^{RSD} (each line represents a different probability of incompatibility q). Second panel: the corresponding predictions \hat{V}^{RSD} from my model. Third panel: Marx and Schummer (2021) have no analytical expression for V^{IID} , but their simulations study the relative difference $(V^{IID} - V^{RSD})/V^{RSD}$. Fourth panel: my model’s predicted relative difference $(\hat{V}^{IID} - \hat{V}^{RSD})/\hat{V}^{RSD}$, which can be proven to be positive. A similar plot cannot be made using the continuum model of Azevedo and Leshno (2016), which predicts that $V^{IID} = V^{RSD}$.

7. CONCLUSION

Stable matching algorithms are used to assign students to schools in cities across the globe. In theory, the design of school priorities offers a flexible tool for encoding policy objectives. In practice, the benefits of designing priorities are limited by the fact that the relationship between priorities and the final outcome is complex and poorly understood.

This paper offers a new perspective on stable matching, which enables the study of settings – such as those with small and asymmetrical schools – that cannot be readily studied using prior techniques. My model has three desirable features: it is *flexible* enough to accommodate complex preferences and priorities, its numerical predictions are extremely *accurate*, and it is tractable enough to offer new *insights*. We use a novel framework for stable matching to show that the only difference between our model and that of Azevedo and Leshno (2016) is that they assume that interest at each school is deterministic, whereas we assume that it follows a Poisson distribution. This difference allows our model to make probabilistic predictions that reflect the uncertainty in finite random markets.

Much work remains, including the establishment of rigorous accuracy guarantees. The predictions of the model can be shown to be asymptotically valid in the special cases considered by Arnosti (2022), and coincide with those of Azevedo and Leshno (2016) in the limiting regime that they consider (where school capacities grow large). However, numerical results indicate its surprising accuracy outside of these cases. Thus, I am left with a useful tool, but only an incomplete understanding of why it works. Instead of letting the desire for a complete understanding of this tool hold back progress, I take inspiration from theoretical physics, which often uses calculations and formal frameworks as tools to generate insights, even if those calculations have not been rigorously justified. Even before more complete justification is available, the model from this paper can be used to generate new insights about matching markets, as I demonstrate in Section 6.

REFERENCES

- Atila Abdulkadiroglu, Joshua D Angrist, Yusuke Narita, and Parag A Pathak. 2017. Research design meets market design: Using centralized assignment for impact evaluation. *Econometrica* 85, 5 (2017), 1373–1432.
- Atila Abdulkadiroglu, Yeon-Koo Che, and Yosuke Yasuda. 2015. Expanding ‘Choice’ in School Choice. *American Economic Journal: Microeconomics* 7, 1 (2015), 1–42.
- Atila Abdulkadiroglu, Parag A. Pathak, and Alvin E. Roth. 2009. Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the New York City High School Match. *American Economic Review* 99, 5 (2009), 1954–1978.
- Hiroyuki Adachi. 2000. On a characterization of stable matchings. *Economic Letters* 68 (2000), 43–49.
- Nikhil Agarwal. 2015. An Empirical Model of the Medical Match. *Amerian Economic Review* 105, 7 (2015), 1939–1978.
- Maxwell Allman, Itai Ashlagi, and Afshin Nikzad. 2022. On Rank Dominance of Tie-Breaking Rules. (2022).
- Nick Arnosti. 2022. Lottery Design for School Choice. *Management Science* (2022).
- Itai Ashlagi, Yashodhan Kanoria, and Jacob Leshno. 2017. Unbalanced Random Matching Markets: The Stark Effect of Competition. *Journal of Political Economy* 125, 1 (2017), 2 pages. <https://doi.org/10.1145/2482540.2482590>
- Itai Ashlagi and Afshin Nikzad. 2019. What Matters in Tie-Breaking Rules? How Competition Guides Design. (2019).
- Itai Ashlagi, Afshin Nikzad, and Assaf I Romm. 2019. Assigning More Students to Their Top Choices: A Tiebreaking Rule Comparison. *Games and Economic Behavior* 115 (May 2019), 167–187.
- Itai Ashlagi and Peng Shi. 2016. Optimal Allocation Without Money: An Engineering Approach. *Management Science* 62, 4 (2016), 1078–1097.
- Eduardo Azevedo and John Hatfield. 2018. Existence of Equilibrium in Large Matching Markets with Complementarities. (2018).
- Eduardo Azevedo and Jacob Leshno. 2016. A Supply and Demand Framework for Two-Sided Matching Markets. *Journal of Political Economy* 124, 5 (2016).
- Aaron Bodoh-Creed. 2020. Optimizing for Distributional Goals in School Choice Problems. *Management Science* 66, 8 (2020), 3657–3676.
- Aaron Bodoh-Creed and Brent Hickman. 2018. College assignment as a large contest. *Journal of Economic Theory* 175 (2018), 88–126.
- Yeon-Koo Che, Jinwoo Kim, and Fuhito Kojima. 2019. Stable Matching in Large Economies. *Econometrica* 87, 1 (2019), 65–110.
- Monique de Haan, Pieter Gautier, Hessel Oosterbeek, and Bas van der Klaauw. 2018. The Performance of School Assignment Mechanisms in Practice. (2018).

- Umut M. Dur, Scott Duke Kominers, Parag A Pathak, and Tayfun Sönmez. 2018. Reserve Design: Unintended Consequences and the Demise of Boston’s Walk Zones. *Journal of Political Economy* 126, 6 (December 2018).
- Federico Echenique. 2004. Core Many-to-one Matchings by Fixed-Point Methods. *Journal of Economic Theory* 115, 2 (2004), 358–376.
- Tamas Fleiner. 2003. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research* 28, 1 (2003), 103–126.
- D. Gale and L. S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), pp. 9–15. <http://www.jstor.org/stable/2312726>
- Michael Greinecker and Christopher Kah. 2021. Pairwise Stable Matching in large Economies. *Econometrica* 89, 6 (2021), 2929–2974.
- Nicole Immorlica and Mohammad Mahdian. 2005. Marriage, Honesty, and Stability. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (Vancouver, British Columbia) (*SODA ’05*). Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 53–62. <http://dl.acm.org/citation.cfm?id=1070432.1070441>
- Robert W. Irving. 1994. Stable marriage and indifference. *Discrete Applied Mathematics* 48 (1994), 261–272.
- Ravi Jagadeesan and Karolina Vocke. 2021. Stability in Large Markets. (2021).
- Yash Kanoria, Seungki Min, and Pengyu Qian. 2021. In which matching markets does the short side enjoy an advantage? *ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2021).
- Matt Kasman and Jon Valant. 2019. The opportunities and risks of K-12 student placement algorithms.
- Donald E. Knuth. 1976. *Mariages Stables et leurs relations avec d’autres problèmes combinatoires*. Les Presses de l’Université de Montréal.
- Donald E. Knuth. 1996. An exact analysis of stable allocation. *Journal of Algorithms* 20, 2 (1996), 431–442.
- Fuhito Kojima and Parag A Pathak. 2009. Incentives and stability in large two-sided matching markets. *The American Economic Review* (2009), 608–627.
- David F. Manlove. 1999. *Stable Marriage with Ties and Unacceptable Partners*. Technical Report 29. University of Glasgow Department of Computing Science.
- Philip Marx and James Schummer. 2021. Revenue from Matching Platforms. *Theoretical Economics* 16, 3 (2021), 799–824.
- D.G. McVitie and L.B. Wilson. 1970. Stable Marriage Assignment For Unequal Sets. *BIT Numerical Mathematics* 10, 3 (1970), 295–309.
- Marcin Peski. 2011. Smooth Stable Matching. (2011).

- B. Pittel. 1989. The Average Number of Stable Matchings. *SIAM Journal on Discrete Mathematics* 2, 4 (1989), 530–549. <https://doi.org/10.1137/0402048> arXiv:<http://dx.doi.org/10.1137/0402048>
- Alvin Roth. 1986. On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets. *Econometrica* 54, 2 (1986), 425–427.
- Alvin E. Roth and Elliott Peranson. 1999. *The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design*. Working Paper 6963. National Bureau of Economic Research. <https://doi.org/10.3386/w6963>
- Peng Shi. 2015. Guiding School-Choice Reform through Novel Applications of Operations Research Guiding School-Choice Reform through Novel Applications of Operations Research Guiding School-Choice Reform through Novel Applications of Operations Research. *Interfaces* 45, 2 (2015), 117–132.
- Peng Shi. 2022. Optimal Priority-Based Allocation Mechanisms. *Management Science* 68, 1 (2022), 171–188.
- L.B. Wilson. 1972. An Analysis of the Stable Marriage Assignment Algorithm. *BIT Numerical Mathematics* 12 (1972), 569–575.

APPENDIX A. PROOFS FROM SECTION 5

A.1. Proof of Proposition 1. The following Lemma states that for any $(\eta^S, \mathcal{V}^{det})$ -stable outcome, the enrollment at each school will be the minimum of the school’s capacity and the number of interested students.

Lemma 1. *If \mathcal{S} is a finite subset of Θ with strict priorities, and (M, I, A) is $(\eta^S, \mathcal{V}^{det})$ -stable, then for any $h \in \mathcal{H}$ and $p \in [0, 1]$,*

$$\sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > p) M_h(\theta) = \min(I_h(p), C_h).$$

Proof of Lemma 1. Define $\bar{p} = \inf\{p \in [0, 1] : I_h(p) < C_h\}$. Note that (4) and (12) imply that I_h is right-continuous, and therefore $I_h(\bar{p}) < C_h$. It follows from (5) and (6) that for $p \in [0, 1]$,

$$(30) \quad A_h(p) = \mathcal{V}^{det}(I_h(p), C_h) = \mathbf{1}(I_h(p) < C_h) = \mathbf{1}(p \geq \bar{p}).$$

Combining (8) and (30) we see that if $I_h(p) < C_h$ then $p \geq \bar{p}$ and

$$(31) \quad \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > p) M_h(\theta) = \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > p) (1 - \sum_{h' \succ^\theta h} M_{h'}(\theta)) = I_h(p),$$

where the final inequality uses (4) and (12).

Meanwhile, if $I_h(p) \geq C_h$, then $\bar{p} > p$, and (8) and (30) imply that

$$\begin{aligned} \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > p) M_h(\theta) &= \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > \bar{p}) (1 - \sum_{h' \succ^{\theta} h} M_{h'}(\theta)) + \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta = \bar{p}) (1 - \sum_{h' \succ^{\theta} h} M_{h'}(\theta)), \\ (32) \quad &= I_h(\bar{p}) + \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta = \bar{p}) (1 - \sum_{h' \succ^{\theta} h} M_{h'}(\theta)), \end{aligned}$$

where the second line also follows from (30).

Note that because $\mathcal{V}^{det}(\lambda, C) \in \{0, 1\}$, (5) and (7) imply that $M_h(\theta) \in \{0, 1\}$ and therefore (4) implies that $I_h(p) \in \mathbb{N}$. Furthermore, the fact that \mathcal{S} has strict priorities implies that at discontinuities of I_h , it decreases by exactly one. We know from the definition of \bar{p} that $I_h(\bar{p}) < C_h$ but $I_h(p) \geq C_h$ for all $p < \bar{p}$, so $I_h(\bar{p}) = C_h - 1$ and $\sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta = \bar{p}) (1 - \sum_{h' \succ^{\theta} h} M_{h'}(\theta)) = 1$. This implies that the expression in (32) is equal to C_h , completing the proof. \square

Proof of Proposition 1. We first suppose that (M, I, A) is $(\eta^{\mathcal{S}}, \mathcal{V}^{det})$ -stable, and show that μ^M has no blocking pairs. Note that the definition of \mathcal{V}^{det} in (6) implies that for all $\lambda \in \mathbb{R}_+$, $C \in \mathbb{N}$ we have $\mathcal{V}^{det}(\lambda, C) \in \{0, 1\}$, so by (5) and (7), M is deterministic and μ^M is well-defined. We now show that μ is feasible. Note that

$$\begin{aligned} \sum_{\theta \in \mathcal{S}} M_h(\theta) &= \sum_{\theta \in \mathcal{S}} M_h(\theta) \mathbf{1}(p_h^\theta > 0) + \sum_{\theta \in \mathcal{S}} M_h(\theta) \mathbf{1}(p_h^\theta = 0) \\ &\leq \sum_{\theta \in \mathcal{S}} M_h^\theta \mathbf{1}(p_h^\theta > 0) + \eta^{\mathcal{S}}(\Theta_h(0)) A_h(I_h(0)) \\ &= \min(I_h(0), C_h) + \eta^{\mathcal{S}}(\Theta_h(0)) \mathbf{1}(I_h(0) < C_h) \\ &\leq \min(I_h(0), C_h) + \mathbf{1}(I_h(0) < C_h) \\ &\leq C_h. \end{aligned}$$

The second line follows from $M_h(\theta) \leq A_h(p_h^\theta)$ (see (7)) and $\sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta = 0) = \eta^{\mathcal{S}}(\Theta_h(0))$ (see (12)); the third from Lemma 1, along with (5) and (6), and the fourth because \mathcal{S} has strict priorities.

Finally, we show that μ^M has no blocking pairs. By definition, if $h \succ^{\theta'} \mu^M(\theta')$ then $M_h(\theta') = 0$, and $A_h(p_h^{\theta'}) = 0$ by (7). From this, (5) and (6) imply that $I_h(p_h^{\theta'}) \geq C_h$, so by Lemma 1,

$$|\{\theta \in \mathcal{S} : \mu^M(\theta) = h, p_h^\theta > p_h^{\theta'}\}| = \sum_{\theta \in \mathcal{S}} M_h(\theta) \mathbf{1}(p_h^\theta > p_h^{\theta'}) = \min(I_h(p_h^{\theta'}), C_h) = C_h.$$

That is, (11) holds, so μ^M has no blocking pairs.

Next, we assume that μ is an \mathcal{S} -matching with no blocking pairs, and show that

i) M^μ “agrees” with μ : for $\theta \in \mathcal{S}, h \in \mathcal{H}_0$, we have

$$(33) \quad M_h^\mu(\theta) = \mathbf{1}(\mu(\theta) = h).$$

ii) M^μ is $(\eta^{\mathcal{S}}, \mathcal{V}^{det})$ -stable.

We start by showing (33). Fix $\theta' \in \mathcal{S}$. Then for any $h \succ^{\theta'} \mu(\theta')$, the fact that μ has no blocking pairs implies that (11) holds, from which the definition of A^μ in (14) implies that $A_h^\mu(p_h^{\theta'}) = 0$, so $M_h^\mu(\theta') = 0$. Meanwhile, for $h' = \mu(\theta')$, feasibility of μ implies

$$|\{\theta \in \mathcal{S} : \mu(\theta) = h', p_{h'}^\theta > p_{h'}^{\theta'}\}| < |\{\theta \in \mathcal{S} : \mu(\theta) = h'\}| \leq C_{h'}.$$

Therefore, the definition of A^μ in (14) implies that $A_{h'}^\mu(p_{h'}^{\theta'}) = 1$, from which (7) implies that $M_{h'}(\theta') = 1$ (and that $M_h(\theta) = 0$ for all h such that $\mu(\theta') \succ^{\theta'} h$). Thus, (33) holds, implying that $\mu = \mu^{M^\mu}$.

We now show that M^μ is $(\eta^\mathcal{S}, \mathcal{V}^{det})$ -stable. Define $I^\mu = \mathcal{I}(M^\mu)$. Then we have

$$(34) \quad I_h^\mu(p) = \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > p) (1 - \sum_{h' \succ^{\theta} h} M_h^\mu(\theta))$$

$$(35) \quad \geq \sum_{\theta \in \mathcal{S}} \mathbf{1}(p_h^\theta > p) \mathbf{1}(\mu(\theta) = h) = |\{\theta \in \mathcal{S} : p_h^\theta > p, \mu(\theta) = h\}|,$$

where the first equality follows from (4) and the definition of $\eta^\mathcal{S}$ in (12), and the inequality from (33). We claim that for $h \in \mathcal{H}_0$, $p \in [0, 1]$,

$$(36) \quad A_h^\mu(p) = \mathbf{1}(I_h^\mu(p) < C_h) = \mathcal{V}^{det}(I_h^\mu(p), C_h),$$

implying that $A^\mu = \mathcal{A}(I^\mu) = \mathcal{A}(\mathcal{I}(\mathcal{M}(A^\mu)))$, so A^μ is $(\eta^\mathcal{S}, \mathcal{V}^{det})$ -stable, and therefore so is M^μ . To establish (36), we show that $A_h^\mu(p) = 0$ implies $I_h^\mu(p) \geq C_h$, and $A_h^\mu(p) = 1$ implies $I_h^\mu(p) < C_h$.

If $A_h^\mu(p) = 0$, then by definition of A^μ in (14),

$$|\{\theta \in \mathcal{S} : p_h^\theta > p, \mu(\theta) = h\}| \geq C_h.$$

By (35), this implies that $I_h^\mu(p) \geq C_h$. Conversely, if $A_h^\mu(p) = 1$, then by definition

$$|\{\theta \in \mathcal{S} : p_h^\theta > p, \mu(\theta) = h\}| < C_h.$$

This implies that for each $\theta \in \mathcal{S}$ that contributes to the sum in (34), $\mu(\theta) = h$ (otherwise, (11) would be violated). Therefore, the inequality in (35) is an equality, implying that $I_h^\mu(p) < C_h$. \square

A.2. Proof of Proposition 2. We start by establishing an analog to Lemma 1, which says that for any $(\mathcal{E}, \mathcal{V})$ -stable outcome, the measure of students matched to school h can be determined by the measure of interest in h .

Lemma 2. *If η has strict priorities and \mathcal{V} is weakly decreasing in its first argument, then for any $(\mathcal{E}, \mathcal{V})$ -stable outcome (M, I, A) , any school $h \in \mathcal{H}$, and any $p \in [0, 1]$,*

$$\int \mathbf{1}(p_h^\theta > p) M_h(\theta) d\eta(\theta) = \int_0^{I_h(p)} \mathcal{V}(\lambda, C_h) d\lambda.$$

If $\mathcal{V} = \mathcal{V}^{det}$, then the expression on the right is $\min(I_h(p), C_h)$, matching that in Lemma 1. However, Lemma 2 provides a more general expression that holds for any monotone vacancy function.¹¹

Proof of Lemma 2. Fix $n \in \mathbb{N}$, define $m = \lceil nI_h(p) \rceil$, and define $\{L_i\}_{i=0}^m$ by $L_i = i/n$ for $i < m$, and $L_m = I_h(p)$. Note that if η is a continuous measure with strict priorities, then (4) implies that I is continuous and decreasing. In particular, this implies that we can choose $1 = P_0 > P_1 > \dots > P_m = p$ such that $I_h(P_i) = L_i$ for each i . We claim the following chain of inequalities:

$$\begin{aligned}
 \int \mathbf{1}(p_h^\theta > p) M_h(\theta) d\eta(\theta) &= \int \mathbf{1}(p_h^\theta > p) A_h(p_h^\theta) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta) \\
 &= \int \mathbf{1}(p_h^\theta > p) \mathcal{V}(I_h(p_h^\theta), C_h) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta) \\
 &= \int \sum_{i=1}^m \mathbf{1}(P_{i-1} \geq p_h^\theta > P_i) \mathcal{V}(I_h(p_h^\theta), C_h) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta) \\
 (37) \quad &\geq \sum_{i=1}^m \mathcal{V}(L_i, C_h) \int \mathbf{1}(P_{i-1} \geq p_h^\theta > P_i) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta).
 \end{aligned}$$

The first equality holds from (7), the second from (5), and the third by definition of P_i . The final inequality comes from the fact that \mathcal{V} is weakly decreasing in its first argument and I_h is weakly decreasing, and thus $p_h^\theta > P_i$ implies $\mathcal{V}(I_h(p_h^\theta), C_h) \geq \mathcal{V}(I_h(P_i), C_h) = \mathcal{V}(L_i, C_h)$. Note that

$$\begin{aligned}
 &\int \mathbf{1}(P_{i-1} \geq p_h^\theta > P_i) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta) \\
 &= \int \mathbf{1}(p_h^\theta > P_i) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta) - \int \mathbf{1}(p_h^\theta > P_{i-1}) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta(\theta) \\
 &= I_h(P_i) - I_h(P_{i-1}) \\
 (38) \quad &= L_i - L_{i-1}.
 \end{aligned}$$

The second equality follows from (4) and the third from the choice of P_i . Combining (37) and (38), and noting that $L_i - L_{i-1} = 1/n$ for $i < m$, we get

$$\int \mathbf{1}(p_h^\theta > p) M_h(\theta) d\eta(\theta) \geq \frac{1}{n} \sum_{i=1}^{m-1} \mathcal{V}(L_i, C_h).$$

¹¹The result does not hold when η is a discrete measure. To see this, consider a simple case with one school with capacity one, and two students, with priority $p_H > p_L$. Regardless of what matching we start with, both students are interested in the school, so the resulting interest function is $I(p) = \mathbf{1}(p \leq p_H) + \mathbf{1}(p \leq p_L)$ (a step function that steps down at p_L and p_H). For this interest function, if we use \mathcal{V}^{pois} , the resulting admissions function is $A(p) = 1$ for $p \in [p_H, 1]$, $A(p) = 1/e$ for $p \in [p_L, p_H]$ and $A(p) = 1/e^2$ for $p \in [0, p_L]$. Therefore, the top student gets in for sure, and the lower student with probability $1/e$. The total number of matches predicted by the stable matching M is $1 + 1/e$; this is more than $\int_0^\infty \mathcal{V}^{pois}(\lambda, 1) d\lambda = 1$.

This holds for any $n \in \mathbb{N}$. Taking the limit as $n \rightarrow \infty$ yields

$$(39) \quad \int \mathbf{1}(p_h^\theta > p) M_h(\theta) d\eta(\theta) \geq \int_0^{I_h(p)} \mathcal{V}(\lambda, C_h) d\lambda.$$

The inequality in (37) can be reversed if we replace $\mathcal{V}(L_i, C_h)$ with $\mathcal{V}(L_{i-1}, C_h)$. From there, an analogous argument implies (39) with the inequality reversed, completing the proof. \square

Proof of Proposition 2. Given cutoffs $P \in [0, 1]^{\mathcal{H}}$, we define

$$(40) \quad A_h^P(p) = \mathbf{1}(p > P_h).$$

$$(41) \quad M^P = \mathcal{M}(A^P).$$

We begin by noting two equalities that will repeatedly prove useful. Note that the definition of \mathcal{A} through (5) and (6), and the definition of A^P and $\mathcal{P}(I)$ in (40) and (17) imply that for any interest function $I \in \mathfrak{I}$,

$$(42) \quad \mathcal{A}(I) = A^{\mathcal{P}(I)}.$$

Furthermore, the definition of I^P , A^P and M^P in (16), (40) and (41), along with (4) and (7), imply that for any cutoff vector P ,

$$(43) \quad I^P = \mathcal{I}(\mathcal{M}(A^P)).$$

We first assume that I is $(\mathcal{E}, \mathcal{V}^{det})$ -stable, and show that $I = I^{\mathcal{P}(I)}$. By Definition 1, if I is $(\mathcal{E}, \mathcal{V}^{det})$ -stable, then so is

$$(44) \quad \mathcal{M}(\mathcal{A}(I)) = \mathcal{M}(A^{\mathcal{P}(I)}) = M^{\mathcal{P}(I)},$$

where the equalities follow from (42) and (41), respectively. Furthermore, stability of I implies that

$$(45) \quad I = \mathcal{I}(\mathcal{M}(\mathcal{A}(I))) = \mathcal{I}(\mathcal{M}(A^{\mathcal{P}(I)})) = I^{\mathcal{P}(I)},$$

where the second and third equalities follow from (44) and (43), respectively.

Next, we show that $\mathcal{P}(I)$ is market-clearing. We claim that

$$(46) \quad D_h(\mathcal{P}(I)) = \int M_h^{\mathcal{P}(I)}(\theta) d\eta(\theta) = \int_0^{I_h(0)} \mathcal{V}^{det}(\lambda, C_h) d\lambda = \min(I_h(0), C_h) \leq C_h.$$

The first equality follows from (15), the second from Lemma 2 and the fact that $\mathcal{I}(M^{\mathcal{P}(I)}) = I$ (by (44) and (45)), and the third from the definition of \mathcal{V}^{det} in (6). Furthermore, if $\mathcal{P}_h(I) > 0$, it follows from definition of \mathcal{P} in (17) that $I_h(P_h) \geq C_h$ (this also uses the fact that I_h is continuous, which follows from (4) and the fact that η is a continuous measure with strict priorities). Because (4) implies that I_h is weakly decreasing, it follows that $I_h(0) \geq C_h$, and therefore the inequality in (46) is tight. Therefore, $\mathcal{P}(I)$ is market-clearing.

Finally, we show that if $P \in [0, 1]^{\mathcal{H}}$ is market-clearing, then I^P is $(\mathcal{E}, \mathcal{V}^{det})$ -stable. By (42) and (43),

$$(47) \quad \mathcal{I}(\mathcal{M}(\mathcal{A}(I^P))) = \mathcal{I}(\mathcal{M}(A^{\mathcal{P}(I^P)})) = I^{\mathcal{P}(I^P)}.$$

We wish to show that this is equal to I^P . For less cumbersome notation, we define the cutoff vector

$$(48) \quad \tilde{P} = \mathcal{P}(I^P).$$

The steps to prove that $I^{\tilde{P}} = I^P$ are as follows:

I. Show that $I_h^P(P_h) = D_h(P)$. Conclude that

$$(49) \quad \tilde{P}_h \geq P_h.$$

II. Define

$$(50) \quad \Delta_h = \{\theta : p_h^\theta \in (P_h, \tilde{P}_h], p_{h'}^\theta \leq P_{h'} \text{ for all } h' \succ^\theta h\}$$

to be the η -measure of students who have priority between P_h and \tilde{P}_h at h , and are not admitted to any school preferred to h under either P or \tilde{P} . Establish that I_h^P is constant on $(P_h, \tilde{P}_h]$, and therefore that

$$(51) \quad 0 = I_h^P(P_h) - I_h^P(\tilde{P}_h) = \eta(\Delta_h).$$

III. Conclude that for all $h \in \mathcal{H}$ and $p \in [0, 1]$,

$$(52) \quad I_h^{\tilde{P}}(p) = I_h^P(p).$$

We now establish step I. Note that

$$\begin{aligned} I_h^P(p) &= \int \mathbf{1}(p_h^\theta > p) \left(1 - \sum_{h' \succ^\theta h} M_{h'}^P(\theta)\right) d\eta(\theta) \\ &= \int \mathbf{1}(p_h^\theta > p) \prod_{h' \succ^\theta h} (1 - A_{h'}^P(p_{h'}^\theta)) d\eta(\theta), \end{aligned}$$

where the first line follows from (43) and (4), and the second from (8). It follows that

$$\begin{aligned} I_h^P(P_h) &= \int A_h^P(p_h^\theta) \prod_{h' \succ^\theta h} (1 - A_{h'}^P(p_{h'}^\theta)) d\eta(\theta) \\ (53) \quad &= \int M_h^P d\eta(\theta) = D_h(P), \end{aligned}$$

where the second line follows from (7) and the definitions of A^P and $D_h(P)$ in (40) and (15). From (53) and the definition of $\mathcal{P}(\cdot)$ in (17), (49) follows.

Next, we move to step II. Because η is a continuous measure with strict priorities, (4) implies that I_h^P is continuous for each $h \in \mathcal{H}_0$. Therefore, the definition of $\mathcal{P}(\cdot)$ in (17) implies that either $\tilde{P}_h = 0$ (in which case (49) implies that $P_h = \tilde{P}_h$), or $I_h^P(\tilde{P}) = C_h$. But

then (49) and the fact that I^P is decreasing imply that

$$I_h^P(P_h) \geq I_h^P(\tilde{P}_h) = C_h.$$

Because P is market-clearing, $D_h(P) \leq C_h$, implying that the inequality above must hold with equality. Therefore, I_h^P is constant on $(P_h, \tilde{P}_h]$. In particular, applying the definition of I^P in (16) reveals that (51) holds.

Finally, we move to step III. By (16), for any $h \in \mathcal{H}$ and $p \in [0, 1]$ we have

$$(54) \quad I_h^{\tilde{P}}(p) - I_h^P(p) = \eta(\Delta),$$

where

$$\Delta = \{\theta : p_h^\theta > p, \prod_{h' \succ^\theta h} \mathbf{1}(p_{h'}^\theta \leq \tilde{P}_{h'}) - \prod_{h' \succ^\theta h} \mathbf{1}(p_{h'}^\theta \leq P_{h'}) = 1\}$$

That is, the difference $I_h^{\tilde{P}}(p) - I_h^P(p)$ is the measure of students who have priority above p at h , and are admitted to a school preferred to h under cutoffs \tilde{P} , but are not admitted to any such school under P . For any $\theta \in \Delta$, there is some most-preferred school h' where θ is admitted under P but not under \tilde{P} . Then (50) implies that $\theta \in \Delta_{h'}$. In other words, $\Delta \subseteq \bigcup_{h \in \mathcal{H}} \Delta_h$. From this, (51) implies that $\eta(\Delta) = 0$, and (54) implies that (52) holds. This establishes that $I^P = I^{\tilde{P}} = I^{\mathcal{P}(I^P)}$, from which (47) implies that I^P is stable. \square

A.3. Proof of Theorems 2 and 3.

Proof of Theorem 2. By Theorem 1, there exist maximal and minimal stable outcomes, corresponding to the school-optimal and student-optimal stable outcomes, respectively. Denote these outcomes by $(M^H, I^H, A^H) \succeq (M^L, I^L, A^L)$, respectively. It is enough to prove the result for these outcomes. Note that

$$(55) \quad \begin{aligned} \int \sum_{h \succ^\theta \emptyset} M_h^L(\theta) d\eta(\theta) &= \sum_{h \in \mathcal{H}} \int M_h^L(\theta) d\eta(\theta) = \sum_{h \in \mathcal{H}} \int_0^{I_h^L(0)} \nu(\lambda, C_h) d\lambda \\ &\geq \sum_{h \in \mathcal{H}} \int_0^{I_h^H(0)} \nu(\lambda, C_h) = \sum_{h \in \mathcal{H}} \int M_h^H(\theta) d\eta(\theta) = \int \sum_{h \succ^\theta \emptyset} M_h^H(\theta) d\eta(\theta). \end{aligned}$$

The first and last equalities hold because $A_\emptyset(p) = 1$ for all p , so by (7), $M_h(\theta) = 0$ if $\emptyset \succ^\theta h$. The second and second-to-last equalities hold by Lemma 2. The inequality follows from the fact that $I^L \succeq^J I^H$. But $M^H \succeq^M M^L$ implies

$$(56) \quad \sum_{h \succ^\theta \emptyset} M_h^L(\theta) \leq \sum_{h \succ^\theta \emptyset} M_h^H(\theta) \quad \forall \theta \in \Theta.$$

Therefore, the inequality in (55) must hold with equality. In particular, this implies that (19) holds for each $h \in \mathcal{H}$. Furthermore, this implies that (20) holds for all θ except possibly a set of η -measure zero. \square

Proof of Theorem 3. It suffices to show that there is a unique stable interest function: that is, if I^H and I^L are the largest and smallest stable interest functions according to \succeq^J , then $I^H = I^L$. We let A^H, M^H be the admissions function and matching associated with I^H , and define A^L, M^L analogously. We note that by (5) and the fact that \mathcal{V} is decreasing in its first argument, $I^H \succeq^J I^L$ implies that

$$(57) \quad A^L \succeq^{\mathfrak{A}} A^H.$$

The proof proceeds by contradiction, showing that $I^H \succ^I I^L$ implies that Theorem 2 does not hold. That is, $I^H \succ^I I^L$ implies the existence of a set $\tilde{\Theta}$ with $\eta(\tilde{\Theta}) > 0$ such that

$$(58) \quad \sum_{h \in \mathcal{H}} M_h^H(\theta) < \sum_{h \in \mathcal{H}} M_h^L(\theta) \text{ for all } \theta \in \tilde{\Theta}.$$

We establish existence of such a $\tilde{\Theta}$ in three steps.

- I. Note that Definition 2 and (4) imply that
 - a) the stable interest functions I^H and I^L are component-wise continuous, and
 - b) for all $h \in \mathcal{H}$, $I_h^H(1) = I_h^L(1) = 0$.
- II. These jointly imply that if $I_h^H(p) > I_h^L(p)$ for some $p \in [0, 1]$ and $h \in \mathcal{H}$, then there must exist an interval (\underline{p}, \bar{p}) such that:
 - a) $I_h^H(p) > I_h^L(p)$ for $p \in (\underline{p}, \bar{p})$, and
 - b) $I_h^H(\bar{p}) > I_h^H(p)$.

That is, I_h^H is not constant and strictly larger than I_h^L on this interval.

III. Define

$$(59) \quad S = \{\theta : M_\emptyset^H(\theta) = 0\}$$

to be the set of student types who are sure to be matched. Define

$$(60) \quad \tilde{\Theta} = \{\theta : h \succ^\theta \emptyset, p_h^\theta \in (\underline{p}, \bar{p})\} \setminus S.$$

We will show that (58) holds, and that $\eta(\tilde{\Theta}) > 0$.

To see that (58) holds, note that if $\theta \in \tilde{\Theta}$,

$$(61) \quad A_h^H(p_h^\theta) = \mathcal{V}(I^H(p_h^\theta), C_h) < \mathcal{V}(I^L(p_h^\theta), C_h) = A_h^L(p_h^\theta),$$

where the equalities hold by (5) and the inequality follows from II.a), the fact that $p_h^\theta \in (\underline{p}, \bar{p})$, and the fact that $\mathcal{V}(\cdot, C_h)$ is strictly decreasing. Thus, when comparing A^H to A^L , each student in $\tilde{\Theta}$ is

- i. weakly less likely to be admitted to each school under A^H by (57),
- ii. strictly less likely to be admitted to h under A^H by (61), and
- iii. not certain to be admitted to any school by definition of $\tilde{\Theta}$ in (60).

From this, (7) implies that each $\theta \in \tilde{\Theta}$ is strictly less likely to match under A^H . That is, (58) holds.

Finally, we show that $\eta(\tilde{\Theta}) > 0$. By definition of $\tilde{\Theta}$ in (60)

$$\begin{aligned}
 (62) \quad \eta(\tilde{\Theta} \cup S) &\geq \int \mathbf{1}(h \succ^\theta \emptyset, p_h^\theta \in (\underline{p}, \bar{p})) d\eta(\theta) \\
 &\geq \int \mathbf{1}(h \succ^\theta \emptyset, p_h^\theta \in (\underline{p}, \bar{p})) (1 - \sum_{h' \succ^\theta h} M_{h'}^H(\theta)) d\eta(\theta) \\
 &= I_h^H(\underline{p}) - I_h^H(\bar{p}) \\
 &> 0,
 \end{aligned}$$

where the equality follows from (4) and stability of I^H , and the final line follows from II.b). We complete the proof by showing that $\eta(S) = 0$.

Because $\mathcal{V}(\cdot, C_h)$ is strictly decreasing, (5) implies that $A_h^H(p) = 1$ if and only if $I_h^H(p) = 0$. Define $p_h = \inf\{p : I_h^H(p_h^\theta) = 0\}$ to be the lowest priority at school h that guarantees admission, and note that

$$(63) \quad 0 = I_h^H(p_h) = \int \mathbf{1}(p_h^\theta > p_h) (1 - \sum_{h' \succ^\theta h} M_{h'}^H(\theta)) d\eta(\theta),$$

where the second equality comes from stability of I^H . Define

$$S_h = \{\theta : p_h^\theta > p_h, \sum_{h' \succ^\theta h} M_{h'}^H(\theta) < 1\}$$

to be the set of agents who are certain to be admitted to h , and not certain to be admitted to any option that they prefer to h . Note that

$$\eta(S_h) = \int \mathbf{1}(p_h^\theta > p_h) \mathbf{1}\left(\sum_{h' \succ^\theta h} M_{h'}^H(\theta) < 1\right) d\eta(\theta) = 0,$$

where the second equality follows from (63). Because $S = \bigcup_{h \in \mathcal{H}} S_h$, it follows that $\eta(S) = 0$ and thus $\eta(\tilde{\Theta}) > 0$ by (62). \square

APPENDIX B. PROOFS FROM SECTION 6

Define

$$(64) \quad \text{AcceptanceRate}(\lambda, C) = \text{Enrollment}(\lambda, C) / \lambda = \frac{1}{\lambda} \int_0^\lambda \mathcal{V}(x, C) dx,$$

where the second equality follows from the definition of *Enrollment* in (25). The following result implies that *AcceptanceRate* is decreasing in its first argument.

Lemma 3. *Given $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, define $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by*

$$g(y) = \frac{1}{y} \int_0^y f(x) dx.$$

If f is weakly increasing, then so is g . If f is weakly decreasing, then so is g .

Proof of Lemma 3. Note that

$$g'(y) = \frac{yf(y) - \int_0^y f(x)dx}{y^2}.$$

If f is weakly increasing, then $yf(y) \geq \int_0^y f(x)dx$; if f is weakly decreasing, the inequality reverses. \square

Proof of Proposition 3. Let $\ell(\theta) = |\{h : h \succ^\theta \emptyset\}|$ be the number of schools listed by type θ . We note that for any individually rational matching M and any θ ,

$$\begin{aligned} \sum_{h' \succ^\theta \emptyset} M_{h'}(\theta) R_{h'}(\theta) &\leq \ell(\theta)(1 - \sum_{h' \succ^\theta \emptyset} M_{h'}(\theta)) + \sum_{h' \succ^\theta \emptyset} M_{h'}(\theta) R_{h'}(\theta) \\ &= \ell(\theta) - \sum_{h' \succ^\theta \emptyset} M_{h'}(\theta)(\ell(\theta) - R_{h'}(\theta)) \\ &= \ell(\theta) - \sum_{h' \succ^\theta \emptyset} \sum_{h' \succ^\theta h \succ^\theta \emptyset} M_{h'}(\theta) \\ (65) \quad &= \sum_{h \succ^\theta \emptyset} (1 - \sum_{h' \succ^\theta h} M_{h'}(\theta)) \end{aligned}$$

$$(66) \quad = \sum_{h \in \mathcal{H}} (1 - \sum_{h' \succ^\theta h} M_{h'}(\theta))$$

where the third line follows from the fact that $\ell(\theta) - R_{h'}(\theta)$ is the number of acceptable schools that rank below h' , the fourth follows by exchanging the order of summation, and the last uses the fact that M is individually rational.

From (66) and the definition of *AverageRank* in (23), it follows that if (M, I, A) is an $(\mathcal{E}, \mathcal{V}^{pois})$ stable outcome,

$$(67) \quad \text{AverageRank}(M) \leq \frac{\int \sum_{h \in \mathcal{H}} M_h(\theta)(1 - \sum_{h' \succ^\theta h} M_{h'}(\theta)) d\eta(\theta)}{\int \sum_{h \in \mathcal{H}} M_h(\theta) d\eta(\theta)} = \frac{\sum_{h \in \mathcal{H}} I_h(0)}{\sum_{h \in \mathcal{H}} \int_0^{I_h(0)} \mathcal{V}^{pois}(\lambda, C_h) d\lambda}.$$

Note that the final equality follows by the fact that M is stable, (4) and Lemma 2. In a symmetric iid market, $C_h = C_{h'}$ and $I_h = I_{h'}$ for all $h, h' \in \mathcal{H}$, so implies that for any $h \in \mathcal{H}$,

$$(68) \quad \text{AverageRank}(M) \leq \frac{I_h(0)}{\int_0^{I_h(0)} \mathcal{V}(\lambda, C_h) d\lambda} = \frac{1}{\text{AcceptanceRate}(I_h(0), C_h)}.$$

Lemma 3 implies that *AcceptanceRate* is decreasing in its first argument, so we can obtain an upper bound on this expression by obtaining an upper bound on $I_h(0)$. But it is clear that the denominator in (67) is at most $\eta(\Theta)$, from which symmetry implies that $\text{Enrollment}(I_h(0), C_h) \leq \rho = \text{Enrollment}(\Lambda(\rho, C_h), C_h)$. Because *Enrollment* is increasing in its first argument, this implies that $I_h(0) \leq \Lambda(\rho, C)$. Plugging this into (68) completes the proof. \square

Lemma 4. *The function $AR : (0, 1] \rightarrow \mathbb{R}_+$ defined by¹²*

$$(69) \quad AR(q) = \frac{1}{q} - \frac{\ell(1-q)^\ell}{1-(1-q)^\ell},$$

is decreasing in q .

Proof of Proposition 4. Let (M, I, A) be the unique $(\eta^{IID}, \mathcal{V}^{pois})$ -stable outcome. Note that by symmetry, $I_h(p) = I_{h'}(p)$ and $A_h(p) = A_{h'}(p)$ for all $h, h' \in \mathcal{H}$ and $p \in [0, 1]$, so in what follows, we write $I(p)$ and $A(p)$ in place of $I_h(p)$ and $A_h(p)$.

Define

$$(70) \quad q = \int_0^1 \mathcal{V}^{pois}(I(p), C) dp.$$

Note that

$$(71) \quad \begin{aligned} \int \sum_{h \in \mathcal{H}} M_h(\theta) d\eta(\theta) &= \int (1 - \prod_{h \succ^\theta \emptyset} (1 - A(p_h^\theta))) d\eta^{IID}(\theta) \\ &= \int (1 - \prod_{h \succ^\theta \emptyset} (1 - \mathcal{V}^{pois}(I(p_h^\theta), C))) d\eta^{IID}(\theta) \\ &= \eta^{IID}(\Theta)(1 - (1 - q)^\ell), \end{aligned}$$

where the first equality follows from (7), the second from (5), and the last from the fact that we assume all students list ℓ schools, and in an iid market, the priorities p_h are drawn iid $U[0, 1]$.

Furthermore, we have

$$(72) \quad \begin{aligned} \int \sum_{h \in \mathcal{H}} M_h(\theta) R_h(\theta) d\eta(\theta) &= \int \sum_{h \in \mathcal{H}} R_h(\theta) A_h(p_h^\theta) \prod_{h' \succ^\theta h} (1 - A_{h'}(p_{h'}^\theta)) d\eta^{IID}(\theta) \\ &= \eta^{IID}(\Theta) \sum_{k=1}^{\ell} k q (1 - q)^{k-1} \\ &= \eta^{IID}(\Theta) (1 - (1 - q)^\ell) AR(q). \end{aligned}$$

Jointly, (71) and (72) imply that

$$(73) \quad \text{AverageRank}(M) = AR(q).$$

Note that (26) implies that

$$(74) \quad \text{Enrollment}(\lambda, C) \leq C \text{ for all } \lambda \in \mathbb{R}_+,$$

¹²For a general list length distribution, the appropriate definition is

$$AR(\alpha) = \frac{\mu(\alpha) - \mathbb{E}[\ell(1-\alpha)^\ell]}{\mathbb{E}[1 - (1-\alpha)^\ell]} = 1 - \frac{(1-\alpha)\mu'(\alpha)}{\mu(\alpha)}.$$

and therefore

$$(75) \quad \eta^{IID}(\Theta)(1 - (1 - q)^\ell) = |\mathcal{H}| \text{Enrollment}(I(0), C) \leq |\mathcal{H}|C.$$

Define α and α' as the solutions to

$$(76) \quad \eta^{IID}(\Theta)(1 - (1 - \alpha)^\ell) = |\mathcal{H}|C = \eta^{IID}(\Theta)(1 - e^{-\alpha'\ell}).$$

Then it follows that

$$(77) \quad q \leq \alpha \leq \alpha',$$

with the last inequality holding because $e^{-\alpha'\ell} \geq (1 - \alpha')^\ell$. Because AR is decreasing by Lemma 4, equations (73) and (77) imply that

$$(78) \quad \text{AverageRank}(M) = AR(q) \geq AR(\alpha) = 1/\alpha - \ell(\rho/C - 1) \geq 1/\alpha' + \ell(1 - \rho/C),$$

where the second equality follows from the definition of α . Finally, noting that $\alpha' = \frac{-\log(1-C/\rho)}{\ell}$ completes the proof.

We now turn to the case where priorities are identical across schools. We define

$$(79) \quad q(u) = \mathcal{V}^{pois}(\Lambda(u, C), C).$$

We will prove that the following chain of inequalities hold:

$$(80) \quad \text{AverageRank}(M) = \frac{1}{\text{Enrollment}(I(0), C)} \int_0^{\text{Enrollment}(I(0), C)} AR(q(u)) du$$

$$(81) \quad \leq \frac{1}{C} \int_0^C AR(q(u)) du$$

$$(82) \quad \leq \int_0^1 AR(q) dq$$

$$(83) \quad \leq 1 + \log(\ell).$$

To evaluate $\text{AverageRank}(M)$, we note that

$$(84) \quad \int \sum_{h \in \mathcal{H}} M_h(\theta) R_h(\theta) d\eta^{RSD}(\theta) = |\mathcal{H}| \text{Enrollment}(I(0), C).$$

$$\int \sum_{h \in \mathcal{H}} M_h(\theta) R_h(\theta) d\eta^{RSD}(\theta) = \eta^{RSD}(\Theta) \int_0^1 (1 - (1 - \mathcal{V}^{pois}(I(p), C))^\ell) AR(\mathcal{V}^{pois}(I(p), C)) dp.$$

We apply u -substitution to the latter integral, with $u = \text{Enrollment}(I(p), C)$, so that

$$\frac{du}{dp} = \mathcal{V}^{pois}(I(p), C) I'(p) = -\frac{\eta^{RSD}(\Theta)}{|\mathcal{H}|} (1 - (1 - \mathcal{V}^{pois}(I(p), C))^\ell).$$

This yields

$$\eta^{RSD}(\Theta) \int_0^1 (1 - (1 - \mathcal{V}^{pois}(I(p), C))^\ell) AR(\mathcal{V}^{pois}(I(p), C)) dp = |\mathcal{H}| \int_0^{\text{Enrollment}(I(0), C)} AR(q(u)) du,$$

which when combined with (84) yields (80).

We now establish (81). The function q given in (79) is decreasing, as is AR by Lemma 4. Therefore, (74) and Lemma 3 imply (81).

We move on to establishing (82). Define $f : [0, 1] \rightarrow \mathbb{R}_+$ implicitly by

$$(85) \quad \mathcal{V}^{pois}(f(q), C) = q.$$

We note that

$$(86) \quad -\frac{d}{d\lambda} \mathcal{V}^{pois}(\lambda, C) = \mathcal{V}^{pois}(\lambda, C) - \mathcal{V}^{pois}(\lambda, C-1) \leq \mathcal{V}^{pois}(\lambda, C).$$

Therefore, by (79) and (85), we have

$$(87) \quad \frac{1}{C} \int_0^C AR(q(u)) du = \frac{1}{C} \int_0^1 AR(q) \frac{\mathcal{V}^{pois}(f(q), C)}{\mathcal{V}^{pois}(f(q), C) - \mathcal{V}^{pois}(f(q), C-1)} dq.$$

(88)

The function

$$h(\lambda) = \frac{1}{C} \frac{\mathcal{V}^{pois}(\lambda, C)}{\mathcal{V}^{pois}(\lambda, C) - \mathcal{V}^{pois}(\lambda, C-1)}$$

is decreasing in λ , from which it follows that $h(f(q))$ is increasing in q . Meanwhile, AR is decreasing in q by Lemma 4. It follows that

$$(89) \quad \int_0^1 AR(q) h(f(q)) dq \leq \int_0^1 AR(q) dq \int_0^1 h(f(q)) dq = \int_0^1 AR(q) dq.$$

The final inequality follows because by (85) and (26) we have

$$\int_0^1 h(f(q)) dq = \frac{1}{C} \int_0^\infty \mathcal{V}^{pois}(\lambda) d\lambda = 1.$$

Finally, we show (83). We note that by u -substitution with $1 - u = (1 - q)^\ell$,

$$\int_\varepsilon^1 \frac{\ell(1-q)^\ell}{1 - (1-q)^\ell} dq = \int_{1-(1-\varepsilon)^\ell}^1 \frac{(1-u)^{1/\ell}}{u} du.$$

Thus, we can write

$$\begin{aligned} \int_\varepsilon^1 AR(q) dq &= \int_\varepsilon^1 \frac{1}{q} - \frac{\ell(1-q)^\ell}{1 - (1-q)^\ell} dq = \int_\varepsilon^{1-(1-\varepsilon)^\ell} \frac{1}{q} dq + \int_{1-(1-\varepsilon)^\ell}^1 \frac{1 - (1-u)^{1/\ell}}{u} du \\ &\leq \log \left(\frac{1 - (1-\varepsilon)^\ell}{\varepsilon} \right) + 1, \end{aligned}$$

where the second line follows by evaluating the first integral and bounding the second using the fact that $(1 - (1 - u)^{1/\ell})/u \leq (1 - (1 - u))/u = 1$. Combining this with the fact that¹³

$$\int_0^1 AR(q) dq = \lim_{\varepsilon \rightarrow 0} \int_\varepsilon^1 AR(q) dq.$$

implies (83).

□

¹³Despite appearances, AR is well-behaved at zero: for $q > 0$,

$$1 \leq \frac{1}{q} - \frac{\ell(1-q)^\ell}{1-(1-q)^\ell} \leq \frac{\ell+1}{2}.$$