Theoretical Economics 16 (2021), 911–942 1555-7561/20210911

Matching with floor constraints

Sumeyra Akin
Instituto de Economía, Pontificia Universidad Católica de Chile

Floor constraints are a prominent feature of many matching markets, such as medical residency, teacher assignment, and military cadet matching. We develop a theory of matching markets under floor constraints. We introduce a stability notion, which we call floor respecting stability, for markets in which (hard) floor constraints must be respected. A matching is floor respecting stable if there is no coalition of doctors and hospitals that can propose an alternative matching that is feasible and an improvement for its members. Our stability notion imposes the additional condition that a coalition cannot reassign a doctor outside the coalition to another hospital (although she can be fired). This condition is necessary to guarantee the existence of stable matchings. We provide a mechanism that is strategy-proof for doctors and implements a floor respecting stable matching.

Keywords. Matching, floor constraints, efficiency, stability, strategy-proofness.

JEL classification. C78, D61, D63.

1. Introduction

The theory of two-sided matching has been developed and applied in many assignment markets. As it has been applied in various types of environments, designers have encountered distributional constraints that require special consideration. In this paper, we study hard floor constraints. This type of constraint arises when the designer has to guarantee the assignment of a minimum number of agents to individual institutions or to groups of institutions.

To fix ideas, consider the Turkish public school teacher assignment market. The assignment of teachers to public schools is determined by a centralized process overseen by the Turkish Ministry of Education. The education ministry targets a class size maximum of 30 students for elementary schools. This means that every school in the country must have a minimum number of teachers. We call these minimum requirements floor constraints. Thus, the government must assign teachers to schools in a way that respects preferences as much as possible while also respecting floor constraints. In this market, schools in the southeastern regions of Turkey are particularly difficult to fill due to their

Sumeyra Akin: sumeyra.akin@uc.cl
I thank the three anonymous referees for many helpful and productive suggestions. I am grateful to Lars Ehlers, Tibor Heumann, and Larry Samuelson for their detailed and thoughtful comments. I would like to especially thank Larry Samuelson for his time and invaluable support throughout the project. I would like to thank Dirk Bergemann, Mira Frick, Johannes Hörner, Costas Meghir, Esteban Peralta, Leeat Yariv, and seminar participants at Yale for their many useful suggestions for the paper.

 https://doi.org/10.3982/TE3785

economic and geopolitical conditions. Thus, the education ministry implements various policies to satisfy floor constraints in these regions.

Floor constraints are also prevalent in various other matching markets. For example, Terrier (2014) notes that in the assignment process of teachers to public schools in France, teachers may not be allowed to leave their regions because of geographical imbalances in the demand and supply of vacant positions. Floor constraints also feature in the assignment of newly graduated cadets to positions in U.S. military branches. A memorandum from the Army Deputy Chief of Staff notes that while the Army intends to accommodate the preferences of cadets as much as possible, the assignment process is subject to minimum staffing requirements (Fragiadakis and Troyan 2017). Finally, consider medical residency matching markets. An important concern in these markets is doctor shortages in rural communities, and governments implement various policies to address this problem (see, e.g., Talbott 2007 for the United States, Alcoba 2009 for India, and Nambiar and Bavas 2010 for Australia).

A common theme in all these markets is the desire to satisfy certain floor constraints while maintaining the "stability" of the market. Stability is maintained in the sense that no group of agents (for example, doctors) and institutions (for example, hospitals) can improve their assignments by matching outside the market. This feature ensures that the market is immune to common problems encountered in practice (e.g., instability may lead to unraveling or low participation in the clearinghouse (Roth 1984), and stability is desired from a normative perspective). We develop a theory of matching markets under hard floor constraints. That is, matchings that violate floor constraints are deemed infeasible in our model.[1] We introduce a stability notion that allows for hard floor constraints, and a mechanism that is strategy-proof for doctors and implements a stable matching.

Our matching model consists of a finite set of doctors and hospitals. Each hospital has a capacity constraint: the maximum number of doctors that the hospital can be assigned. We assume that every doctor and hospital prefers being assigned to some agent than to remaining unassigned. This is a reasonable assumption in our applications. For example, only 8% of K-12 students attend private schools in Turkey. Thus, teachers have very limited options outside the public school system. Outside options also play a minor role in most of the other markets mentioned above.

[1]The floor constraints in our model are different from *soft* floor constraints. Under soft floor constraints, the minimum requirements are interpreted as a guideline. A classic market that is subject to soft floor constraints is school choice markets, where schools have quotas or reserved seats for minorities, but these serve only as a benchmark (see, e.g., Echenique and Yenmez 2015 and Hafalir et al. 2013). For example, minority students are not forced to attend a school so that minority quotas are met, and it is rare to see that a student would like to move to a school that has a vacant seat and she cannot do so because the floor constraints at her current school would be violated. In this sense, the floor constraints are soft. In contrast, in the markets that we study, including the Turkish teacher assignment market, the floor constraints are interpreted as legal constraints, and active policies restrict agents' mobility. Thus, we may see agents and institutions that would like to be matched with each other, and even though the institutions' capacity constraints are not binding, such improvements do not realize because of the existing regulatory controls. For example, teachers in Turkey's southeastern regions are not allowed to move to schools in other areas, even when these schools have open slots.

We formalize the distributional constraints by assuming that only a subset of all matchings that respect hospital capacities are feasible. The set of feasible matchings is a subset of all matchings that satisfy the following property. We require that if a matching is feasible, then any matching that assigns weakly more doctors to hospitals is also feasible (as long as it respects the hospital capacities). Our definition of feasibility generalizes the idea of regional floor constraints present in the examples given previously. We provide a formal definition of regions and their floor constraints, and show that the set of feasible matchings, in this case, satisfies the property mentioned above.

Our stability notion is defined as follows. A matching is *floor respecting stable* if there is no group of doctors and hospitals that can propose a feasible matching that is a weak improvement for every member of the coalition and strictly improves some of its member. The key element of this definition is the specification that a doctor outside the coalition cannot be reassigned to another hospital without her consent (although a doctor *can* be fired without her consent). This condition is absent from the stability notion introduced previously in the literature (Ehlers et al. 2014) and is the key element of our definition that allows us to guarantee the existence of stable matchings.

Our first result states that any floor respecting stable matching is (constrained) Pareto efficient. A matching is Pareto efficient if there is no other feasible matching that (weakly) improves the assignment of every doctor and hospital. This result is significant because Pareto efficiency is an essential criterion for a stability notion from both a positive and a normative perspective. Interestingly, the absence of blocking pairs (of doctors and hospitals) is not sufficient to verify Pareto efficiency in markets under floor constraints. Thus, we conclude that a proper stability notion for markets under floor constraints should consider coalitions of any size.

We propose the following two-stage mechanism. We give priority numbers to doctors, with lower numbers indicating higher priorities. In the first stage, the deferred acceptance algorithm (DA, henceforth) is run on doctors with priority numbers lower than some cutoff priority number and all hospitals. In the second stage, we determine the assignment of the rest of the doctors by running a variant of the serial dictatorship mechanism (SD, henceforth); doctors sequentially choose from a restricted set of hospitals. The cutoff priority number is endogenously determined by the floor constraints and the preference profile. The rest of our results show that our mechanism has desirable properties.

We first show that our mechanism produces the outcome of the DA (run on all doctors and hospitals) if the floor constraints are not binding in the matching implemented by the DA. Thus, when there are no binding floor constraints, our mechanism has all the desirable properties of the DA and departs from the DA only when the DA would not produce a feasible outcome. This feature is desirable because the DA has been successfully applied in many settings, and departing from the DA is desirable only if doing so is strictly necessary.

Second, we show that our mechanism produces a floor respecting stable matching. From a practical perspective, this result is important because, as noted above, stability plays a crucial role in the success of matching market design. From a theoretical perspective, we obtain the existence of floor respecting stable matchings as an immediate corollary.

Finally, we show that our mechanism is strategy-proof for doctors. The main challenge to proving this result is that the mechanism uses the reported preference profile to determine the number of doctors who are assigned in the first stage. Therefore, it is conceivable that doctors may have incentives to misreport their preferences to receive more preferred assignments by influencing who is assigned in each stage. Nevertheless, this problem does not arise in our case because of the way that we define the cutoff priority number.

2. Literature review

There is a rich literature studying matching markets under distributional constraints. However, few papers study matching markets under hard floor constraints. We believe that one of the reasons for this is that there is no suitable notion of stability for markets under floor constraints. Thus, the literature is largely divided between papers finding mechanisms that have desirable properties and papers using ceiling constraints to indirectly satisfy floor constraints. Our contribution is to provide a stability notion that accommodates floor constraints and a strategy-proof mechanism that implements a (floor respecting) stable matching.

To the best of our knowledge, the study of hard floor constraints in matching markets was pioneered by Ehlers (2010). An example developed by Ehlers et al. (2014) (which supersedes Ehlers 2010) show that an immediate extension of the stability notion to one that accommodates floor constraints may result in the nonexistence of a stable matching.[2] We show that this nonexistence problem disappears when one imposes a condition on the reassignment of doctors outside the coalition. A similar solution appears in the context of an allocation problem. Ehlers (2014) shows that an efficient allocation may fail to exist when we allow for indifferences in agents' preferences, but its existence may be guaranteed by imposing restrictions on the reassignment of objects.[3]

Fragiadakis et al. (2016) take the impossibility result of Ehlers et al. (2014) and propose two mechanisms that offer different trade-offs between fairness and nonwastefulness. Tomoeda (2018) adopts the stability notion introduced by Ehlers et al. (2014) and finds sufficient conditions on hospital preferences such that a stable matching exists.

Kamada and Kojima (2017b) propose stability notions for a broad class of matching markets under ceiling constraints. That is, there is a cap on the number of doctors who can be assigned to hospitals or groups of hospitals. Kamada and Kojima (2015) offer a strategy-proof stable mechanism in the Japanese medical residency matching context. Kamada and Kojima (2018) find necessary and sufficient conditions on the structure of ceiling constraints such that a strategy-proof stable mechanism always exists.[4] Goto

[2]In their paper, fairness and non-wastefulness, which together comprise stability, are defined separately. They introduce a fairness notion in the presence of floor constraints and show that there may not exist a fair and non-wasteful matching.

[3]In Ehlers (2014), the weaker efficiency notion is called normal efficiency. After Proposition 2, they remark that the nonexistence result of Proposition 2 is reversed when they require normal efficiency instead of (strong) efficiency.

[4]A summary of their findings can be found in Kamada and Kojima (2017a).

et al. (2017) propose a strategy-proof mechanism that is non-wasteful and weakly Pareto efficient, but the implemented outcome may not be stable.

Recognizing that in many situations, ceiling constraints are introduced to satisfy floor constraints, Fragiadakis and Troyan (2017) offer a strategy-proof mechanism that, by endogenously adjusting the ceiling constraints, finds a fair matching that satisfies the floor constraints. They show that doctors unanimously prefer their mechanism over the mechanism with fixed ceiling constraints.[5] In contrast to our mechanism, their mechanism is wasteful.

We also note that the tools developed for the study of ceiling constraints cannot be applied to the study of matching markets subject to floor constraints. Intuitively, a stability notion that accommodates ceiling constraints prevents doctors from moving to an overdemanded region, while a stability notion that respects floor constraints prevents doctors from moving away from an underserved region.

A different strand of the literature studies diversity concerns in the school choice context.[6] In these models, distributional constraints serve only as a guideline rather than a requirement. That is, the constraints might be violated in a feasible matching.

Matching markets with distributional constraints also garner interest in the computer science literature. For example, papers by Biró et al. (2010) and Huang (2010) show the NP-completeness of checking whether a stable matching exists in the presence of distributional constraints and find a class of distributional constraints in which this problem does not arise. Goto et al. (2014) report a similar result and offer two strategy-proof fair mechanisms.

3. Model

Doctors and hospitals There is a finite set of doctors D and hospitals H. Each doctor $d \in D$ has a strict preference $\succ_d$ over hospitals and remaining unassigned, denoted by $\varnothing$. Each hospital h is endowed with a physical capacity $q_h > 0$ and has a strict preference $\succ_h$ over the subsets of doctors. A preference profile $(\succ_i)_{i \in D \cup H}$ is denoted by $\succ$. For any $h, h' \in H \cup \{\varnothing\}$, we write $h \succeq_d h'$ if and only if $h \succ_d h'$ or $h = h'$; the corresponding notation is used for hospital preferences.

Assumptions on the preferences We assume throughout the paper that the hospital preferences are *responsive* (cf. Roth 1985a). Formally, hospital h's preference $\succ_h$ is responsive if for every $D' \subseteq D$ with $|D'| < q_h$ and $i, j \in D \setminus D'$,

(i) $D' \cup \{i\} \succ_h D' \cup \{j\}$ if and only if $i \succ_h j$

(ii) $D' \cup \{i\} \succ_h D'$ if and only if $i \succ_h \varnothing$.

In other words, if two subsets of doctors differ by only one doctor, then the hospital prefers the subset of doctors containing the more preferred doctor. This condition rules out any complementarities in hospital preferences.

[5]The mechanism with fixed ceiling constraints is called artificial caps deferred acceptance.

[6]Recent studies on diversity concerns in matching markets include Abdulkadiroğlu and Sönmez (2003), Abdulkadiroğlu (2005), Hafalir et al. (2013), Echenique and Yenmez (2015), Erdil and Kumano (2013), Dur et al. (2018), and Aygün and Turhan (2020).

We also assume that there are no individual rationality constraints. Formally, $h \succ_d \varnothing$ and $d \succ_h \varnothing$ for every $d \in D$ and $h \in H$. That is, every hospital is acceptable to every doctor and every doctor is acceptable to every hospital.

Matching function A matching μ is a mapping, $\mu : D \cup H \to D \cup H \cup \{\varnothing\}$, that satisfies (i) $\mu(d) \in H \cup \{\varnothing\}$ for all $d \in D$, (ii) $\mu(h) \subseteq D$ for all $h \in H$, and (iii) for any $d \in D$ and $h \in H$, $d \in \mu(h)$ if and only if $\mu(d) = h$.

Feasible matchings The set of feasible matchings, denoted by $\mathbb{M}$, is a nonempty subset of all matchings that satisfy the following properties:

(i) If $\mu \in \mathbb{M}$, then for all $h \in H$, $|\mu(h)| \leq q_h$.

(ii) If $\mu \in \mathbb{M}$, then every matching μ' satisfying $q_h \geq |\mu'(h)| \geq |\mu(h)|$ for every $h \in H$ is feasible (i.e., $\mu' \in \mathbb{M}$).

The first condition states that every matching in $\mu \in \mathbb{M}$ respects the hospital capacities. The second condition states that if a matching is feasible, then every matching that assigns weakly more doctors to each hospital is also feasible (as long as it respects hospital capacities). We say that a matching is feasible if and only if $\mu \in \mathbb{M}$. To understand better the implications of the second condition, we note its contraposition. The contraposition states that if a matching is not feasible (and it respects hospital capacities), then any matching that assigns weakly fewer doctors to each hospital is also not feasible. Intuitively, this feasibility condition requires a certain number of doctors at each hospital and formalizes the general notion of floor constraints in our paper.

Regional floor constraints Our leading motivation is the study of regional floor constraints, which we now formally introduce. We discuss our results as if the set of feasible matchings $\mathbb{M}$ is constructed from the set of matchings that satisfy some regional floor constraints (as we introduce next), but our results carry through to any set of matchings $\mathbb{M}$ that satisfy the two aforementioned properties. Hence, the introduction of regions is only for the interpretation of our results.

There is a set of regions R and each hospital h resides in a subset of regions $r(h) \subseteq R$. We denote the set of hospitals in region r by H_r (i.e., $H_r = \{h \in H : r \in r(h)\}$). There is a minimum number of doctors that need to be assigned to the hospitals of region r, which we denote by $\underline{q}_r$. Matching μ satisfies the floor constraints if for all $r \in R$,

$$\underline{q}_r \leq \sum_{h \in H_r} |\mu(h)|.$$

The set of feasible matchings $\mathbb{M}$ contains all the matchings that satisfy hospital capacities and the floor constraints. We note that in this case, $\mathbb{M}$ satisfies the conditions stated above.

A particular case of interest is when each hospital belongs to only one region, in which case we say that we have *disjoint regions*. This case is the simplest one to study because the assignment of each doctor can relax the floor constraint of at most one region. Additionally, most of the papers in the literature on distributional constraints allow for disjoint regions. For these reasons, we consider disjoint regions to illustrate our results.

Throughout the paper, we assume that $\mathbb{M} \neq \varnothing$. For the study of regional floor constraints, a sufficient condition for the existence of feasible matchings is $\sum_{r\in R} \underline{q}_r \leq |D|$. That is, the number of doctors is greater than the total number of floor constraints. We note that this condition is necessary when regions are disjoint.

Examples of regional floor constraints We now provide some examples of different regional floor constraints. Suppose that $H = \{h_1, h_2, h_3, h_4\}$. We omit the reference to the hospital capacities and the floor constraints. We simply describe different sets of regions and provide an interpretation of them.

As a first example, suppose that there are two regions and hospitals are distributed across regions as

$$r(h_1) = r(h_2) = \{r_1\} \quad \text{and} \quad r(h_3) = r(h_4) = \{r_2\}.$$

That is, hospitals h_1 and h_2 reside in region r_1 and hospitals h_3 and h_3 reside in region r_2. In this case, the regions are disjoint.

As a second example, suppose that there are six regions and hospitals are distributed across regions as

$$r(h_1) = \{r_1, r_3\} \quad \text{and} \quad r(h_2) = \{r_1, r_4\}$$
$$r(h_3) = \{r_2, r_5\} \quad \text{and} \quad r(h_4) = \{r_2, r_6\}.$$

That is, hospitals are distributed to regions r_1 and r_2 in the same way as in the first example, and, additionally, each hospital belongs to a distinct region. This distributional constraint can accommodate a situation in which we have regional floor constraints presented in the first example. Additionally, there are floor constraints on the number of doctors assigned to each hospital (e.g., because hospitals need a minimum number of doctors to operate). The floor constraints on regions r_3–r_6, respectively, indicate the minimum number of doctors who need to be assigned to hospitals h_1–h_4.

As a third example, suppose that there are three regions and hospitals are distributed across regions as

$$r(h_1) = \{r_1\} \quad \text{and} \quad r(h_2) = \{r_1, r_3\}$$
$$r(h_3) = \{r_2, r_3\} \quad \text{and} \quad r(h_4) = \{r_2\}.$$

That is, hospitals are distributed to regions r_1 and r_2 in the same way as in the first example and, additionally, hospitals h_2 and h_3 reside in region r_3. This distributional constraint can accommodate a situation where hospitals are located on a line, patients live between the hospitals, and patients need care at an adjacent hospital. That is, patients living between hospital h_1 and h_2 are those living in region r_1; patients living between hospital h_2 and h_3 are those living in region r_3; patients living between hospital h_3 and h_4 are those living in region r_2.

As a fourth example, suppose that there are three regions and hospitals are distributed across regions as

$$r(h_1)=\{r_1\} \quad \text{and} \quad r(h_2)=\{r_1,r_2,r_3\}$$
$$r(h_3)=\{r_2\} \quad \text{and} \quad r(h_4)=\{r_3\}.$$

That is, every hospital except hospital h_2 is located in a distinct region and hospital h_2 shares regions with others. This distributional constraint can accommodate a situation where hospital h_2 is located at the center of an area and the rest of the hospitals are located on the periphery. Each patient can go either to the central hospital or to their "local" peripheral hospital.

4. Stability notion with floor constraints

We begin this section by introducing our stability notion that accommodates floor constraints: floor respecting stability. We then discuss the key element in the definition of stability that is necessary to guarantee the existence of stable matchings. We also show that every floor respecting stable matching is Pareto efficient and explain why it is necessary to consider coalitions of size larger than 2 to verify efficiency.

4.1 *Floor respecting stability*

We first define our stability notion and then discuss the elements of this definition.

Definition 1 (Floor respecting stability). We say that a matching $\mu \in \mathbb{M}$ is blocked by coalition $A \subseteq D \cup H$ if there is another matching $\mu' \in \mathbb{M}$ such that

(i) $\mu'(d) \in A$ and $\mu'(d) \succeq_d \mu(d)$ for all doctors $d \in A \cap D$

(ii) $\mu'(h) \subseteq A$ and $\mu'(h) \succeq_h \mu(h)$ for all hospitals $h \in A \cap H$

(iii) $\mu'(i) \succ_i \mu(i)$ for some doctor or hospital $i \in A$

(iv) if doctor $d \notin A$, then $\mu'(d)=\mu(d)$ or $\mu'(d)=\varnothing$.

A matching μ is floor respecting stable if it is feasible and there is no coalition blocking it.

A feasible matching is floor respecting stable if there is no coalition that can propose another feasible matching that satisfies the four conditions in Definition 1. The first three conditions state that all doctors and hospitals that belong to the coalition are better off under the alternative matching, with at least one of them being strictly better off. If there are no distributional constraints (that is, $\mathbb{M}$ coincides with the set of all matchings that respect hospital capacities), then our definition coincides with the definition of the (weakly dominating) core (cf. Roth 1985b) and the classic stability notion (cf.

Gale and Lloyd 1962).[7] This is because when there are no floor constraints, a blocking coalition can always fire doctors, so condition (iv) becomes irrelevant.

The key novelty in our stability definition is condition (iv). This condition requires that if a doctor is assigned to a hospital other than her current assignment, then she belongs to the coalition. In other words, a coalition can only propose a matching in which doctors are either reassigned to more preferred hospitals or fired (or assigned to their original hospitals). This condition prevents a coalition from forcefully reassigning a doctor to a less preferred hospital. For example, suppose that a hospital and a doctor would like to block some status quo matching, but the doctor's original assignment is located in a region with a binding floor constraint. That is, if doctor d were to leave her current region, the region's floor constraint would be violated. Then, to meet the floor constraint, another doctor needs to be assigned to this region. At this point, condition (iv) stipulates that the change in the assignment of the other doctor should be an improvement for her. In other words, if there is no doctor willing to fill the vacancy created, then the hospital and the doctor cannot block the matching.

The implicit assumption behind our stability definition is that there is a government or a central authority that prevents blocking coalitions from being formed if floor constraints are violated. Indeed, matching markets under floor constraints seem to be appropriately regulated. For example, in Turkey, the Ministry of Education assigns newly graduated teachers predominantly to underserved areas every year. Furthermore, if a teacher wants to change her school, she has to submit a relocation application, and her request needs the ministry's approval even if there is a vacancy at the other school.[8] Although a priori the government could implement any matching without regard to the preferences, this would lead to the perception that the assignment process is potentially inefficient or unfair. Hence, our stability notion is a criterion that the authorities can adopt on efficiency grounds or to eliminate blocking coalitions that are not justified by the floor constraints, which we see as desirable from a normative perspective.[9]

The blocking coalitions that are regarded as illegitimate only because they violate condition (iv) (i.e., they require reassigning a doctor against her consent) are unlikely to materialize in markets subject to floor constraints. Specifically, the blocking coalitions that reassign a doctor to another hospital without her consent are inadmissible in most of the markets that we consider (which we mentioned in the Introduction). For example, in Turkey, teachers who are currently teaching cannot be reassigned against their consent.[10] Similarly, hospitals usually have the power to fire doctors, but not the right

[7] Roth (1985b) proves that the (weakly dominating) core coincides with the set of stable matchings.

[8] Another example of a heavily regulated matching market is the Japanese medical residency market (cf. Kamada and Kojima 2015).

[9] In the school choice context, for example, our stability notion corresponds to a "no justified envy" condition, and the literature usually views it as a normative criterion.

[10] A forceful reassignment may occur if a teacher does not fulfill her compulsory service duty in underserved regions on time and does not submit a preference list over the schools located in these regions. However, this situation does not arise because the teacher's original school hires a more preferred teacher; rather, it arises to satisfy floor constraints.

to reassign them to a different hospital, which is consistent with the fact that coalitions cannot violate condition (iv).

If we drop condition (iv) from our stability definition, a stable matching may fail to exist (as shown by Ehlers 2010). Ensuring the existence of stable matchings is important from a theoretical perspective. Thus, relaxing the stability notion is unavoidable in the presence of floor constraints. In Section 4.2, we explore alternatives to condition (iv) and demonstrate its necessity for the existence of stable matchings. In Section 4.3, we show that floor respecting stable matchings are Pareto efficient. Interestingly, in the presence of floor constraints, considering only blocking pairs (that is, pairs of doctors and hospitals) is not sufficient to verify Pareto efficiency, which we explain in Section 4.4.

4.2 *Condition (iv) and the existence of stable matchings*

Condition (iv) in Definition 1 is critical to guarantee the existence of floor respecting stable matchings. Ehlers et al. (2014) adopt a stability notion similar to ours, but they allow coalitions to reassign doctors who are not in the coalition to different hospitals. They show that a stable matching may not exist.[11]

Alternatively, we could consider a stability notion as in Definition 1, but change condition (iv) to

$$\text{if doctor} \quad d \notin A, \quad \text{then} \quad \mu'(d) = \mu(d) \quad \text{or} \quad \mu'(d) \neq \varnothing. \tag{1}$$

This stability notion allows a blocking coalition to reassign doctors who are not in the coalition without any restriction, but prevents them from being fired. However, such a stable matching may not exist (shown in Example 1).

We now explain why condition (iv) is necessary to guarantee the existence of a floor respecting stable matching with an example.

Example 1. There are two doctors and three hospitals. The hospitals' capacities are given by $q_{h_i} = 1$ for $i = 1, 2, 3$. The preference profile is

$$\begin{aligned} &\succ_{d_1} : h_2, h_1, h_3, && \succ_{h_1} : d_1, d_2 \\ &\succ_{d_2} : h_1, h_2, h_3, && \succ_{h_2}, \succ_{h_3} : d_2, d_1. \end{aligned}$$

There are three regions, r_1, r_2, and r_3, the floor constraints of which are given by $\underline{q}_{r_1} = \underline{q}_{r_2} = 0$ and $\underline{q}_{r_3} = 1$. The hospitals are located across regions as $r(h_i) = \{r_i\}$ for $i = 1, 2, 3$. Because we revisit this example later, we summarize the market in Table 1.

[11] In their paper, the stability notion is separated into fairness and non-wastefulness. They show that a feasible fair and non-wasteful matching does not exist.

Hospitals				
	Preferences	Hospitals' capacities	Region of the hospital	Region's floor constraint
h_1	$d_1 \succ_{h_1} d_2$	1	$r(h_1) = \{r_1\}$	0
h_2	$d_2 \succ_{h_2} d_1$	1	$r(h_2) = \{r_2\}$	0
h_3	$d_2 \succ_{h_3} d_1$	1	$r(h_3) = \{r_3\}$	1
Doctors				
d_1	$h_2 \succ_{d_1} h_1 \succ_{d_1} h_3$			
d_2	$h_1 \succ_{d_2} h_2 \succ_{d_2} h_3$			

TABLE 1. Preferences, capacities, regions, and floor constraints.

To satisfy the floor constraints, one of the doctors has to be assigned to hospital h_3. Consider the feasible matchings[12]

$$\mu_1 = \begin{pmatrix} h_1 & h_2 & h_3 \\ d_1 & \varnothing & d_2 \end{pmatrix}, \qquad \mu_2 = \begin{pmatrix} h_1 & h_2 & h_3 \\ \varnothing & d_1 & d_2 \end{pmatrix}$$

$$\mu_3 = \begin{pmatrix} h_1 & h_2 & h_3 \\ \varnothing & d_2 & d_1 \end{pmatrix}, \qquad \mu_4 = \begin{pmatrix} h_1 & h_2 & h_3 \\ d_2 & \varnothing & d_1 \end{pmatrix}.$$

Matching μ_1 (resp. μ_3) is not floor respecting stable. Coalition $\{d_1, h_2\}$ (resp. $\{d_2, h_1\}$) would block matching μ_1 (resp. μ_3) with matching μ_2 (resp. μ_4).

Matching μ_2 is floor respecting stable. Coalition $\{d_2, h_2\}$ would like to block matching μ_2, but it cannot do so because its alternative matching would violate the floor constraints. To satisfy the floor constraints, coalition $\{d_2, h_2\}$ needs to assign doctor d_1 to hospital h_3. By condition (iv) of our stability notion, doctor d_1 needs to be included in the coalition. Coalition $\{d_1, d_2, h_2, h_3\}$ (with alternative matching μ_3) cannot block matching μ_2 because doctor d_1 is strictly worse off in the alternative matching. Similar arguments can be used to show that matching μ_4 is also floor respecting stable.

If we consider a definition of stability as in Definition 1 but drop condition (iv) or change it to the condition in (1), then no stable matching would exist under this new definition. Specifically, matching μ_2 (resp. μ_4) would be blocked by coalition $\{d_2, h_2\}$ (resp. $\{d_1, h_1\}$) with alternative matching μ_3 (resp. μ_1). Matchings μ_1 and μ_3 would be blocked by the same coalitions that arise when we assume Definition 1. ◊

4.3 *Stability and Pareto efficiency*

One important reason why stability is a desirable criterion in matching markets is that it implies Pareto efficiency. In this section, we show that any floor respecting stable matching is Pareto efficient. A matching $\mu \in \mathbb{M}$ is *Pareto efficient* if there is no other matching

[12]There are two more feasible matchings in which some doctor remains unassigned: $\mu_5 = \left(\begin{smallmatrix} h_1 & h_2 & h_3 \\ \varnothing & \varnothing & d_2 \end{smallmatrix}\right)$ and $\mu_6 = \left(\begin{smallmatrix} h_1 & h_2 & h_3 \\ \varnothing & \varnothing & d_1 \end{smallmatrix}\right)$. We exclude them from the discussion because it is clear that these matchings are not stable. The unassigned doctors and any hospital with a vacant position can block these matchings.

$\mu' \in \mathbb{M}$ such that $\mu'(i) \succeq_i \mu(i)$ for all $i \in D \cup H$ and $\mu'(i) \succ_i \mu(i)$ for some $i \in D \cup H$. We now state the efficiency of floor respecting stable matchings.

Proposition 1 (Pareto Efficiency). *Every floor respecting stable matching is Pareto efficient.*

If a matching is not Pareto efficient, then there exists another feasible matching that is a weak improvement for every agent and strictly improves some agent. Thus, the grand coalition can block a Pareto inefficient matching with a matching that Pareto dominates it.

4.4 *Blocking pairs versus larger coalitions*

In (unconstrained) matching markets, the stability notion of Gale and Lloyd (1962) considers only two-member coalitions, that is, pairs of doctors and hospitals, and every stable matching is Pareto efficient.[13] However, to guarantee the efficiency of stable matchings in markets under floor constraints, an appropriate stability notion should consider any-size coalitions. This is because a pair of a doctor and a hospital may not be powerful enough to achieve an improvement for themselves without violating some floor constraints. A group of doctors and hospitals, however, might be able to achieve an improvement for themselves while satisfying the floor constraints (we illustrate this feature below in Example 2).

Given that coalitions of size 2 are not sufficient to verify whether a matching is floor respecting stable (or even efficient), we now characterize in greater detail when it is necessary to consider coalitions of size larger than 2.

Proposition 2 (Size and type of blocking coalitions). *Suppose that matching μ is not floor respecting stable. Then there exists a coalition A that blocks matching μ with matching μ' that satisfies one of the following statements:*

(i) Coalition A does not fire any doctor, i.e., for every doctor $d \in D$, $\mu'(d) = \varnothing$ implies that $\mu(d) = \varnothing$.

(ii) Coalition A consists of one doctor and one hospital.

Proposition 2 states that we can verify whether a matching is floor respecting stable by checking for two types of blocking coalitions. The first is blocking coalitions that do not fire any doctor. By condition (iv) in Definition 1, such coalitions propose a matching that is a Pareto improvement for doctors. The second type of blocking coalitions is pairs of doctors and hospitals.

We now provide an example to illustrate the extent to which the absence of blocking pairs may be insufficient to verify that a matching is stable or even Pareto efficient. We first provide a matching that cannot be blocked by a pair of a doctor and a hospital, but

[13]In (unconstrained) matching markets, matching μ is stable if there is no pair of a doctor and a hospital that prefers another one over their current assignments at μ.

that can be blocked by the grand coalition. In other words, a matching may be Pareto inefficient despite there being no blocking pairs. We then provide a matching that can be blocked by only one coalition that does not fire any doctor, but this coalition leaves a hospital worse off. In other words, the *only* blocking coalition proposes a matching that is not a Pareto improvement for all agents in the economy (only for doctors).

Example 2. There are two doctors and three hospitals. The hospitals' capacities are given by $q_{h_i} = 1$ for $i = 1, 2, 3$. The preference profile is

$$\begin{array}{ll} \succ_{d_1} : h_1, h_3, h_2, & \succ_{h_1} : d_1, d_2 \\ & \succ_{h_2} : d_2, d_1 \\ \succ_{d_2} : h_3, h_2, h_1, & \succ_{h_3} : d_2, d_1. \end{array}$$

There are two regions with floor constraints given by $\underline{q}_{r_i} = 1$ for $i = 1, 2$. The hospitals are located across regions as $r(h_1) = \{r_1\}$ and $r(h_2) = r(h_3) = \{r_2\}$.

Consider the matching

$$\mu' = \begin{pmatrix} h_1 & h_2 & h_3 \\ d_2 & \varnothing & d_1 \end{pmatrix}.$$

It is easy to verify that this matching is not Pareto efficient. The grand coalition blocks matching μ' with the matching

$$\begin{pmatrix} h_1 & h_2 & h_3 \\ d_1 & \varnothing & d_2 \end{pmatrix}.$$

Nonetheless, there is no pair of a doctor and a hospital that can block matching μ'. For example, if hospital h_1 hires doctor d_1 and fires doctor d_2, this would violate region r_2's floor constraint (analogously, (d_2, h_2) or (d_2, h_3) cannot block this matching). Thus, it is clear that allowing for coalitions larger than one doctor and one hospital is natural in our environment.

Now suppose that the preference of doctor d_1 is $\succ_{d_1}: h_1, h_2, h_3$ and that of doctor d_2 is $\succ_{d_2}: h_3, h_1, h_2$, while everything else is the same as previously specified. Consider the matching

$$\mu'' = \begin{pmatrix} h_1 & h_2 & h_3 \\ d_2 & d_1 & \varnothing \end{pmatrix}.$$

It is easy to verify that this matching is not floor respecting stable (although Pareto efficient). The only blocking coalition is $A = \{d_1, d_2, h_1, h_3\}$, and it proposes a matching that leaves hospital h_2 worse off than under μ''. The example illustrates that to verify whether a matching is floor respecting stable, it is necessary to consider blocking coalitions of size larger than 2, and these coalitions may leave some hospitals worse off. $\diamond$

5. Implementation

In this section, we propose a strategy-proof (for doctors) mechanism that produces a floor respecting stable matching and discuss its properties.

5.1 *Preliminary definitions*

A mechanism φ is a function that maps preference profiles to matchings. We assume that every doctor must report a full ranking of hospitals.[14] Throughout this section, we use $\succ$ and $\tilde{\succ}$ to refer to the true preference profile and a reported preference profile, respectively. Given a preference profile $\tilde{\succ}$, we denote doctor d's assignment by $\varphi(\tilde{\succ})(d)$.

A mechanism φ is strategy-proof for doctors if there do not exist a preference profile $\succ$, a doctor $d \in D$, and a preference $\tilde{\succ}_d$ of doctor d such that

$$\varphi(\tilde{\succ}_d, \succ_{-d})(d) \succ_d \varphi(\succ_d, \succ_{-d})(d).$$

In other words, doctor d never profits from reporting a preference different from her true preference.

We now describe the (doctor-proposing) deferred acceptance algorithm.

Algorithm 1: Deferred acceptance

Step 1. Each doctor applies to her most preferred hospital. Each hospital rejects any unacceptable doctor and the lowest ranked doctors in excess of its capacity, while it holds the rest of the doctors temporarily.

Step k. Each doctor who was rejected in step $k-1$ applies to her most preferred hospital that has not yet rejected her. Each hospital considers these new doctors and all the doctors who were temporarily held in step $k-1$, rejects the lowest ranked doctors in excess of its capacity and all unacceptable doctors, and temporarily holds the rest of the doctors.

The DA stops at the step in which no rejection occurs and assigns hospitals to the doctors that they still hold.

We denote by $\mathrm{DA}(\tilde{\succ}|Y, H)$ the outcome of the DA when the reported preference profile is $\tilde{\succ}$ and only the subset of doctors $Y \subseteq D$ is present in the market. The assignment of doctor $d \in Y$ is denoted by $\mathrm{DA}(\tilde{\succ}|Y, H)(d)$.

5.2 *Description of the mechanism*

In this section, we introduce a mechanism that is strategy-proof for doctors and implements a floor respecting stable matching.

To run mechanism φ, we label doctors from 1 to N, i.e., $D = \{1, \dots, N\}$, which is interpreted as doctors' priority numbers. Lower labels indicate higher priorities (e.g., $d = 1$ is the highest priority doctor). Given the priority order and preference profile $\tilde{\succ}$, mechanism φ first determines a cutoff priority number $\tilde{\nu}$ and then runs a two-stage algorithm.

Before we continue, we make the following preliminary definitions. Suppose that the reported preference profile is $\tilde{\succ}$. To simplify the notation, we define

$$\tilde{\mu} \triangleq \varphi(\tilde{\succ}).$$

[14] If a doctor declares some hospital unacceptable, then she would obviously be misreporting and the mechanism would leave her unmatched.

That is, $\tilde{\mu}$ is the outcome of mechanism φ when the reported preference profile is $\tilde{\succsim}$. We also define, for any $d \leq n$,

$$\tilde{m}(d; n) \triangleq \mathrm{DA}\big(\tilde{\succsim}|\{1, \ldots, n\}, H\big)(d). \tag{2}$$

That is, $\tilde{m}(d; n)$ is the assignment of doctor d under the DA when only the subset of doctors $\{1, \ldots, n\} \subseteq D$ is present in the market and the reported preference profile is $\tilde{\succsim}$.

We first explain the two stages for a given cutoff $\tilde{\nu}$ and then define the cutoff priority.

Stage 1 Mechanism φ assigns the subsets of doctors with priority number weakly lower than the cutoff $\tilde{\nu}$ (i.e., $\{1, \ldots, \tilde{\nu}\}$) using the DA. In the first stage, the DA is run only on doctors in $\{1, \ldots, \tilde{\nu}\}$ and all hospitals in H. Formally, for all $d \in \{1, \ldots, \tilde{\nu}\}$,

$$\tilde{\mu}(d) = \tilde{m}(d; \tilde{\nu}). \tag{3}$$

This concludes the mechanism's first stage.

Stage 2 To determine the assignment of doctors in $\{\tilde{\nu} + 1, \ldots, N\}$, we use a variant of the SD. Doctors in $\{\tilde{\nu} + 1, \ldots, N\}$ are sequentially assigned according to their priority numbers. Each doctor in $\{\tilde{\nu} + 1, \ldots, N\}$ is matched to her most preferred hospital in $\tilde{H}(d) \subseteq H$. The set $\tilde{H}(d)$ is constructed inductively as

$$\tilde{H}(d) \triangleq \big\{h \in H : \exists \mu \in \mathbb{M} \text{ such that for all } d' < d, \mu\big(d'\big) = \tilde{\mu}\big(d'\big) \text{ and } \mu(d) = h\big\}. \tag{4}$$

In other words, for every doctor $d \in \{\tilde{\nu} + 1, \ldots, N\}$, to evaluate whether a hospital $h \in \tilde{H}(d)$, we verify whether there exists a feasible matching when doctor d is matched to h and all doctors $d' < d$ are matched according to $\tilde{\mu}$.

Cutoff $\tilde{\nu}$ We now define the cutoff priority $\tilde{\nu}$ as

$$\tilde{\nu} = \min \left\{ n \in \{0, \ldots, N-1\} : \begin{array}{l} \nexists \mu \in \mathbb{M} \text{ such that for all } d \leq n, \\ \mu(d) = \tilde{m}(d; n) \text{ and } \mu(n+1) = \varnothing \end{array} \right\}. \tag{5}$$

If there is no number $\tilde{\nu}$ that satisfies the expression in (5), we let $\tilde{\nu} = N$. That is, $\tilde{\nu}$ is the minimum number such that when doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned by the DA, every doctor in $\{\tilde{\nu}+1, \ldots, N\}$ is needed to find a feasible matching. Assuming that doctors in $\{1, \ldots, \tilde{\nu}\}$ are assigned via the DA, doctor $\tilde{\nu} + 1$ becomes critical in the sense that if she were to remain unassigned, one could not find a way to assign the rest of the doctors (i.e., $\{\tilde{\nu}+2, \ldots, N\}$) that leads to a feasible outcome.

5.3 *Description of the mechanism in disjoint regions*

As an illustration, we provide explicit expressions for $\tilde{\nu}$ and $\tilde{H}(\cdot)$ when the regions are disjoint, i.e., each hospital belongs to only one region.

We first explain how to find the cutoff $\tilde{\nu}$. The cutoff $\tilde{\nu}$ is the smallest number n of doctors such that when doctors in $\{1, \ldots, n\}$ are assigned by the DA in the first stage, the number of the rest of the doctors, which is $N - n$, is just enough to satisfy all the regional floor constraints. When regions are disjoint, the number of doctors that is necessary to meet the floor constraints at the end of the first stage is the difference between regions' floor constraints and the number of doctors that are assigned to the regions by the DA. Thus, we can write the cutoff $\tilde{\nu}$ formally as

$$\tilde{\nu} = \min\left\{ n \in \{0, 1, 2, \dots, N\} : \sum_{r \in R} \left\{ \underline{q}_r - \sum_{h \in H_r} \left|\tilde{m}(h; n)\right| \right\}^+ \geq N - n \right\}. \tag{6}$$

The number of doctors assigned to hospital h through the DA is $|\tilde{m}(h; n)|$ and, hence, the (positive part of the) difference between the floor constraint $\underline{q}_r$ of region r and $\sum_{h \in H(r)} |\tilde{m}(h; n)|$ yields the number of doctors that region r lacks at the end of the first stage. For any number $n < \tilde{\nu}$, the inequality in (6) fails to hold under responsive hospital preferences, and at $n = \tilde{\nu}$, it is satisfied by equality. Thus, the expression in (6) provides the smallest number of doctors that can be assigned by the DA such that the floor constraints can be satisfied (at the end of the algorithm).

Mechanism φ assigns doctors in $\{1, \dots, \tilde{\nu}\}$ by the DA. That is, for every $d \leq \tilde{\nu}$,

$$\tilde{\mu}(d) = \tilde{m}(d; \tilde{\nu}).$$

The assignments of doctors in $\{\tilde{\nu} + 1, \dots, N\}$ are determined by applying the SD for a restricted subset of hospitals. To ensure the feasibility of the final matching, doctors in $\{\tilde{\nu} + 1, \dots, N\}$ should be assigned to the hospitals that (i) have a remaining vacant position and (ii) are in an underserved region (i.e., a region with a floor constraint strictly larger than the total number of doctors that have already been assigned to its hospitals). Formally, the subset of hospitals that are available to doctor d, $\tilde{H}(d)$, can be written formally, for all $d \in \{\tilde{\nu} + 1, \dots, N\}$, as

$$\tilde{H}(d) = \left\{ h \in H : \begin{array}{c} \left|\tilde{m}(h; \tilde{\nu})\right| + \displaystyle\sum_{j=\tilde{\nu}+1}^{d-1} \mathbb{1}\left\{\tilde{\mu}(j) = h\right\} < q_h, \\ \displaystyle\sum_{h' \in H_{r(h)}} \left(\left|\tilde{m}\left(h'; \tilde{\nu}\right)\right| + \sum_{j=\tilde{\nu}+1}^{d-1} \mathbb{1}\left\{\tilde{\mu}(j) = h'\right\} \right) < \underline{q}_{r(h)} \end{array} \right\}. \tag{7}$$

Note that $\tilde{H}(d)$ is constructed inductively using $\tilde{\mu}(j)$, with $\nu + 1 \leq j \leq d - 1$, which denotes the assignment of doctor j who is assigned via the SD before doctor d. The two inequalities in (7) formally express the two criteria that determine whether $h \in \tilde{H}(d)$. Doctor d is assigned to her most preferred hospital in $\tilde{H}(d)$; that is, $\tilde{\mu}(d) \succeq_d h$, for every $h \in \tilde{H}(d)$.

5.4 *Main elements of the mechanism*

Our mechanism has two stages. The first stage consists of the DA. This is a natural starting point because the floor respecting stability notion reduces to the classic stability notion of Gale and Lloyd (1962) when there are no floor constraints. Hence, if we are to look for a mechanism that is strategy-proof and produces a floor respecting stable matching, then it should coincide with the DA when the floor constraints are not present. In fact, in the following section, we show that our mechanism coincides with the DA when the floor constraints are slack.

We then have a second stage in which we assign doctors via the SD. The reason why we use the SD is less obvious. The SD satisfies two important properties: (a) it produces a doctor Pareto optimal matching and (b) is strategy-proof. The first property means that there is no other feasible matching that can improve the assignments of every doctor. This property is important to guarantee that the matching produced by the mechanism is floor respecting stable. The second property is important to guarantee that the mechanism is strategy-proof. We explain the role of each of these properties in the following sections when we discuss the properties of the mechanism.

The cutoff $\tilde{\nu}$ plays an important role in ensuring that the mechanism produces a floor respecting stable matching. If we assign too many doctors through the DA, we may not be able to ensure the feasibility of the final matching. Alternatively, if we assign too few (e.g., zero) doctors through the DA, then the resulting matching may not be stable because the SD does not produce a stable matching. Thus, the cutoff $\tilde{\nu}$ must be chosen such that the number of doctors assigned in the second stage is just enough to meet the floor constraints.

The cutoff $\tilde{\nu}$ also plays a crucial role in guaranteeing that the mechanism is strategy-proof. Although our mechanism consists of two strategy-proof mechanisms, it is not evident that the combination of the two will be strategy-proof, because the number of doctors assigned through the DA in the first stage depends on doctors' reported preferences. Thus, it is conceivable that doctors can change the cutoff priority in their favor by submitting different reports. Nevertheless, as we show in Section 7.1, mechanism φ is strategy-proof. We provide intuition for how we achieve this result and explain the role of $\tilde{\nu}$ in Sections 7.1 and 7.2.

5.5 *When the mechanism coincides with the DA*

Our mechanism coincides with the DA (runs on all doctors and hospitals) when the floor constraints are slack. Specifically, if the matching implemented by the DA runs on all doctors except doctor N is feasible, then our mechanism implements the same matching as the DA. The floor constraints are slack in the sense that we can exclude a doctor and still implement a feasible matching. We now state the result formally.

Proposition 3 (Mechanism Coincides With the DA When Constraints Do Not Bind). *Suppose that matching $\hat{\mu}$ defined as*

$$\hat{\mu}(d) = \begin{cases} \tilde{m}(d; N-1) & \text{if } d \neq N \\ \varnothing & \text{if } d = N \end{cases} \tag{8}$$

is feasible (i.e., $\hat{\mu} \in \mathbb{M}$). Then mechanism φ described by (3)–(5) implements the same matching as the DA (that is, $\mathrm{DA}(\succ|D, H) = \varphi(\succ)$*).*

This proposition states that if the DA implements a feasible matching without considering doctor N, then the outcome of mechanism φ will be the same as the outcome of the DA. In most markets under distributional constraints, the main reason for departing from the DA is that its outcome may not satisfy the constraints. Thus, we believe that the property stated in Proposition 3 is desirable.

We note that our mechanism may implement a different matching than the DA even though the DA may implement a feasible matching. This situation arises when doctor N is needed to implement a feasible matching by running the DA. Nevertheless, we expect that this nuance is unlikely to arise in markets with large numbers of participants because whether the DA would produce a feasible matching is independent of the assignment of any specific doctor (and, in particular, of doctor N). In other words, the event that the DA produces a feasible matching and the DA runs with doctors $\{1, \ldots, N-1\}$ does not produce a feasible matching is unlikely when the number of participants is large. Thus, our mechanism departs from the DA only when necessary.

6. Floor respecting stability

6.1 *Mechanism produces a floor respecting matching*

We now state that mechanism φ implements a floor respecting stable matching.

Proposition 4 (Floor respecting stable outcome). *Mechanism φ described by (3)–(5) produces a floor respecting stable matching.*

An important implication of Proposition 4 is that a floor respecting stable matching exists. We formally state this as a corollary.

Corollary 1 (Existence of Floor Respecting Stable Matchings). *Floor respecting stable matchings exist.*

We note that the result stated in Corollary 1 is not trivial. As shown in Section 4.2, condition (iv) of Definition 1 is crucial in guaranteeing the existence of floor respecting stable matchings.[15]

6.2 *Discussion on stability*

In this section, we provide intuition for the stability of the matching $\tilde{\mu}$ produced by mechanism φ. This discussion also aids a better understanding of why we chose the SD for the second stage of the mechanism.

The first stage of the mechanism is the DA run on doctors $\{1, \ldots, \tilde{\nu}\}$ and hospitals H. Since the DA produces a stable matching in the absence of floor constraints, there is no blocking coalition that contains only doctors from $\{1, \ldots, \tilde{\nu}\}$. Thus, any potential blocking coalition must include a doctor in $\{\tilde{\nu}+1, \ldots, N\}$.

[15]In an earlier version of this paper (Akin 2020), assuming disjoint regions, we show the existence of floor respecting stable matchings under less stringent assumptions on the preferences. This result extends Corollary 1 to markets where doctors and hospitals may have outside options.

We now remark that the distributional constraints are binding in matching $\tilde{\mu}$ in the following sense. By the definition of $\tilde{\nu}$, all the doctors $\{\tilde{\nu}+1,\dots,N\}$ are necessary to meet the distributional constraints when doctors $\{1,\dots,\tilde{\nu}\}$ are assigned by the DA. That is, if a doctor $d \geq \tilde{\nu}+1$ were unassigned, then the distributional constraints could not be satisfied. For the same reason, if a doctor $d \geq \tilde{\nu}+1$ wants to be reassigned to a hospital not in $\tilde{H}(d)$, she must find another doctor who is willing to fill a vacancy at some hospital in $\tilde{H}(d)$.

At this point, the fact that the SD produces a (constrained) doctor-optimal matching is important to guarantee stability. This feature of the SD implies that a doctor in $\{\tilde{\nu}+1,\dots,N\}$ joins a blocking coalition only if her assignment is not in $\tilde{H}(d)$. However, such a reassignment would require finding a doctor $d' < d$ who is willing to fill a vacancy at some hospital in $\tilde{H}(d)$. Ultimately, the coalition would have to assign some doctor in $\{1,\dots,\tilde{\nu}\}$ to a hospital in $\tilde{H}(\tilde{\nu}+1)$. However, since every doctor $d \leq \tilde{\nu}$ is assigned through the DA, they find every hospital in $\tilde{H}(\tilde{\nu}+1)$ worse than their own assignment. It follows that no doctor $d \leq \tilde{\nu}$ would be willing to be reassigned to a hospital in $\tilde{H}(\tilde{\nu}+1)$. We conclude that there is no blocking coalition to matching $\tilde{\mu}$.

7. Strategy-proofness

7.1 *The mechanism is strategy-proof*

We now state the strategy-proofness of the mechanism.

Proposition 5 (Strategy-proofness). *Mechanism φ described by (3)–(5) is strategy-proof for doctors.*

We now provide intuition for why mechanism φ is strategy-proof. The DA and the SD are strategy-proof mechanisms (see Dubins and Freedman 1981). Hence, if the set of doctors assigned in each mechanism were determined independently of the reported preferences, then it would be straightforward that our mechanism is strategy-proof. However, in our mechanism, the number of doctors assigned through the DA in the first stage is given by $\tilde{\nu}$ in (5), which is determined by the reported preference profile.

We first remark that $\tilde{\nu}$ is determined by the preferences of doctors $d \leq \tilde{\nu}$, so no doctor $d > \tilde{\nu}$ can manipulate the cutoff priority. Furthermore, reporting true preferences is a weakly dominant strategy in the SD. Therefore, no doctor $d > \tilde{\nu}$ can obtain a more preferred assignment by misreporting her preferences. Thus, if a doctor finds a profitable misreporting, then she has a priority number $d \leq \tilde{\nu}$.

We now note that no profitable misreporting by doctors $d \leq \tilde{\nu}$ can increase the number of doctors who are assigned through the DA. This is because the outcome of the DA is less preferred by a doctor when the set of doctors in the market expands (see Crawford 1991). Hence, increasing the cutoff priority results in a weakly less preferred assignment for each doctor $d \leq \tilde{\nu}$. Thus, we conclude that only doctors $d \leq \tilde{\nu}$ may have incentives to misreport, and any profitable misreporting involves a decrease in the cutoff $\tilde{\nu}$.

The intuition for the rest of the proof is as follows. Suppose that by misreporting, doctor d receives a more preferred assignment, hospital $\tilde{h}$. We now consider a hypothetical cutoff priority found as follows. We consider the maximum of numbers $\hat{n}$ such that doctor d receives an assignment weakly more preferred than $\tilde{h}$ when the number of doctors assigned via the DA is fixed at $\hat{n}$ (and she reports truthfully). We now compare the outcome of the DA run on the first $\hat{n}$ doctors with their true preferences to the outcome of the DA with the misreport. The proof is by contradiction and consists of showing two statements. We first show that every hospital must receive a weakly larger number of doctors when we run the DA on doctors $\{1, \dots, \hat{n}\}$. Intuitively, a profitable misreport should reduce the competition in the first stage. Thus, if we were to set the cutoff priority exogenously (and, thus, doctors report truthfully), hospitals should receive more applications. At this point, we use this result and recall the definition of $\tilde{\nu}$ in (5) to conclude that there cannot be any feasible matching such that doctors $\{1, \dots, \hat{n}\}$ are assigned via the DA (and doctor $\hat{n}+1$ remains unassigned). This observation implies that the cutoff priority under truthful reporting should be smaller than $\hat{n}$, leading to a contradiction to the assumption of profitable misreporting. We conclude that doctor d should receive a weakly more preferred outcome under truthful reporting.

The intuition for the strategy-proofness of our mechanism echoes that found in the literature. For example, Goto et al. (2017) and Fragiadakis and Troyan (2017) propose iterated versions of the DA for markets under ceiling constraints. Although they study different models, the idea behind maintaining the strategy-proofness of these mechanisms is similar. The general idea is that whenever a doctor can manipulate the mechanism to change the updated constraints in her favor, the designer ensures that she will receive a (weakly) less preferred outcome. While the general ideas behind the strategy-proofness results resemble each other, we cannot apply any of their findings or methodologies to our setting.

7.2 *Discussion of strategy-proofness*

The definition of $\tilde{\nu}$ in (5) is crucial for maintaining the strategy-proofness of mechanism φ. To better understand the role of $\tilde{\nu}$, we provide an alternative definition of $\tilde{\nu}$ that would result in a mechanism that produces a floor respecting stable matching but is not strategy-proof.

Consider the alternative definition of $\tilde{\nu}$:

$$\tilde{\kappa} \triangleq \max\left\{ n \in \{0, 1, \dots, N\} : \begin{array}{l} \exists \mu \in \mathbb{M} \text{ such that for all } d \leq n, \\ \mu(d) = \tilde{m}(d; n) \end{array} \right\}.$$

That is, $\tilde{\kappa}$ is defined as the maximum number of doctors that can be assigned via the DA such that the floor constraints can be satisfied with the remaining subset of doctors. We analyze the mechanism by which doctors $\{1, \dots, \tilde{\kappa}\}$ are assigned via the DA, while doctors $\{\tilde{\kappa}+1, \dots, N\}$ are assigned via the SD.

We first note a crucial difference between $\tilde{\nu}$ and $\tilde{\kappa}$. A doctor d satisfies $d \leq \tilde{\nu}$ only if the floor constraints can be satisfied when doctors $d' < d$ are assigned via the DA and she remains unassigned. In contrast, a doctor d satisfies $d \leq \tilde{\kappa}$ if the floor constraints can be satisfied when she and doctors $d' < d$ are assigned via the DA. For this reason, $d \leq \tilde{\nu}$ implies that $d \leq \tilde{\kappa}$. Therefore, we also have that $\tilde{\nu} \leq \tilde{\kappa}$.

The mechanism with cutoff $\tilde{\kappa}$ implements a floor respecting stable matching. The proof of Proposition 4 goes through line-by-line when we use $\tilde{\kappa}$ instead of $\tilde{\nu}$. As we argued in Section 6.2, a key part of the proof uses the fact that a doctor $d > \tilde{\nu}$ cannot join a blocking coalition without finding another doctor $d' < d$ who is willing to be reassigned to some hospital in $\tilde{H}(d)$. This feature is still satisfied when we change the cutoff priority to $\tilde{\kappa}$ because $\tilde{\nu} \leq \tilde{\kappa}$.

We now explain why the two-stage mechanism with cutoff $\tilde{\kappa}$ is not strategy-proof. The mechanism determines $d \leq \tilde{\kappa}$ by verifying that the floor constraints can be satisfied by running the DA with the reported preferences of doctors in $\{1, \dots, d\}$. Only if the floor constraints cannot be satisfied, $d > \tilde{\kappa}$ and doctor d is assigned in the second stage. For this reason, the mechanism is prone to two types of manipulations. A doctor may increase $\tilde{\kappa}$ by submitting a different report, thereby assigned via the DA instead of the SD. Alternatively, a doctor may want to misreport her preferences to decrease $\tilde{\kappa}$ so that some other doctor is assigned via the SD instead of the DA (i.e., she reduces the competition in the first stage). We illustrate this point in the following example.

Example 3. There are two doctors and three hospitals. The hospitals' capacities are given by $q_{h_i} = 1$ for $i = 1, 2, 3$. The preference profile is

$$\begin{aligned} &\succ_{d_1}: h_3, h_1, h_2, \qquad &&\succ_{h_i}: d_2, d_1 \quad \text{for } i = 1, 2, 3 \\ &\succ_{d_2}: h_2, h_3, h_1. && \end{aligned}$$

There are three regions with floor constraints given by $\underline{q}_{r_1} = 1$ and $\underline{q}_{r_i} = 0$ for $i = 2, 3$. The hospitals are located across regions as $r(h_i) = \{r_i\}$ for $i = 1, 2, 3$. Suppose that doctor d_1 has a higher priority, i.e., $d_1 < d_2$.

Under our mechanism, $\tilde{\nu} = 1$. The total number of floor constraints is 1, so if doctor d_1 were to remain unassigned, there is a feasible matching (which assigns doctor d_2 to hospital h_1). Thus, $\tilde{\nu} > 0$. When only doctor d_1 is assigned via the DA, she would be assigned to hospital h_3. In this case, region r_1 lacks one doctor to satisfy its floor constraint, and if doctor d_2 were to remain unassigned, there would not be any feasible matching. Thus, $\tilde{\nu} = 1$.

Suppose that we have adopted the definition of $\tilde{\kappa}$ to determine the number of doctors in the first stage. Under truthful reporting, $\tilde{\kappa} = 1$. If all the doctors were assigned via the DA, doctor d_1 would be assigned to hospital h_3 and doctor d_2 would be assigned to hospital h_2. This matching is not feasible and, thus, $\tilde{\kappa} < 2$. When only doctor d_1 is assigned via the DA, doctor d_1 is assigned to hospital h_3 and doctor d_2 is assigned to hospital h_1. Such matching is feasible and, thus, $\tilde{\kappa} = 1$.

If doctor d_1 reports $\succ_{d_1}$, doctor d_2 profits from reporting $\succ'_{d_2}: h_3 \succ h_2 \succ h_1$. Under the preference profile $(\succ_{d_1}, \succ'_{d_2})$, $\tilde{\kappa} = 2$. When all the doctors are assigned via the DA, doctor d_1 is assigned to hospital h_1 and doctor d_2 is assigned to hospital h_3. This matching is feasible and, thus, $\tilde{\kappa} = 2$. In this case, doctor d_2 profits from misreporting by increasing the cutoff $\tilde{\kappa}$, thereby being assigned via the DA instead of the SD.

We note that if doctor d_2 reports preference $\succ'_{d_2}$, then doctor d_1 profits from reporting $\succ'_{d_1}$: $h_3 \succ h_2 \succ h_1$. Under the preference profile $(\succ'_{d_1}, \succ'_{d_2})$, $\tilde{\kappa} = 1$. When all the doctors are assigned via the DA, doctor d_1 is assigned to hospital h_2 and doctor d_2 is assigned to hospital h_3. This matching is not feasible and, thus, $\tilde{\kappa} < 2$. When only doctor d_1 is assigned via the DA, doctor d_1 is assigned to hospital h_3 and doctor d_2 is assigned to hospital h_1. This matching is feasible and, thus, $\tilde{\kappa} = 1$. In this case, doctor d_1 profits from misreporting by lowering the cutoff $\tilde{\kappa}$, thereby reducing the competition in the first stage of the mechanism. $\Diamond$

7.3 *Strategy-proofness and outside options*

If doctors find some hospitals unacceptable, then a strategy-proof (for doctors) mechanism that produces a floor respecting stable matching may not exist. Intuitively, a mechanism is required to produce a feasible individually rational matching given a reported preference profile. To do so, the mechanism fills the hospitals that reside in a region with positive floor constraints with doctors who find these hospitals acceptable. In response, a doctor may have an incentive to declare some of these hospitals unacceptable so that she is not assigned to one of them. The following example illustrates this point.

Example 4. Consider the market described in Example 1. The floor respecting stable matchings are μ_2 and μ_4. If a mechanism chooses matching μ_2 (resp. μ_4) with positive probability, then doctor d_2 (resp. d_1) profits from declaring hospital h_3 unacceptable. In this way, doctor d_2 (resp. d_1) guarantees being assigned to hospital h_1 (resp. h_2). Hence, there is no mechanism that is strategy-proof for doctors and produces a floor respecting stable matching. $\Diamond$

8. Conclusion

One foundational issue in the study of matching markets under floor constraints is to define an appropriate stability notion. The present paper addresses this challenge. We defined a stability notion—floor respecting stability—which accommodates hard floor constraints. We provided a mechanism that implements a floor respecting stable matching and is strategy-proof (for doctors). A natural direction for future work is the study of hard floor constraints when doctors have types.

Appendix

Proof of Proposition 1

Suppose that a matching $\mu \in \mathbb{M}$ is not Pareto efficient. Then there exists a matching $\mu' \in \mathbb{M}$ such that every doctor and hospital is weakly better off, with some of them being strictly better off. Hence, the grand coalition can block matching μ with matching μ'.

Proof of Proposition 2

Suppose that there exists a blocking coalition A' that blocks matching μ with matching μ'. If there is only one doctor in coalition A', then (without loss) coalition A' is a blocking pair and we obtain the result. Hence, we assume that coalition A' has at least two doctors.

Suppose that there is a doctor d' who is fired under matching μ' (i.e., $\mu'(d') = \varnothing$ and $\mu(d') \neq \varnothing$). Without loss of generality, there exists a doctor d'' who is in the coalition and is assigned to $\mu(d')$ (i.e., $\mu'(d'') \neq \mu(d'')$ and $\mu'(d'') = \mu(d')$) (otherwise, one can construct a new blocking coalition that does not fire doctor d'). As a shorthand notation, we denote

$$h' = \mu(d') \quad \text{and} \quad h'' = \mu(d'').$$

Note that $\mu'(d'') = h'$. Suppose that d' is the most preferred doctor among all doctors who are fired by hospital h', and that d'' is the least preferred doctor among all doctors who are in the coalition and assigned to h'. Furthermore, without loss of generality, suppose that $|\mu'(h')| = q_{h'}$; that is, doctor d' is fired from μ' because h' is at capacity (otherwise, one can construct a new blocking coalition that does not fire doctor d').

Now, consider blocking coalition $A'' = A' \backslash \{d''\}$ with matching μ'', which is constructed as follows. If $h'' \notin A'$, then

$$\mu''(d) = \begin{cases} \mu'(d) & \text{if } d \notin \{d', d''\} \\ \mu(d) & \text{if } d = d'' \\ \mu(d) & \text{if } d = d'. \end{cases}$$

If $h'' \in A'$, then

$$\mu''(d) = \begin{cases} \mu'(d) & \text{if } d \notin \{d', d''\} \\ \varnothing & \text{if } d = d'' \\ \mu(d) & \text{if } d = d'. \end{cases}$$

In other words, every doctor except d' and d'' is assigned the same way as in blocking matching μ'. The fired doctor d' is reassigned back to her original hospital h'. Depending on whether hospital h'' is in coalition A', doctor d'' is either reassigned back to her original hospital h'' or fired.

Matching μ'' is feasible since all hospitals are assigned weakly more doctors than under matching μ' and, by construction, matching μ'' respects hospital capacity constraints. We note that if h'' is not in coalition A', then h'' does not admit another doctor under μ', and so the hospital can hire back doctor d'' without violating its capacity.

We now show that matching μ'' improves the outcome of every agent in coalition A''. Every doctor $d \in A''$ is assigned to the same hospital under matching μ'' and μ', so every doctor $d \in A''$ is weakly better off under matching μ'' than μ. Every hospital $h \in H \setminus \{h'', h'\}$ is assigned to the same subset of doctors under matchings μ'' and μ'. Therefore, every hospital $h \in A'' \setminus \{h'', h'\}$ is weakly better off under matching μ'' than μ.

We now show that hospital h'' is better off under matching μ'' than μ whenever $h'' \in A''$ (otherwise, Definition 1 does not require that hospital h'' is better off). If $h'' \in A''$, then $h'' \in A'$ and this hospital is assigned to the same subset of doctors under matchings μ' and μ''. Since $h'' \in A'$, it must be that this hospital is better off under matching μ' than matching μ.

Finally, we show that hospital h' is weakly better off under matching μ'' than μ. Since hospital h' hires a new doctor under matching μ', conditions (i) and (iv) of Definition 1 imply that $h' \in A'$. Then we have

$$\mu'(h') \succeq_{h'} \mu(h').$$

We also note that

$$\mu''(h') = \mu'(h') \backslash \{d''\} \cup \{d'\}.$$

Now suppose that $d' \succ_{h'} d''$; then, from responsive hospital preferences, it follows that

$$\mu''(h') = \mu'(h') \backslash \{d''\} \cup \{d'\} \succ_{h'} \mu'(h') \succ_{h'} \mu(h').$$

Hence, we obtain the result.

If $d'' \succ_{h'} d'$, then hospital h' prefers every doctor who is reassigned to it under matching μ' to doctor d'. Suppose that $\mu''(h') \setminus \mu(h') \neq \varnothing$. We have that if $d'' \succ_{h'} d'$, then for all $d \in \mu''(h') \setminus \mu(h')$ and $\hat{d} \in \mu(h') \setminus \mu''(h')$, $d \succ_{h'} \hat{d}$. Since $|\mu'(h')| = q_{h'}$, we have $|\mu''(h')| = q_{h'}$. This implies that under matching μ'', the number of doctors reassigned to hospital h' is weakly larger than that of doctors fired by hospital h', i.e., $|\mu''(h') \setminus \mu(h')| \geq |\mu(h') \setminus \mu''(h')|$. Then, from responsive hospital preferences, we have that $\mu''(h') \succ_{h'} \mu(h')$. If $\mu''(h') \setminus \mu(h') = \varnothing$, then $\mu''(h') = \mu(h')$. We conclude that the assignment of hospital h' weakly improves under matching μ''. Thus, every hospital $h \in A''$ is weakly better off under matching μ'' than μ. This implies that coalition A'' with matching μ'' blocks matching μ.

To summarize, either A' has only one doctor (in which case the result is trivially satisfied) or we can construct a coalition A'' with one fewer doctor than coalition A' that blocks matching μ. If no doctor is fired under μ'', then condition (iv) of Definition 1 implies that coalition A'' proposes a matching that is a Pareto improvement for doctors. If A'' has only one doctor, we also obtain the result because this would be a blocking pair. Otherwise, we can repeat the same argument as before to obtain a new blocking coalition A''' with one fewer doctor than coalition A''. Repeating inductively, we obtain a blocking coalition that consists of one doctor and one hospital or that proposes a matching that is a Pareto improvement for doctors.

Proof of Proposition 3

Suppose that $\hat{\mu}$ defined in (8) is feasible. Given any $n \leq N-1$, suppose that doctors in $\{1, \dots, n\}$ are assigned by the DA. Since hospitals have responsive preferences, for every $n \leq N-1$, the number of doctors assigned to each hospital in matching $\tilde{m}(\cdot; n)$ is weakly lower than in matching $\tilde{m}(\cdot; N-1)$ (recall that $\tilde{m}$ is defined in (2)). This means that given any $n \leq N-1$ and assuming that the assignment of doctors $d \leq n$ is determined by the DA, it is possible to assign doctors $\{n+2, \dots, N\}$ (leaving doctor $n+1$ unassigned) in such a way that the number of doctors in each hospital is equal to that in matching $\hat{\mu}$. Furthermore, by assumption, matching $\hat{\mu}$ is feasible. Then, for every $n \leq N-1$, there is a feasible matching in which doctors $d \leq n$ are assigned through the DA and doctor $n+1$ is unassigned. Thus, we conclude that $\tilde{\nu} = N$, i.e., $\varphi(\succ) = \mathrm{DA}(\succ |D, H)$.

Proof of Proposition 4

Let $\tilde{\mu}$ be the matching produced by mechanism φ. We first prove that $\tilde{\mu}$ is a feasible matching and then show that there is no coalition blocking it.

Feasibility We first prove that there exists a feasible matching in which doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned through the DA. The case $\tilde{\nu} = 0$ is trivial because we assume that feasible matchings exist (i.e., $\mathbb{M} \neq \varnothing$). If $\tilde{\nu} > 0$, then there exists a feasible matching $\mu' \in \mathbb{M}$ such that doctors $\{1, \ldots, \tilde{\nu}-1\}$ are assigned through the DA and doctor $\tilde{\nu}$ remains unassigned (if such a matching μ' did not exist, then we would have that $\tilde{\nu} - 1$ is the minimum number satisfying the condition in (5)). We now observe that when doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned through the DA, the number of doctors at each hospital is weakly larger than when doctors $\{1, \ldots, \tilde{\nu} - 1\}$ are assigned through the DA, and it is strictly larger for one hospital. This observation follows from the fact that hospitals have responsive preferences and that there are no individual rationality constraints. Let $\tilde{h}$ be the hospital for which the number of doctors strictly increases when doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned through the DA.

We now construct a matching μ'' that is feasible, and doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned through the DA. This step ensures that $\tilde{H}(\tilde{\nu} + 1) \neq \varnothing$. Let μ'' be the resulting matching when doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned through the DA and doctors $\{\tilde{\nu} + 1, \ldots, N\}$ are assigned as follows:

(a) If there exists a doctor $\hat{d} > \tilde{\nu}$ such that $\mu'(\hat{d}) = \tilde{h}$, then doctors $\{\tilde{\nu} + 1, \ldots, N\}$ are assigned in the same way as in matching μ' (i.e., $\mu''(d) = \mu'(d)$), except for doctor $\hat{d}$, who remains unassigned (i.e., $\mu''(\hat{d}) = \varnothing$). In this case, by construction, μ'' is feasible because $|\mu''(h)| = |\mu'(h)|$ for all $h \in H$.

(b) If there does not exist a doctor $\hat{d} > \tilde{\nu}$ such that $\mu'(\hat{d}) = \tilde{h}$, then doctors $\{\tilde{\nu} + 1, \ldots, N\}$ are assigned in the same way as in matching μ'. First, note that the hospital capacities are satisfied because every hospital receives the same number of doctors as in μ', except for hospital $\tilde{h}$. However, hospital $\tilde{h}$'s capacity constraint is also respected because all of its doctors are assigned through the DA. Additionally, $|\mu''(h)| \geq |\mu'(h)|$ for all $h \in H$ and μ' is feasible, so μ'' is feasible.

Thus, we can construct a matching μ'' that is feasible, and doctors $\{1, \ldots, \tilde{\nu}\}$ are assigned through the DA. We conclude that $\tilde{H}(\tilde{\nu} + 1) \neq \varnothing$. The rest of the proof follows from the definition of $\tilde{H}(\cdot)$.

No blocking coalitions We now prove that there is no blocking coalition. Suppose that there exists a blocking coalition A proposing matching μ' that blocks matching $\tilde{\mu}$ (implemented by mechanism φ).

Step 1 We show that for every $h \in H$, $|\mu'(h) \cap \{1, \dots, \tilde{\nu}\}| \leq |\tilde{\mu}(h) \cap \{1, \dots, \tilde{\nu}\}|$. That is, every hospital is assigned to fewer doctors from set $\{1, \dots, \tilde{\nu}\}$ in matching μ'.

As a shorthand notation, we denote by F the set of hospitals that are at capacity in matching $\tilde{\mu}$ only with doctors in $\{1, \dots, \tilde{\nu}\}$:

$$F \triangleq \{h \in H : |\tilde{\mu}(h) \cap \{1, \dots, \tilde{\nu}\}| = q_h\}.$$

That is, F is the subset of hospitals that are at capacity at the end of the first stage of the mechanism.

Since there are no individual rationality constraints and doctors $d \leq \tilde{\nu}$ are assigned via the DA, doctors $d \leq \tilde{\nu}$ find every hospital with a vacancy less preferred than their assignments. In other words, for every doctor $d \leq \tilde{\nu}$ and every hospital $h \notin F$, $\tilde{\mu}(d) \succeq_d h$. By condition (iv) of Definition 1, if $\mu'(d) \neq \tilde{\mu}(d)$, then either $\mu'(d) \succ_d \tilde{\mu}(d)$ or $\mu'(d) = \varnothing$. Hence, there is no doctor $d \leq \tilde{\nu}$ who is reassigned to one of the hospitals with vacant positions. In other words, for every doctor $d \leq \tilde{\nu}$ such that $\mu'(d) \neq \tilde{\mu}(d)$, either $\mu'(d) \in F$ or $\mu'(d) = \varnothing$. Therefore, for every $h \in H$, $|\mu'(h) \cap \{1, \dots, \tilde{\nu}\}| \leq |\tilde{\mu}(h) \cap \{1, \dots, \tilde{\nu}\}|$. This concludes Step 1.

Step 2 We now show that for every doctor $d > \tilde{\nu}$, $\mu'(d) = \tilde{\mu}(d)$. In other words, for every doctor $d > \tilde{\nu}$, $d \notin A$. We begin by proving that $\mu'(\tilde{\nu}+1) = \tilde{\mu}(\tilde{\nu}+1)$ and do so by contradiction. That is, we assume that $\mu'(\tilde{\nu}+1) \neq \tilde{\mu}(\tilde{\nu}+1)$. By condition (iv) of Definition 1, if $\mu'(\tilde{\nu}+1) \neq \tilde{\mu}(\tilde{\nu}+1)$, then $\mu'(\tilde{\nu}+1) \succ_{\tilde{\nu}+1} \tilde{\mu}(\tilde{\nu}+1)$ or $\mu'(\tilde{\nu}+1) = \varnothing$, which implies that $\mu'(\tilde{\nu}+1) \notin \tilde{H}(\tilde{\nu}+1)$ or $\mu'(\tilde{\nu}+1) = \varnothing$. We now reach a contradiction.

Consider matching $\tilde{\mu}'$ constructed inductively as

$$\tilde{\mu}'(d) = \begin{cases} \tilde{\mu}(d) & \text{if } d \leq \tilde{\nu} \\ \mu'(h) & \text{if } d > \tilde{\nu} \text{ and } |\tilde{\mu}'(h) \cap \{1, \dots, d-1\}| < q_h \text{ for } h = \mu'(d) \\ \varnothing & \text{otherwise.} \end{cases}$$

In other words, $\tilde{\mu}'$ is constructed such that doctors $d \leq \tilde{\nu}$ are assigned in the same way as in matching $\tilde{\mu}$ and doctors $d > \tilde{\nu}$ are assigned sequentially in the same way as in μ' as long as this does not violate hospital capacities. By construction, matching $\tilde{\mu}'$ respects hospital capacities. In Step 1, we showed that for every $h \in H$, $|\mu'(h) \cap \{1, \dots, \tilde{\nu}\}| \leq |\tilde{\mu}(h) \cap \{1, \dots, \tilde{\nu}\}|$. So, for every $h \in H$, $|\mu'(h) \cap \{1, \dots, \tilde{\nu}\}| \leq |\tilde{\mu}'(h) \cap \{1, \dots, \tilde{\nu}\}|$ as well. Finally, doctors $d > \tilde{\nu}$ are assigned sequentially in the same way as in μ' as long as doing so would not violate hospital capacities. Hence, $|\mu'(h)| \leq |\tilde{\mu}'(h)|$ for every $h \in H$. Furthermore, since $\mu' \in \mathbb{M}$ and $|\mu'(h)| \leq |\tilde{\mu}'(h)|$ for every $h \in H$ (and $\tilde{\mu}'$ respects hospital capacities), we have that $\tilde{\mu}' \in \mathbb{M}$. We now observe that $\tilde{\mu}'(\tilde{\nu}+1) = \mu'(\tilde{\nu}+1)$ or $\tilde{\mu}'(\tilde{\nu}+1) = \varnothing$, which implies that $\tilde{\mu}'(\tilde{\nu}+1) \notin \tilde{H}(\tilde{\nu}+1)$ or $\tilde{\mu}'(\tilde{\nu}+1) = \varnothing$. However, $\tilde{H}(\tilde{\nu}+1)$ is constructed such that if $\tilde{\mu}'(d) = \tilde{m}(d; \tilde{\nu})$ for all doctors $d \leq \tilde{\nu}$ and $\tilde{\mu}'(\tilde{\nu}+1) \notin \tilde{H}(\tilde{\nu}+1)$, then $\tilde{\mu}' \notin \mathbb{M}$. If $\tilde{\mu}'(\tilde{\nu}+1) = \varnothing$, then from definition of $\tilde{\nu}$, we have that $\tilde{\mu}' \notin \mathbb{M}$. Therefore, we reach a contradiction and conclude that $\mu'(\tilde{\nu}+1) = \tilde{\mu}(\tilde{\nu}+1)$.

Since $\mu'(\tilde{\nu}+1) = \tilde{\mu}(\tilde{\nu}+1)$, we can now repeat the same arguments for doctor $\tilde{\nu}+2$ and continue inductively to prove that for all $d > \tilde{\nu}$, $\mu'(d) = \tilde{\mu}(d)$. This concludes the second step of the proof.

The rest of the proof We now note that mechanism φ determines the assignment of doctors $d \le \tilde{\nu}$ by running the DA without including the rest of the doctors. Since the DA produces (unconstrained) stable matchings, there is no coalition that satisfies the first three conditions in Definition 1. Hence, there does not exist a blocking coalition $A \subseteq \{1, \dots, \tilde{\nu}\} \cup H$. Then there exists a doctor $d \in \{\tilde{\nu}+1, \dots, N\}$ such that $d \in A$. We reach a contradiction because in Step 2, we proved that $d \notin A$. We conclude that there is no coalition blocking matching $\tilde{\mu}$.

Proof of Proposition 5

Fix an ordering of doctors $D = \{1, \dots, N\}$. Suppose that mechanism φ is not strategy-proof for doctors. Then there is a preference profile $\succ$ and a doctor $d^* \in D$ who can profitably misreport. Throughout the proof, $\tilde{\succ}_{d^*}$ refers to the profitable misreport and $\succ_{d^*}$ refers to the true preferences of doctor d^*. By definition, $\tilde{\succ}_{d^*}$ is a profitable misreport if

$$\varphi(\tilde{\succ}_{d^*}, \succ_{-d^*})(d^*) \succ_{d^*} \varphi(\succ_D)(d^*).$$

We denote by $\tilde{\succ}$ the preference profile when doctor d^* reports $\tilde{\succ}_{d^*}$ and others report their true preferences, i.e.,

$$\tilde{\succ} \triangleq (\tilde{\succ}_{d^*}, \succ_{-d^*}).$$

Let ν (resp. $\tilde{\nu}$) be the number of doctors assigned in the first stage when the submitted preference profile is $\succ$ (resp. $\tilde{\succ}$). We note that the assignments of the doctors in the first stage do not change with the preferences reported by doctors $d > \nu$, and doctors $d > \nu$ are assigned by the SD, which is strategy-proof. Thus, $d^* \le \nu$. We now show that $d^* \le \tilde{\nu}$. From the definition of ν, it follows that for any $n < d^*$, there is a feasible matching in which the assignments of doctors $\{1, \dots, n\}$ are determined via the DA (and doctor $n+1$ remains unassigned). Thus, such a feasible matching exists for every $n < d^*$ when doctor d^* misreports as well. Thus, $\tilde{\nu} \ge d^*$.

We now note that $\tilde{\nu} < \nu$. By responsive hospital preferences, doctor d^* is worse off when the DA is run with more doctors (see Crawford 1991) and she cannot profitably misreport if the number of doctors assigned by DA is fixed to be ν (see Dubins and Freedman 1981). Thus, a profitable misreporting by doctor d^* implies that $\tilde{\nu} < \nu$.

Throughout the proof, to keep the notation compact, we let

$$\mathrm{DA}(\succ; n) \triangleq \mathrm{DA}\big(\succ |\{1, \dots, n\}, H\big).$$

We now denote by $\hat{n}$ the maximum number of doctors such that doctor d^* prefers the matching produced by the DA when she reports $\succ_{d^*}$ to the matching produced by the DA when she reports $\tilde{\succ}_{d^*}$. That is,

$$\hat{n} = \max\big\{n \in \{1, \dots, N-1\} : \mathrm{DA}(\succ; n)\big(d^*\big) \succeq_{d^*} \mathrm{DA}(\tilde{\succ}; \tilde{\nu})\big(d^*\big) \succ_{d^*} \mathrm{DA}(\succ; n+1)\big(d^*\big)\big\}. \quad (9)$$

The existence of $\hat{n}$ is proved as follows. Since the DA is strategy-proof and we assume that $\tilde{\succ}_{d^*}$ is a profitable misreporting, it follows that

$$\mathrm{DA}(\succ; \tilde{\nu})\big(d^*\big) \succeq_{d^*} \mathrm{DA}(\tilde{\succ}; \tilde{\nu})\big(d^*\big) \succ_{d^*} \mathrm{DA}(\succ, \nu)\big(d^*\big).$$

In other words, reporting truthfully yields a more preferred hospital than reporting $\tilde{\succsim}_{d^*}$ when the number of doctors in the first stage is fixed to be $\tilde{\nu}$ (see Dubins and Freedman 1981). Since $\tilde{\succsim}_{d^*}$ is a profitable misreport, $\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(d^*) \succ_{d^*} \mathrm{DA}(\succ; \nu)(d^*)$. Thus, the largest number $n \in \{\tilde{\nu}, \dots, \nu - 1\}$ such that $\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(d^*) \succ_{d^*} \mathrm{DA}(\succ, n+1)(d^*)$ satisfies the condition in (9), which implies the existence of $\hat{n}$. Moreover, it is clear that

$$\tilde{\nu} \leq \hat{n} < \nu.$$

We now reach a contradiction by proving the following two claims.

Claim 1. *There exists $h \in H$ such that $|\mathrm{DA}(\succ; \hat{n})(h)| < |\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(h)|$.*

Claim 2. *For all $h \in H$, $|\mathrm{DA}(\succ; \hat{n})(h)| \geq |\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(h)|$.*

Proof of Claim 1. We denote by $\tilde{\mu}$ the matching implemented by mechanism φ under preference profile $\tilde{\succsim}$. We prove the result by contradiction, so we suppose that $\forall h \in H$, $|\mathrm{DA}(\succ; \hat{n})(h)| \geq |\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(h)|$. We reach a contradiction by showing that there does not exist $\mu \in \mathbb{M}$ such that $\mu(d) = \mathrm{DA}(\succ; \hat{n})(d)$ for all $d \leq \hat{n}$ and $\mu(\hat{n}+1) = \varnothing$. Since there does not exist $\mu \in \mathbb{M}$ such that for all $d \leq \hat{n}$, $\mu(d) = \mathrm{DA}(\succ; \hat{n})(h)$ and $\mu(\hat{n}+1) = \varnothing$, this implies that $\nu \leq \hat{n}$. Therefore, we reach a contradiction and conclude that $\exists h \in H$ such that $|\mathrm{DA}(\succ; \hat{n})(h)| < |\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(h)|$.

Suppose that there exists $\mu \in \mathbb{M}$ such that $\mu(d) = \mathrm{DA}(\succ; \hat{n})(d)$ for all $d \leq \hat{n}$ and $\mu(\hat{n}+1) = \varnothing$. Let μ' be defined as follows: (a) for all $d \leq \tilde{\nu}$, $\mu'(d) = \tilde{\mu}(d)$, (b) for $d = \tilde{\nu}+1$, $\mu'(d) = \varnothing$, (c) for all doctors $d > \hat{n}+1$, $\mu'(d) = \mu(d)$, and (d) all doctors $d \in \{\tilde{\nu}+2, \dots, \hat{n}+1\}$ are assigned such that $|\mu'(h) \cap \{1, \dots, \hat{n}+1\}| \geq |\mu(h) \cap \{1, \dots, \hat{n}+1\}|$ for all $h \in H$, and μ' respects hospital capacities.[16] In other words, μ' is the same as $\tilde{\mu}$ for doctors $d \leq \tilde{\nu}$, doctor $\tilde{\nu}+1$ remains unassigned (as opposed to doctor $\hat{n}+1$), doctors $d > \hat{n}+1$ are assigned in the same way as in μ, and doctors $d \in \{\tilde{\nu}+2, \dots, \hat{n}+1\}$ are assigned such that all hospitals have weakly more doctors from subset $\{1, \dots, \hat{n}+1\}$ under μ' than under μ.

We show that doctors in $\{\tilde{\nu}+2, \dots, \hat{n}+1\}$ can be assigned in a way that matching μ' satisfies condition (d) stated above. Since $\forall h \in H$, $|\mathrm{DA}(\succ; \hat{n})(h)| \geq |\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(h)|$ and doctor $\hat{n}+1$ remains unassigned under μ, it follows that

$$\sum_{h \in H} \bigl|\mu(h) \cap \{1, \dots, \hat{n}+1\}\bigr| - \bigl|\tilde{\mu}(h) \cap \{1, \dots, \tilde{\nu}\}\bigr| \leq \hat{n} - \tilde{\nu}.$$

In other words, doctors $\{\tilde{\nu}+2, \dots, \hat{n}+1\}$ are sufficient to fill in the necessary positions such that $|\mu'(h) \cap \{1, \dots, \hat{n}+1\}| \geq |\mu(h) \cap \{1, \dots, \hat{n}+1\}|$ for all $h \in H$. Moreover, we can do the assignment of these doctors such that hospital capacities are respected as well, because matching μ respects hospital capacities.

We note that $\mu' \in \mathbb{M}$. This follows from the fact that $\mu \in \mathbb{M}$ and $q_h \geq |\mu'(h)| \geq |\mu(h)|$ for every $h \in H$. Moreover, $\mu'(\tilde{\nu}+1) = \varnothing$ and $\mu'(d) = \tilde{\mu}(d)$ for all $d \leq \tilde{\nu}$. However, the definition of $\tilde{\nu}$ implies that there does not exist $\mu' \in \mathbb{M}$ such that for all $d \leq \tilde{\nu}$, $\mu'(d) = \tilde{\mu}(d)$ and $\mu'(\tilde{\nu}+1) = \varnothing$. Hence, we reach a contradiction. We conclude that there does not exist $\mu \in \mathbb{M}$ such that $\mu(d) = \mathrm{DA}(\succ; \hat{n})(d)$ for all $d \leq \hat{n}$ and $\mu(\hat{n}+1) = \varnothing$. □

[16] We note that when $\hat{n} = \tilde{\nu}$, $\{\tilde{\nu}+2, \dots, \hat{n}+1\} = \varnothing$. However, in this case, we have $\forall h \in H$, $|\mathrm{DA}(\succ; \hat{n})(h)| = |\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})(h)|$ and, thus, $|\mu'(h) \cap \{1, \dots, \hat{n}+1\}| = |\mu(h) \cap \{1, \dots, \hat{n}+1\}|$ for all $h \in H$.

Proof of Claim 2. Let $\tilde{h} = \text{DA}(\tilde{\succsim}; \tilde{\nu})(d^*)$ and $\tilde{\succsim}'_{d^*}$ be the preference that is equal to $\tilde{\succsim}_{d^*}$ truncated at hospital $\tilde{h}$. In other words, $\tilde{h}$ is the hospital assigned to doctor d^* under preference profile $\tilde{\succsim}$, and every hospital less preferred than $\tilde{h}$ becomes unacceptable under $\tilde{\succsim}'_{d^*}$. We also denote $\tilde{\succsim}' = (\tilde{\succsim}'_{d^*}, \succ_{-d^*})$. Let $\succ'_{d^*}$ be doctor d^*'s true preference truncated at hospital $\text{DA}(\succ; \hat{n})(d^*)$. That is, every hospital less preferred than $\text{DA}(\succ; \hat{n})(d^*)$ under $\succ_{d^*}$ becomes unacceptable under $\succ'_{d^*}$. We denote $\succ' = (\succ'_{d^*}, \succ_{-d^*})$.

We now note that for every $h \in H$,

$$\left|\text{DA}(\succ; \hat{n})(h)\right| = \left|\text{DA}\big(\succ'; \hat{n}\big)(h)\right| = \left|\text{DA}\big(\succ'; \hat{n}+1\big)(h)\right|. \tag{10}$$

The first equality follows from the fact that $\succ'_{d^*}$ is equal to $\succ_{d^*}$ truncated at $\text{DA}(\succ; \hat{n})(d^*)$, so the truncation will not change the matching produced by the DA. To see why the second equality holds, note that $\text{DA}(\succ; \hat{n})(d^*) \neq \text{DA}(\succ; \hat{n}+1)(d^*)$ (which follows from the construction of $\hat{n}$), and matching $\text{DA}(\succ'; \hat{n}+1)$ can be obtained by considering matching $\text{DA}(\succ'; \hat{n})$ and adding doctor $\hat{n}+1$ (see McVitie and Wilson 1971). Adding doctor $\hat{n}+1$ to matching $\text{DA}(\succ'; \hat{n})$ creates a rejection chain that ends when d^* is rejected by hospital $\text{DA}(\succ'; \hat{n})(d)$ and remains unassigned because she finds (according to $\succ'_{d^*}$) every remaining hospital unacceptable. This proves that the total number of doctors assigned in matching $\text{DA}(\succ'; \hat{n}+1)$ is the same as in matching $\text{DA}(\succ'; \hat{n})$. By responsive hospital preferences, for every $h \in H$, $|\text{DA}(\succ'; \hat{n}+1)(h)| \geq |\text{DA}(\succ'; \hat{n})(h)|$; thus, the second equality follows.

We now prove that for every $h \in H$,

$$\left|\text{DA}\big(\succ'; \hat{n}+1\big)(h)\right| = \left|\text{DA}\big(\tilde{\succsim}'; \hat{n}+1\big)(h)\right|. \tag{11}$$

To prove (11), we first show that $\text{DA}(\tilde{\succsim}'; \hat{n}+1)(d^*) = \{\varnothing\}$. By construction of $\hat{n}$,

$$\text{DA}(\tilde{\succsim}; \tilde{\nu})\big(d^*\big) \succ_{d^*} \text{DA}(\succ; \hat{n}+1)\big(d^*\big).$$

Therefore, hospital $\tilde{h}$ is not available to d^* when there are $\hat{n}+1$ doctors in the first stage of the mechanism and she reports $\succ_{d^*}$. Since DA is strategy-proof for markets with a fixed number of doctors, hospital $\tilde{h}$ is not available to doctor d^* as well when she reports $\tilde{\succsim}'_{d^*}$ and there are $\hat{n}+1$ doctors, and, thus, she remains unassigned. This proves that $\text{DA}(\tilde{\succsim}'; \hat{n}+1)(d^*) = \varnothing$. Recall that we previously showed that $\text{DA}(\succ'; \hat{n}+1)(d^*) = \varnothing$. Thus, the total number of doctors assigned in both matchings $\text{DA}(\succ'; \hat{n}+1)$ and $\text{DA}(\tilde{\succsim}', \hat{n}+1)$ is $\hat{n}$.

Last, we note that matchings $\text{DA}(\succ'; \{1, \dots, \hat{n}+1\} \setminus \{d^*\})$ and $\text{DA}(\tilde{\succsim}'; \{1, \dots, \hat{n}+1\} \setminus \{d^*\})$ are the same and the total number of doctors assigned is $\hat{n}$. When we add doctor d^* to both of these markets, the number of doctors assigned to each hospital weakly increases by responsive hospital preferences. Since the total number of doctors assigned in both markets remains at $\hat{n}$ when we add doctor d^*, we conclude that the number of doctors assigned to each hospital remains the same, i.e., $|\text{DA}(\succ'; \hat{n}+1)(h)| = |\text{DA}(\tilde{\succsim}'; \hat{n}+1)(h)|$ for every $h \in H$. This concludes the proof of (11).

We now note that for every $h \in H$,

$$\left|\text{DA}\big(\tilde{\succsim}'; \hat{n}+1\big)(h)\right| \geq \left|\text{DA}\big(\tilde{\succsim}'; \tilde{\nu}\big)(h)\right| = \left|\text{DA}(\tilde{\succsim}; \tilde{\nu})(h)\right|. \tag{12}$$

The first inequality follows from the fact that $\tilde{\nu} < \hat{n} + 1$ and hospital preferences are responsive. The last equality follows from the fact that doctor d^* is assigned to $\tilde{h}$ in matching $\mathrm{DA}(\tilde{\succsim}; \tilde{\nu})$ (by the definition of $\tilde{h}$). In matching $\mathrm{DA}(\tilde{\succsim}'; \tilde{\nu})$, she will also be assigned to $\tilde{h}$ because doctor d^*'s preference for hospitals preferred less than $\tilde{h}$ does not affect the outcome of the DA. Thus, matchings $\mathrm{DA}(\tilde{\succsim}'; \tilde{\nu})$ and $\mathrm{DA}(\tilde{\succsim}'; \tilde{\nu})$ are the same.

By combining (10), (11), and (12), we obtain the desired result. □

Claim 1 and Claim 2 are clearly a contradiction, which implies that doctor d^* does not have profitable misreporting. This concludes the proof for the strategy-proofness of mechanism φ.

References

Abdulkadiroğlu, Atila (2005), "College admissions with affirmative action." *International Journal of Game Theory*, 33, 535–549. [915]

Abdulkadiroğlu, Atila and Tayfun Sönmez (2003), "School choice: A mechanism design approach." *American Economic Review*, 93, 729–747. [915]

Akin, Sumeyra (2020), "Matching with floor constraints." Unpublished paper, Pontifical Catholic University of Chile - Institute of Economics, SSRN 3713005. [928]

Alcoba, Natalie (2009), "Hitting the road to alleviate India's rural doctor shortage." *Canadian Medical Association Journal*, 180, E34–E36. [912]

Aygün, Orhan and Bertan Turhan (2020), "Dynamic reserves in matching markets." *Journal of Economic Theory*, 188, 105069. [915]

Biró, Peter, Tamás Fleiner, Robert W. Irving, and David F. Manlove (2010), "The college admissions problem with lower and common quotas." *Theoretical Computer Science*, 411, 3136–3153. [915]

Crawford, Vincent P. (1991), "Comparative statics in matching markets." *Journal of Economic Theory*, 54, 389–400. [929, 937]

Dubins, Lester E. and David A. Freedman (1981), "Machiavelli and the Gale–Shapley algorithm." *American Mathematical Monthly*, 88, 485–494. [929, 937, 938]

Dur, Umut, 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, 2457–2479. [915]

Echenique, Federico and M. Bumin Yenmez (2015), "How to control controlled school choice." *American Economic Review*, 105, 2679–2694. [912, 915]

Ehlers, Lars (2010), "School choice with control." Working Paper, CIREQ Cahier 13-2010, Université de Montréal. [914, 920]

Ehlers, Lars (2014), "Top trading with fixed tie-breaking in markets with indivisible goods." *Journal of Economic Theory*, 151, 64–87. [914]

Ehlers, Lars, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim (2014), "School choice with controlled choice constraints: Hard bounds versus soft bounds." *Journal of Economic Theory*, 153, 648–683. [913, 914, 920]

Erdil, Aytek and Taro Kumano (2013), "Prioritizing diversity in school choice." Unpublished paper, University of Cambridge, Faculty of Economics. [915]

Fragiadakis, Daniel, Atsushi Iwasaki, Peter Troyan, Suguru Ueda, and Makoto Yokoo (2016), "Strategyproof matching with minimum quotas." *ACM Transactions on Economics and Computation*, 4, 6. [914]

Fragiadakis, Daniel and Peter Troyan (2017), "Improving matching under hard distributional constraints." *Theoretical Economics*, 12, 863–908. [912, 915, 930]

Gale, David and S. Shapley Lloyd (1962), "College admissions and the stability of marriage." *American Mathematical Monthly*, 69, 9–15. [919, 922, 926]

Goto, Masahiro, Naoyuki Hashimoto, Atsushi Iwasaki, Yujiro Kawasaki, Suguru Ueda, Yosuke Yasuda, and Makoto Yokoo (2014), "Strategy-proof matching with regional minimum quotas." In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems* (Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns, eds.), 1225–1232, International Foundation for Autonomous Agents and Multiagent Systems. [915]

Goto, Masahiro, Fuhito Kojima, Ryoji Kurata, Akihisa Tamura, and Makoto Yokoo (2017), "Designing matching mechanisms under general distributional constraints." *American Economic Journal: Microeconomics*, 9, 226–262. [914, 915, 930]

Hafalir, Isa E., M. Bumin Yenmez, and Muhammed A. Yildirim (2013), "Effective affirmative action in school choice." *Theoretical Economics*, 8, 325–363. [912, 915]

Huang, Chien-Chung (2010), "Classified stable matching." In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, 1235–1253, SIAM. [915]

Kamada, Yuichiro and Fuhito Kojima (2015), "Efficient matching under distributional constraints: Theory and applications." *American Economic Review*, 105, 67–99. [914, 919]

Kamada, Yuichiro and Fuhito Kojima (2017a), "Recent developments in matching with constraints." *American Economic Review*, 107, 200–204. [914]

Kamada, Yuichiro and Fuhito Kojima (2017b), "Stability concepts in matching under distributional constraints." *Journal of Economic Theory*, 168, 107–142. [914]

Kamada, Yuichiro and Fuhito Kojima (2018), "Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition." *Theoretical Economics*, 13, 761–793. [914]

McVitie, David G. and Leslie B. Wilson (1971), "The stable marriage problem." *Communications of the ACM*, 14, 486–490. [939]

Nambiar, Meera and Josh Bavas (2010), "Rudd plan won't fix rural doctor shortage: RDA." ABC News, URL https://www.abc.net.au/news/2010-03-16/rudd-plan-wont-fix-rural-doctor-shortage-rda/365524. [912]

Roth, Alvin E. (1984), "The evolution of the labor market for medical interns and residents: A case study in game theory." *Journal of Political Economy*, 92, 991–1016. [912]

Roth, Alvin E. (1985a), "The college admissions problem is not equivalent to the marriage problem." *Journal of Economic Theory*, 36, 277–288. [915]

Roth, Alvin E. (1985b), "Common and conflicting interests in two-sided matching markets." *European Economic Review*, 27, 75–96. [918, 919]

Talbott, Chris (2007), "Shortage of doctors affects rural US." The Oklahoman, URL https://www.oklahoman.com/article/3088024/shortage-of-doctors-affects-rural-us. [912]

Terrier, Camille (2014), "Matching practices for secondary public school teachers—France." Matching in Practice, URL https://www.matching-in-practice.eu/wp-content/uploads/2014/06/MiP_-Profile_No.20-1.pdf. [912]

Tomoeda, Kentaro (2018), "Finding a stable matching under type-specific minimum quotas." *Journal of Economic Theory*, 176, 81–117. [914]

Co-editor Thomas Mariotti handled this manuscript.

Manuscript received 29 May, 2019; final version accepted 28 October, 2020; available online 3 November, 2020.