

EFFICIENT MATCHING UNDER DISTRIBUTIONAL CONSTRAINTS: THEORY AND APPLICATIONS

YUICHIRO KAMADA AND FUHITO KOJIMA

ABSTRACT. Many real matching markets are subject to distributional constraints. These constraints often take the form of restrictions on the numbers of agents on one side of the market matched to certain subsets of the other side. Real-life examples include restrictions imposed on regions in medical residency matching, academic master's programs in graduate school admission, and state-financed seats for college admission. Motivated by these markets, we study the design of matching mechanism under distributional constraints. We show that the existing matching mechanisms around the world may result in avoidable inefficiency and instability, and propose a better mechanism that has desirable properties in terms of efficiency, stability, and incentives while respecting the distributional constraints.

JEL Classification Numbers: C70, D61, D63.

Keywords: medical residency matching, college and graduate school admission, distributional constraints, efficiency, stability, strategy-proofness, matching with contracts, the rural hospital theorem

Date: July 30, 2014.

Kamada: Haas School of Business, University of California, Berkeley, Berkeley, CA 94720, y.cam.24@gmail.com. Kojima: Department of Economics, Stanford University, Stanford, CA 94305, fkojima@stanford.edu. We are grateful to Mustafa Oguz Afacan, Péter Biró, Erich Budish, Sylvain Chassang, Vincent Crawford, Hisao Endo, Clayton Featherstone, Tamás Fleiner, Drew Fudenberg, Tadashi Hashimoto, John William Hatfield, Toshiaki Iizuka, Rob Irving, Ryo Jinnai, Onur Kesten, Scott Duke Kominers, Hideo Konishi, Mihai Manea, David Manlove, Taisuke Matsubae, Aki Matsui, Yusuke Narita, Muriel Niederle, Parag Pathak, Al Roth, Dan Sasaki, Tayfun Sönmez, Satoru Takahashi, William Thomson, Alexis Akira Toda, Kentaro Tomoeda, Utku Ünver, Jun Wako, Alex Westkamp, Yosuke Yasuda, Bumin Yenmez, and participants at numerous seminars and conferences for helpful comments. Doctors Keisuke Izumi, Yoshiaki Kanno, Masataka Kawana, and Masaaki Nagano answered our questions about medical residency in Japan and introduced us to the relevant medical literature. We are grateful to officials at the Ministry of Health, Labor and Welfare and the Japan Residency Matching Program for discussion. Jin Chen, Irene Hsu, Seung Hoon Lee, Stephen Nei, Bobak Pakzad-Hurson, Neil Prasad, Fanqi Shi, Pete Troyan, and Rui Yu provided excellent research assistance. We thank Daniel Fragiadakis, Jiajia Gao, Mar Reguant, and Kun Wang for their help in simulation. Kojima gratefully acknowledges financial support from the Sloan Foundation, as well as financial support from the National Research Foundation through its Global Research Network Grant (NRF- 2013S1A2A2035408).

1. INTRODUCTION

Many real matching markets are subject to distributional constraints. In health care, there is often a concern that certain medical specialties attract too many doctors while others suffer from shortage, and regulations on the number of doctors in certain specialties are imposed or proposed.¹ In some public school districts, multiple school programs often share one school building, so there is a bound on the total number of students in these programs in addition to each program’s capacity because of the building’s physical size.² Achieving demographic balance of the incoming class is often one of the most important goals in school and college admission, which may lead to distributional constraints on the matching outcome.

An interesting example of a concrete matching market with such distributional constraints is the one for Japanese medical residency in which around 8,000 doctors (mostly consisting of graduating medical students) are matched to about 1,500 residency programs each year. In 2008, the Japanese government introduced a “regional cap” which, for each of the 47 prefectures that partition the country, restricts the total number of residents matched within the prefecture. This measure was taken to regulate the geographical distribution of doctors, which would otherwise be concentrated too much in urban areas at the expense of rural areas.³ Since the introduction of the regional caps, they have used a mechanism that respects the caps, which is a modification of the standard deferred acceptance mechanism that they used before 2008. Specifically, in the modified mechanism, which we call the Japan Residency Matching Program (JRMP) mechanism, if the sum of the hospital capacities in a region exceeds its regional cap, then the capacity of each

¹In the United States, the plaintiff of *Jung v. Association of American Medical Colleges* alleged that the association called Accreditation Council for Graduate Medical Education regulates the number of medical residents in each specialty, although the allegation was dismissed. Nicholson (2003) finds some suggesting evidence that Residency Review Committees, which works closely with ACGME, essentially has complete control over the number of residents who train in each specialty.

²See Abdulkadiroğlu and Sönmez (2003) for the introduction to school choice problems.

³Regional imbalance of doctors is a serious concern in many other parts of the world as well. For instance, a Washington Post article entitled “Shortage of Doctors Affects Rural U.S.” describes a dire situation in the United States (Talbot, 2007) as follows: “The government estimates that more than 35 million Americans live in underserved areas, and it would take 16,000 doctors to immediately fill that need, according to the American Medical Association.” Similar problems are present around the world. For example, one can easily find reports of doctor shortages in rural areas in the United Kingdom, India, Australia, and Thailand (see Shallcross (2005), Alcoba (2009), Nambiar and Bavas (2010), and Wongruang (2010)).

hospital is reduced to equalize the total capacity with the regional cap. Then the deferred acceptance algorithm is implemented under the reduced capacities.⁴

A similar policy is taken in the context of graduate school admission in China, which has placed more than 400,000 students every year since 2009, where master’s programs are categorized as either academic or professional. To enlarge the labor force with professional master’s degrees, in 2010 the Chinese government started restricting the number of admissions to academic master’s programs. More specifically, the government decided to reduce the number of available seats of each academic master’s program by about 25 percent by 2015, just as in Japan where a rigid restriction is imposed on each hospital. A clearinghouse mechanism is run given the reduced capacities.

In the context of college admission, the Ukrainian government provides a limited number of “state-financed” seats in public universities, for which tuition is free for the students. The number of state-financed seats and that of privately financed seats in each college are determined first, and then a matching procedure assigns students to these college-specific state-financed or privately financed seats. The situation resembles the market for Japanese medical residents if the total number of state-financed seats is interpreted as a regional cap.

A different type of mechanism is used in the U.K. medical match in which about 7,000 new doctors participate. Their matching procedure has two rounds, in which a doctor is assigned to one of the country’s 25 regions first, and then to a program within their assigned region. A similar two-round procedure is employed in the matching of new teachers in Scotland as well.

An interesting observation is that distributional constraints appear in many different contexts (we discuss more examples in the Appendix), and many different policies have been tried to accommodate these constraints, but it has been unclear which mechanism in practice, if any, achieves appealing properties such as efficiency, stability, and incentive compatibility under such constraints.

Motivated by these real-life policies, we study the design of matching markets under constraints on the doctor distribution. This paper shows that each of those existing mechanisms may result in avoidable instability and inefficiency. We further find that

⁴ The capacity of a hospital is reduced proportionately to its original capacity in principle (subject to integrality constraints) although there are a number of fine adjustments and exceptions. This rule might suggest that hospitals have incentives to misreport their true capacities, but in Japan, the government regulates how many positions each hospital can offer so that the capacity can be considered exogenous. More specifically, the government decides the physical capacity of a hospital based on verifiable information such as the number of beds in it.

a mechanism that appears to be intuitively appealing (and is similar to a mechanism that policy makers often informally suggest to us) suffers from incentive problems. We propose an alternative mechanism that overcomes these shortcomings while respecting the distributional goals. More specifically, we first introduce a stability concept and (constrained) efficiency that take distributional constraints into account. We point out that none of the aforementioned existing mechanisms always produces a stable or efficient matching, and present a new mechanism that we call the flexible deferred acceptance mechanism. We show that, unlike other mechanisms, this mechanism finds a stable and efficient matching and is (group) strategy-proof for doctors.⁵ In addition, the flexible deferred acceptance mechanism matches weakly more doctors to hospitals (in the sense of set inclusion) and makes every doctor weakly better off than the JRMP mechanism. These results suggest that replacing the current mechanisms with the flexible deferred acceptance mechanism will improve the performance of the matching market.

We also find that the structural properties of the stable matchings with distributional constraints are strikingly different from those in the standard matching models. First, there does not necessarily exist a doctor-optimal stable matching (a stable matching unanimously preferred to every stable matching by all doctors). Neither do there exist hospital-optimal nor doctor-pessimal nor hospital-pessimal stable matchings. Second, different stable matchings can leave different hospitals with unfilled positions, implying that the conclusion of the rural hospital theorem fails in our context. Based on these observations, we investigate whether the government can design a reasonable mechanism that selects a particular stable matching based on its policy goals such as minimizing the number of unmatched doctors.

Let us emphasize that analyzing abstract technical issues associated with distributional constraints is *not* the primary purpose of this paper. On the contrary, we study a model motivated by various real markets and offer practical solutions for these markets. Improving the existing markets is important by itself, such as the Japanese residency markets which produces around 8,000 medical doctors, the Chinese graduate admission which admits more than 400,000 new master students, and the UK medical match which involves around 7,000 doctors. Moreover, this paper tries to provide a general framework in which one can tackle problems arising in practical markets which have yet to be recognized or

⁵A mechanism being (group) strategy-proof for doctors means that telling the truth is a dominant strategy for each doctor (and even a coalition of doctors cannot jointly misreport preferences and benefit).

addressed. In these senses, our paper contributes to the general research agenda of market design, advocated by Roth (2002) for instance, that emphasizes the importance of addressing issues arising in practical allocation problems.

The rest of this paper proceeds as follows. In Section 2, we present the model of matching with regional caps. In Section 3 where we define the JRMP mechanism, we show that it does not necessarily produce an efficient matching and there is a sense in which the produced matching is not stable. Section 4 introduces and analyzes a stability concept under distributional constraints. In Section 5 we propose a new mechanism, the flexible deferred acceptance mechanism, and show that it produces a stable and efficient matching and is group strategy-proof. Section 6 discusses a number of further topics, Section 7 discusses the related literature, and Section 8 concludes.

Our companion papers (Kamada and Kojima, 2014a,b) consider a model that generalizes the one that we use in this paper, and prove general results in such a context. Unless otherwise stated, the results in this paper are corollaries of the results in these papers.

2. MODEL

This section presents the model of matching with distributional constraints. Motivated by Japanese residency matching, we describe the model in terms of matching between doctors and hospitals, where there is a “regional cap,” that is, an upper bound on the number of doctors that can be matched to hospitals in each region.⁶ Later we discuss other matching problems in practice such as Chinese graduate school admission, U.K. medical matching, Scottish teacher matching, and college admissions in Ukraine and Hungary. The model is also applicable to diverse contexts discussed in the Introduction, such as doctor-hospital matching with specialty constraints, school choice with building size constraints, and matching with affirmative action constraints.

Let there be a finite set of doctors D and a finite set of hospitals H . Each doctor d has a strict preference relation \succ_d over the set of hospitals and being unmatched (being unmatched is denoted by \emptyset). For any $h, h' \in H \cup \{\emptyset\}$, we write $h \succeq_d h'$ if and only if $h \succ_d h'$ or $h = h'$. Each hospital h has a strict preference relation \succ_h over the set of subsets of doctors. For any $D', D'' \subseteq D$, we write $D' \succeq_h D''$ if and only if $D' \succ_h D''$ or $D' = D''$. We denote by $\succ = (\succ_i)_{i \in D \cup H}$ the preference profile of all doctors and hospitals.

⁶In practice, there may be multiple residency programs within a hospital. However, in such a case we use the word “a hospital” to mean a residency program for simplicity.

Doctor d is said to be **acceptable** to h if $d \succ_h \emptyset$.⁷ Similarly, h is acceptable to d if $h \succ_d \emptyset$. It will turn out that only rankings of acceptable partners matter for our analysis, so we often write only acceptable partners to denote preferences. For example,

$$\succ_d: h, h'$$

means that hospital h is the most preferred, h' is the second most preferred, and h and h' are the only acceptable hospitals under preferences \succ_d of doctor d .

Each hospital $h \in H$ is endowed with a (physical) **capacity** q_h , which is a nonnegative integer. We say that preference relation \succ_h is **responsive with capacity** q_h (Roth, 1985) if

- (1) For any $D' \subseteq D$ with $|D'| \leq q_h$, $d \in D \setminus D'$ and $d' \in D'$, $(D' \cup d) \setminus d' \succeq_h D'$ if and only if $d \succeq_h d'$,
- (2) For any $D' \subseteq D$ with $|D'| \leq q_h$ and $d' \in D'$, $D' \succeq_h D' \setminus d'$ if and only if $d' \succeq_h \emptyset$, and
- (3) $\emptyset \succ_h D'$ for any $D' \subseteq D$ with $|D'| > q_h$.

In words, preference relation \succ_h is responsive with a capacity if the ranking of a doctor (or keeping a position vacant) is independent of her colleagues, and any set of doctors exceeding its capacity is unacceptable. We assume that preferences of each hospital h are responsive with capacity q_h throughout the paper.

There is a finite set R which we call the set of **regions**. The set of hospitals H is partitioned into hospitals in different regions, that is, $H_r \cap H_{r'} = \emptyset$ if $r \neq r'$ and $H = \cup_{r \in R} H_r$, where H_r denotes the set of hospitals in region $r \in R$. For each $h \in H$, let $r(h)$ denote the region r such that $h \in H_r$. For each region $r \in R$, there is a **regional cap** q_r , which is a nonnegative integer.

A **matching** μ is a mapping that satisfies (i) $\mu_d \in H \cup \{\emptyset\}$ for all $d \in D$, (ii) $\mu_h \subseteq D$ for all $h \in H$, and (iii) for any $d \in D$ and $h \in H$, $\mu_d = h$ if and only if $d \in \mu_h$. That is, a matching simply specifies which doctor is assigned to which hospital (if any). A matching is **feasible** if $|\mu_r| \leq q_r$ for all $r \in R$, where $\mu_r = \cup_{h \in H_r} \mu_h$. In other words, feasibility requires that the regional cap for every region is satisfied. This requirement distinguishes the current environment from the standard model without regional caps: We allow for (though do not require) $q_r < \sum_{h \in H_r} q_h$, that is, the regional cap can be smaller than the sum of hospital capacities in the region.

⁷We denote singleton set $\{x\}$ by x when there is no confusion.

Since regional caps are a primitive of the environment, we consider a constrained efficiency concept. A feasible matching μ is **(constrained) efficient** if there is no feasible matching μ' such that $\mu'_i \succeq_i \mu_i$ for all $i \in D \cup H$ and $\mu'_i \succ_i \mu_i$ for some $i \in D \cup H$.

To accommodate the regional caps, we introduce a new stability concept that generalizes the standard notion. For that purpose, we first define two basic concepts. A matching μ is **individually rational** if (i) for each $d \in D$, $\mu_d \succeq_d \emptyset$, and (ii) for each $h \in H$, $d \succeq_h \emptyset$ for all $d \in \mu_h$, and $|\mu_h| \leq q_h$. That is, no agent is matched with an unacceptable partner and each hospital's capacity is respected.

Given matching μ , a pair (d, h) of a doctor and a hospital is called a **blocking pair** if $h \succ_d \mu_d$ and either (i) $|\mu_h| < q_h$ and $d \succ_h \emptyset$, or (ii) $d \succ_h d'$ for some $d' \in \mu_h$. In words, a blocking pair is a pair of a doctor and a hospital who want to be matched with each other (possibly rejecting their partners in the prescribed matching) rather than following the proposed matching.

When there are no binding regional caps (in the sense that $q_r \geq \sum_{h \in H_r} q_h$ for every $r \in R$), a matching is said to be stable if it is individually rational and there is no blocking pair. Gale and Shapley (1962) show that there exists a stable matching in that setting. In the presence of binding regional caps, however, there may be no such matching that is feasible (in the sense that all regional caps are respected). Thus in some cases every feasible and individually rational matching may admit a blocking pair.

A **mechanism** φ is a function that maps preference profiles to matchings. The matching under φ at preference profile \succ is denoted $\varphi(\succ)$ and agent i 's match is denoted by $\varphi_i(\succ)$ for each $i \in D \cup H$.

A mechanism φ is said to be **strategy-proof** if there does not exist a preference profile \succ , an agent $i \in D \cup H$, and preferences \succ'_i of agent i such that

$$\varphi_i(\succ'_i, \succ_{-i}) \succ_i \varphi_i(\succ).$$

That is, no agent has an incentive to misreport her preferences under the mechanism. Strategy-proofness is regarded as a very important property for a mechanism to be successful.⁸

⁸One good aspect of having strategy-proofness is that the matching authority can actually state it as the property of the algorithm to encourage doctors to reveal their true preferences. For example, the current webpage of the JRMP (last accessed on May 25, 2010, <http://www.jrmp.jp/01-ryui.htm>) states, as advice for doctors, that "If you list as your first choice a program which is not actually your first choice, the probability that you end up being matched with some hospital does not increase [...] the probability that you are matched with your actual first choice decreases." In the context of student placement in Boston, strategy-proofness was regarded as a desirable fairness property, in the sense that it provides

Unfortunately, however, there is no mechanism that produces a stable matching for all possible preference profiles and is strategy-proof even in a market without regional caps, that is, $q_r > |D|$ for all $r \in R$ (Roth, 1982).⁹ Given this limitation, we consider the following weakening of the concept requiring incentive compatibility only for doctors. A mechanism φ is said to be **strategy-proof for doctors** if there does not exist a preference profile \succ , a doctor $d \in D$, and preferences \succ'_d of doctor d such that

$$\varphi_d(\succ'_d, \succ_{-d}) \succ_d \varphi_d(\succ).$$

A mechanism φ is said to be **group strategy-proof for doctors** if there is no preference profile \succ , a subset of doctors $D' \subseteq D$, and a preference profile $(\succ'_{d'})_{d' \in D'}$ of doctors in D' such that

$$\varphi_d((\succ'_{d'})_{d' \in D'}, (\succ_i)_{i \in D \cup H \setminus D'}) \succ_d \varphi_d(\succ) \text{ for all } d \in D'.$$

That is, no subset of doctors can jointly misreport their preferences to receive a strictly preferred outcome for every member of the coalition under the mechanism.

We do not necessarily regard (group) strategy-proofness for doctors as a minimum desirable property that our mechanism should satisfy (our criticism of the JRMP mechanism in Section 3 does not hinge on (group) strategy-proofness), but it will turn out that the flexible deferred acceptance mechanism we propose in Section 5 does have this property.

As this paper analyzes the effect of regional caps in matching markets, it is useful to compare it with the standard matching model without regional caps. Gale and Shapley (1962) consider a matching model without any binding regional cap, which corresponds to a special case of our model in which $q_r > |D|$ for every $r \in R$. In that model, they propose the following **(doctor-proposing) deferred acceptance algorithm**:

- Step 1: Each doctor applies to her first choice hospital. Each hospital rejects the lowest-ranking doctors in excess of its capacity and all unacceptable doctors among those who applied to it, keeping the rest of the doctors temporarily (so doctors not rejected at this step may be rejected in later steps).

In general,

- Step t : Each doctor who was rejected in Step $(t - 1)$ applies to her next highest choice (if any). Each hospital considers these doctors *and* doctors who are temporarily held from the previous step together, and rejects the lowest-ranking

equal access for children and parents with different degrees of sophistication to strategize (Pathak and Sonmez, 2008).

⁹Remember that a special case of our model in which $q_r > |D|$ for all $r \in R$ is the standard matching model with no binding regional caps.

doctors in excess of its capacity and all unacceptable doctors, keeping the rest of the doctors temporarily (so doctors not rejected at this step may be rejected in later steps).

The algorithm terminates at a step in which no rejection occurs. The algorithm always terminates in a finite number of steps. Gale and Shapley (1962) show that the resulting matching is stable in the standard matching model without any binding regional cap.

Even though there exists no strategy-proof mechanism that produces a stable matching for all possible inputs, the deferred acceptance mechanism is (group) strategy-proof for doctors (Dubins and Freedman, 1981; Roth, 1982).¹⁰ This result has been extended by many subsequent studies, suggesting that the incentive compatibility of the mechanism is quite robust and general.¹¹

3. A MOTIVATING EXAMPLE OF AN INEFFICIENT MECHANISM

In this section we formulate the Japan Residency Matching Program (JRMP) mechanism, one of the mechanisms with distributional constraints used in practice, and point out its deficiencies. As we have mentioned in the Introduction, there are a number of practical markets besides the Japanese one in which distributional constraints are imposed, such as the UK medical match, the admission problem for the Chinese master's programs, and so forth. Here we chose the Japanese case because the JRMP mechanism is easy to formulate so its problem can be transparently seen and a clear-cut comparison with our mechanism can be made. Other mechanisms are analyzed in subsequent sections.

In the JRMP mechanism, there is an exogenously given (government-imposed) non-negative integer $\bar{q}_h \leq q_h$, which we call **target capacity**, for each hospital h such that

¹⁰Ergin (2002) defines a stronger version of group strategy-proofness. It requires that no group of doctors can misreport preferences jointly and make some of its members strictly better off without making any of its members strictly worse off. He identifies a necessary and sufficient condition for the deferred acceptance mechanism to satisfy this version of group strategy-proofness.

¹¹Researches generalizing (group) strategy-proofness of the mechanism include Abdulkadiroğlu (2005), Hatfield and Milgrom (2005), Martinez, Masso, Neme, and Oviedo (2004), Hatfield and Kojima (2009, 2010), and Hatfield and Kominers (2009, 2010).

$\sum_{h \in H_r} \bar{q}_h \leq q_r$ for each region $r \in R$.^{12,13} The **JRMP mechanism** is a rule that produces the matching resulting from the deferred acceptance algorithm except that, for each hospital h , it uses \bar{q}_h instead of q_h as the hospital's capacity.

The JRMP mechanism is based on a simple idea: In order to satisfy regional caps, simply force hospitals to be matched to a smaller number of doctors than their real capacities, but otherwise use the standard deferred acceptance algorithm. Note, however, that target capacities are not feasibility constraints by themselves: the goal of Japanese policy makers is to satisfy regional caps and target capacities were introduced to achieve that goal.

Although the mechanism is a variant of the deferred acceptance algorithm, it suffers from at least one problem. Despite the government's intention, the result of the JRMP mechanism is not necessarily efficient, as seen in the following example.

Example 1 (JRMP mechanism does not necessarily produce an efficient matching). There is one region r with regional cap $q_r = 10$, in which two hospitals, h_1 and h_2 , reside. Each hospital h has a capacity of $q_h = 10$. Suppose that there are 10 doctors, d_1, \dots, d_{10} . Preference profile \succ is as follows:

$$\begin{aligned} \succ_{h_i}: d_1, d_2, \dots, d_{10} \quad \text{for } i = 1, 2, \\ \succ_{d_j}: h_1 \text{ if } j \leq 3 \quad \text{and} \quad \succ_{d_j}: h_2 \text{ if } j \geq 4. \end{aligned}$$

Thus, three doctors prefer hospital h_1 to being unmatched (the option \emptyset) to hospital h_2 , while the other seven doctors prefer hospital h_2 to being unmatched to hospital h_1 . Let the target capacities be $\bar{q}_{h_1} = \bar{q}_{h_2} = 5$.¹⁴

At the first round of the JRMP algorithm, doctors d_1, d_2 and d_3 apply to hospital h_1 , and the rest of the doctors apply to hospital h_2 . Hospital h_1 does not reject anyone at this round, as the number of applicants is less than its target capacity, and all applicants are acceptable. Hospital h_2 rejects d_9 and d_{10} and accepts other applicants, because the

¹²Note that we allow the sum of target capacities to be strictly smaller than the regional cap. This is necessary if the sum of hospital capacities is strictly smaller than the regional cap; we allow this possibility even otherwise. All results, including (counter)examples, hold when we assume that the sum of target capacities is equal to the regional cap.

¹³In our model, \bar{q}_h is exogenously given for each hospital h . In the current Japanese system, if the sum of the hospitals' capacities exceeds the regional cap, then the target \bar{q}_h of each hospital h is set at an integer close to $\frac{q_r}{\sum_{h' \in H_r} q_{h'}} \cdot q_h$. That is, each hospital's target is (roughly) proportional to its capacity. This might suggest that hospitals have incentives to misreport their capacities. As explained in footnote 4, however, the capacity can be considered exogenous in the Japanese context.

¹⁴The specification of target capacities follows the formula used in Japan that we mentioned earlier.

number of applicants exceeds the *target capacity* (not the hospital’s capacity itself!), and it prefers doctors with smaller indices (and all doctors are acceptable). Since d_9 and d_{10} find h_1 unacceptable, they do not make further applications, so the algorithm terminates at this point. Hence the resulting matching μ is such that

$$\mu = \begin{pmatrix} h_1 & h_2 & \emptyset \\ d_1, d_2, d_3 & d_4, d_5, d_6, d_7, d_8 & d_9, d_{10} \end{pmatrix}.$$

Consider a matching μ' defined by,

$$\mu' = \begin{pmatrix} h_1 & h_2 \\ d_1, d_2, d_3 & d_4, d_5, d_6, d_7, d_8, d_9, d_{10} \end{pmatrix}.$$

Since the regional cap is still respected, μ' is feasible. Moreover, every agent is weakly better off with d_9, d_{10} , and h_2 being strictly better off than at μ . Hence we conclude that the JRMP mechanism results in an inefficient matching in this example.¹⁵ \square

Remark 1. We note that there is a sense in which the matching μ is not stable. For example, hospital h_2 and doctor d_9 constitute a blocking pair while the regional cap for r is not binding. That is, even after d_9 is matched with h_2 , the total number of doctors in the region is 9, which is less than the regional cap of 10.¹⁶ Although this argument may appear straightforward, defining stability in the presence of regional caps is not a trivial task. In Section 4, we define the notion of stability and show that the matching μ is not stable. \square

The above example suggests that a problem of the JRMP mechanism is its lack of flexibility: The JRMP mechanism runs as if the target capacity is the actual capacity

¹⁵In this example, not all hospitals are acceptable to all doctors. One may wonder whether this is an unrealistic assumption because doctors may be so willing to work that any hospital is acceptable. However, the example can be easily modified so that all hospitals are acceptable to all doctors while some doctors are unacceptable to some hospitals (which may be a natural assumption because, for instance, typically a hospital only lists doctors who they interviewed). Also, in many markets doctors apply to only a small subset of hospitals. In 2009, for instance, a doctor applied to only 3.3 hospitals on average (Japan Residency Matching Program, 2009a).

¹⁶One may wonder whether the “failure of stability” depends on the assumption that some agents find some of the potential partners unacceptable. However, a similar example can be constructed even if we require every agent finds every potential partner acceptable. For instance, modify the market in the example by introducing another hospital h_3 in another region with regional cap two; let h_3 find every doctor acceptable and have two positions; d_1, d_2 and d_3 prefer h_1 to h_3 to h_2 to being unmatched, while all other doctors prefer h_2 to h_3 to h_1 to being unmatched (thus every doctor finds all hospitals acceptable).

of hospitals, thus rejecting an application of a doctor to a hospital unnecessarily. The mechanism that we propose in Section 5 overcomes problems of inefficiency (and stability) by, intuitively speaking, making the target capacities flexible. Before formally introducing this mechanism, we define and discuss the goals that we try to achieve with the mechanism.

4. GOAL SETTING: STABILITY AND STRATEGY-PROOFNESS

As we discussed earlier, there may be no stable matching in the traditional sense that satisfies feasibility. Given this observation, this section defines a weaker stability concept, in which certain types of blocking pairs are tolerated.

One possible notion of stability in this context may be to require that satisfying any remaining blocking pair would violate the regional cap. Our companion paper, Kamada and Kojima (2014b), studies such a stability concept, called strong stability. One of the main findings of that paper is that there does not necessarily exist a stable matching that satisfies strong stability.

Given that strong stability is “too strong” in the above sense, a weaker stability notion that still takes regional caps into account is hoped for. To define an appropriate stability concept, we restrict blocking pairs that are regarded as legitimate. We do so by using the notion of target capacity. More specifically, we regard target capacities $(\bar{q}_h)_{h \in H}$ as reflecting certain distributional goals (though not feasibility constraints) and define a stability concept that respects target capacities as much as possible.¹⁷

Definition 1. A matching μ is **stable** if it is feasible, individually rational, and if (d, h) is a blocking pair then (i) $|\mu_{r(h)}| = q_{r(h)}$, (ii) $d' \succ_h d$ for all doctors $d' \in \mu_h$, and

$$(iii) \text{ either } \mu_d \notin H_{r(h)} \text{ or } |\mu'_h| - \bar{q}_h > |\mu'_{\mu_d}| - \bar{q}_{\mu_d},$$

where μ' is the matching such that $\mu'_d = h$ and $\mu'_{d'} = \mu_{d'}$ for all $d' \neq d$.

As stated in the definition, only certain blocking pairs are tolerated under stability. Any blocking pair that may remain is in danger to violate the regional cap since condition (i) implies that the cap for the blocking hospital’s region is currently binding, and condition (ii) implies that the only blocking involves filling a vacant position.

There are two possible cases under (iii). The first case implies that the blocking doctor is not currently assigned in the hospital’s region, so the blocking pair violates the regional

¹⁷Depending on the distributional goals, target capacities can be set differently from those specified in the description of the JRMP mechanism. Kamada and Kojima (2014a) discuss alternative ways to allocate target capacities.

cap. The second part of condition (iii) declares that certain types of blocking pairs within a region (note that $\mu_d \in H_{r(h)}$ holds in the remaining case) are not regarded as legitimate deviations (as clarified below, our interpretation of stability concepts is normative). To see this point, consider the inequality in condition (iii),

$$(4.1) \quad |\mu'_h| - \bar{q}_h > |\mu'_{\mu_d}| - \bar{q}_{\mu_d}.$$

The left-hand side is the number of doctors matched to h in excess of its target \bar{q}_h if d actually moves to h , realizing a new matching μ' . The right hand side is the number of doctors matched to the original hospital μ_d in excess of its target \bar{q}_{μ_d} if d moves out of μ_d . This property says that such a movement will not decrease the imbalance of over-target numbers of matching across hospitals. Intuitively, if the movement of the doctor in the blocking pair “equalizes” the excesses over the target capacities compared to the current matching (that is, $|\mu_h| - \bar{q}_h < |\mu'_h| - \bar{q}_h \leq |\mu'_{\mu_d}| - \bar{q}_{\mu_d} < |\mu_{\mu_d}| - \bar{q}_{\mu_d}$), then such a movement should be regarded as a legitimate deviation. Thus, the only blocking pair within a region that can remain under this definition should satisfy condition (4.1).

This concept reduces to the standard stability concept of Gale and Shapley (1962) if there are no binding regional caps.

The implicit idea behind the definition is that the government or some authority can interfere and prohibit a blocking pair to be formed if regional caps are an issue. Indeed, in Japan, participants seem to be effectively forced to accept the matching announced by the clearinghouse because a severe punishment is imposed on deviators.¹⁸ One might then wonder “If the government has the power to prohibit a blocking pair in certain cases, why doesn’t it have the power to do so in all cases, so why do we care about stability in the first place?”

Our view is that even if the clearinghouse has power to enforce a matching (which may be the case in the Japanese residency match), an assignment that completely ignores participants’ preferences would be undesirable. Indeed, as we will discuss in Appendix A.1, the introduction of a stable matching mechanism in this market was motivated by the criticism that the previous assignment system was “unfair” and “inefficient,” rather than by a desire to prevent participants from circumventing the assignment by forming “blocking pairs.”¹⁹ In other words, we view minimizing blocking pairs as a normative

¹⁸For example, violating hospitals can be excluded from participating in the matching mechanism in subsequent years (Japan Residency Matching Program, 2010).

¹⁹Another example of a labor market using a stable mechanism despite being heavily regulated is the labor market for junior academic positions in France (Haeringer and Ichle, 2010).

criterion.²⁰ Given this observation, our stability captures the idea that it is desirable to minimize blocking pairs so that the only blocking pairs are “caused” by regional caps, which may be a legitimate reason to deny a blocking pair.²¹

We note that there may be other natural definitions of stability. For example, it may be desirable to entitle a hospital with capacity 20 to twice as many doctors over the target as a hospital with capacity 10. There may also be other criteria that are deemed desirable, including even cases in which target capacities are not defined. To address this issue, Section 6.5 and Kamada and Kojima (2014a) consider a class of stability concepts that includes the stability in Definition 1 as a special case and accommodates the above ideas. For each stability notion in this class, they present a mechanism that generates a stable matching. In this paper, we assume that the policy goal is expressed as in condition (4.1). This particular policy goal is chosen here because it is expositionally simple and appears to be a reasonable starting point. However, it is not a necessary requirement for our analysis to work, as we will observe in Section 6.5 and Kamada and Kojima (2014a). The choice of a particular variant of stability should be in part the product of society’s preferences. We restrict ourselves to proposing solutions that are flexible enough to meet as wide a range of policy goals as possible. See Kamada and Kojima (2014a) for a partial list of other possible social preferences that we can accommodate.

Stability implies the following desirable property:

Theorem 1. *Any stable matching is (constrained) efficient.*

When there is no regional cap (in which case stability reduces to the standard concept of stability), a matching is stable if and only if it is in the core, and any core outcome is efficient. Without regional caps, Theorem 1 follows straightforwardly from these facts. With regional caps, however, there is no obvious way to define an appropriate cooperative game or a core concept. Theorem 1 states that efficiency of stable matchings still holds in our model.²²

Remark 2. Since the outcome of the JRMP mechanism in Example 1 is not efficient, Theorem 1 implies that it is not stable either. This is easy to check by inspection as well.

²⁰“No justified envy” in the school choice literature corresponds to “no blocking pair” in our context, and it is viewed as a normative criterion.

²¹Another obvious normative criterion is (constrained) efficiency. Indeed, it will turn out that stability implies efficiency (Theorem 1). Thus stability has an additional normative appeal.

²²To overcome the above difficulty, the proof presented in Kamada and Kojima (2014b) shows this result directly instead of associating stability to the core in a cooperative game. In fact, Kamada and Kojima (2014b) observe (constrained) efficiency of “weak stability,” which is weaker than stability.

As hinted in Remark 1, (d_9, h_2) is a blocking pair that does not satisfy condition (i) in the definition of stability (Definition 1). That is, matching d_9 to h_2 does not violate the regional cap. \square

A natural question is whether a stable matching exists in every market. This question will be answered in the affirmative in the next section, where we propose an algorithm that always generates a stable matching.

5. A NEW MECHANISM: FLEXIBLE DEFERRED ACCEPTANCE

We present a new mechanism that, for any given input, results in a stable matching. To do so, we first define the **flexible deferred acceptance algorithm**:

For each $r \in R$, specify an order of hospitals in region r : Denote $H_r = \{h_1, h_2, \dots, h_{|H_r|}\}$ and order h_i earlier than h_j if $i < j$. Given this order, consider the following algorithm.

- (1) Begin with an empty matching, that is, a matching μ such that $\mu_d = \emptyset$ for all $d \in D$.
- (2) Choose a doctor d who is currently not tentatively matched to any hospital and who has not applied to all acceptable hospitals yet. If such a doctor does not exist, then terminate the algorithm.
- (3) Let d apply to the most preferred hospital \bar{h} at \succ_d among the hospitals that have not rejected d so far. Let r be the region such that $\bar{h} \in H_r$.
- (4) (a) For each $h \in H_r$, let D'_h be the entire set of doctors who have applied to but have not been rejected by h so far and are acceptable to h . For each hospital $h \in H_r$, choose the \bar{q}_h best doctors according to \succ_h from D'_h if they exist, and otherwise choose all doctors in D'_h . Formally, for each $h \in H_r$ choose D''_h such that $D''_h \subset D'_h$, $|D''_h| = \min\{\bar{q}_h, |D'_h|\}$, and $d \succ_h d'$ for any $d \in D''_h$ and $d' \in D'_h \setminus D''_h$.
- (b) Start with a tentative match D''_h for each hospital $h \in H_r$. Hospitals take turns to choose (one doctor at a time) the best remaining doctor in their current applicant pool. Repeat the procedure (starting with h_1 , proceeding to h_2, h_3, \dots and going back to h_1 after the last hospital) until the regional quota q_r is filled or the capacity of the hospital is filled or no doctor remains to be matched. All other applicants are rejected.²³

²³Formally, let $\iota_i = 0$ for all $i \in \{1, 2, \dots, |H_r|\}$. Let $i = 1$.

(i) If either the number of doctors already chosen by the region r as a whole equals q_r , or $\iota_i = 1$, then reject the doctors who were not chosen throughout this step and go back to Step 2.

We define the **flexible deferred acceptance mechanism** to be a mechanism that produces, for each input, the matching at the termination of the above algorithm.²⁴

The flexible deferred acceptance mechanism is analogous to the deferred acceptance mechanism and the JRMP mechanism. What distinguishes the flexible deferred acceptance mechanism from the JRMP mechanism is that the former lets hospitals fill their capacities “more flexibly” than the latter. To see this point, first observe that the way that hospitals choose doctors who applied in Step 4a is essentially identical to the one in the JRMP algorithm. As seen before, the JRMP may result in an inefficient and unstable matching because this step does not let hospitals tentatively keep doctors beyond target capacities even if regional caps are not binding. This is addressed in Step 4b. In this step, hospitals in a region are allowed to keep more doctors than their target capacities if doing so keeps the regional caps respected. Thus there is a sense in which this algorithm corrects the deficiency of the JRMP mechanism while following closely the deferred acceptance algorithm.

In the flexible deferred acceptance algorithm, one needs to specify an ordering of hospitals. There are at least two reasons that this requirement may not cause problems such as conflicts among hospitals to get a “desirable position” in the order. The first is that, as we will discuss in Subsection 6.6, the effect of different ways of setting orders on the welfare of hospitals is ambiguous. More specifically, it may be the case that a hospital is better off being ordered later under some specification of preference profiles, while the opposite may be true under other specifications. Second, the flexible deferred acceptance algorithm can be modified without losing its desirable properties, by adding “Step 0” in which a particular ordering is chosen according to some probabilistic rule. The aforementioned problems can be resolved by, for example, choosing an order according to the uniform probability distribution.

The following example illustrates how the flexible deferred acceptance algorithm works.

Example 2 (The flexible deferred acceptance algorithm). Consider the same example as in Example 1. Recall that the JRMP mechanism can produce an inefficient and unstable

-
- (ii) Otherwise, let h_i choose the most preferred (acceptable) doctor in D'_{h_i} at \succ_{h_i} among the doctors that have not been chosen by h_i so far, if such a doctor exists and the number of doctors chosen by h_i so far is strictly smaller than q_{h_i} .
 - (iii) If no new doctor was chosen at Step 4(b)ii, then set $\iota_i = 1$. If a new doctor was chosen at Step 4(b)ii, then set $\iota_j = 0$ for all $j \in \{1, 2, \dots, |H_r|\}$. If $i < |H_r|$ then increment i by one and if $i = |H_r|$ then set i to be 1 and go back to Step 4(b)i.

²⁴We show in Theorem 2 that the algorithm stops in a finite number of steps.

matching. By contrast, the flexible deferred acceptance algorithm selects a matching that is efficient and stable. Precisely, let doctors apply to hospitals in the specified order. For doctors d_1 to d_8 , the algorithm does not proceed to Step 4b, as the number of doctors in each hospital is no larger than its target. When d_9 applies, doctors d_1, \dots, d_8 are still matched to hospitals in Step 4a, and d_9 is matched to h_2 in Step 4b. In the same way, when d_{10} applies, doctors d_1, \dots, d_8 are still matched to hospitals in Step 4a, and d_9 and d_{10} are matched to h_2 in Step 4b. Hence an efficient and stable matching results. Intuitively, the algorithm treats doctors' applications in a more flexible manner than in the JRMP algorithm. This is the idea behind the name “*flexible* deferred acceptance.” \square

The following is the main result of this section.

Theorem 2. *The flexible deferred acceptance algorithm stops in finite steps. The mechanism produces a stable matching for any input and is group strategy-proof for doctors.*

To see an intuition for the stability of the flexible deferred acceptance mechanism, recall that there is a sense in which hospitals fill their capacities “flexibly.” More specifically, at each step of the algorithm hospitals can tentatively accept doctors beyond their target capacities as long as the regional cap is not violated. Then the kind of rejection that causes instability in Example 1 does not occur in the flexible deferred acceptance algorithm.²⁵ Thus an acceptable doctor is rejected from a preferred hospital either because there are enough better doctors in that hospital, or the regional quota is filled by other doctors. So such a doctor cannot form a blocking pair, suggesting that the resulting matching is stable.²⁶

The intuition for strategy-proofness for doctors is similar to the one for the deferred acceptance mechanism. A doctor does not need to give up trying for her first choice because, even if she is rejected, she will be able to apply to her second choice, and so forth. In other words, the “deferred” acceptance guarantees that she will be treated equally if she applies to a position later than others.

Although the above are rough intuitions of the results, the formal proof, presented in Kamada and Kojima (2014a), takes a different approach. It relates our model to the model of “(many-to-many) matching with contracts” (Hatfield and Milgrom, 2005). The basic idea of the proof is to regard each region as a consortium of hospitals that acts as one

²⁵Indeed, in the market of Example 1, the result of the flexible deferred acceptance mechanism matches 10 doctors as in matching μ' . Thus in this example the flexible deferred acceptance mechanism strictly improves upon the JRMP mechanism.

²⁶Since hospitals take turns one by one when they tentatively accept doctors, no blocking pair involving a doctor's movement within a region can “equalize” the distribution of doctors in the region.

agent, and to define its choice function that selects a subset from any given collection of pairs (contracts) of a doctor and a hospital in the region. Once our model is successfully connected to the matching model with contracts, properties of the latter model can be invoked to show the theorem. In fact, the proof shows that a more general result holds which can be applicable to the class of stability concepts mentioned in Section 6.5 and that the current model is indeed a special case of the general model. Theorem 2 then follows as a corollary of these results.

Theorems 1 and 2 imply the following appealing welfare property of the flexible deferred acceptance mechanism.

Corollary 1. *The flexible deferred acceptance mechanism produces an efficient matching for any input.*

Recall that the JRMP mechanism does not necessarily produce an efficient matching. In light of this observation, Corollary 1 implies that the flexible deferred acceptance mechanism has a better efficiency property than the JRMP mechanism.

The next two propositions formalize the idea that the flexible deferred acceptance mechanism respects target capacities and regional caps as much as possible.

Proposition 1. *If the number of doctors matched with $h \in H$ in the flexible deferred acceptance mechanism is strictly less than its target capacity, then for any $d \in D$ who are not matched with h , either d is unacceptable to h or d prefers its current match to h .*

Proof. Assume that d prefers h to her outcome under the flexible deferred acceptance mechanism. Then d has applied to h and was rejected under the flexible deferred acceptance algorithm. If the number of doctors matched with h in the flexible deferred acceptance mechanism is strictly less than its target capacity, then the number of doctors who have ever applied to h and are acceptable to h is strictly smaller than the target capacity of h . This implies that any doctor who applied to h and was rejected in the flexible deferred acceptance algorithm is unacceptable to h . In particular d is unacceptable, completing the proof. \square

In other words, the flexible deferred acceptance mechanism does not prevent a doctor from being matched to an underserved hospital, relative to the target capacity, in the name of respecting the regional caps.²⁷ This result suggests, as we argued informally when defining stability, that the choice of target capacities can be utilized as a means to achieve distributional goals.

²⁷The conclusion of the theorem applies even if the regional cap is already binding, thus this property is not implied by the fact that the outcome of the flexible deferred acceptance algorithm is stable.

- Proposition 2.** (1) *If the number of doctors matched with $h \in H$ in the flexible deferred acceptance mechanism is strictly less than its target capacity, then the set of doctors matched with h under the (unconstrained) deferred acceptance mechanism is a subset of the one under the flexible deferred acceptance mechanism.*
- (2) *If the number of doctors matched in $r \in R$ in the flexible deferred acceptance mechanism is strictly less than its regional cap, then each hospital h in r weakly prefers a matching produced by the flexible deferred acceptance mechanism to the one under the (unconstrained) deferred acceptance mechanism. Moreover, the number of doctors matched to any such h in the former matching is weakly larger than that in the latter.*

This result implies that, whenever a hospital or a region is underserved under the flexible deferred acceptance mechanism, the (unconstrained) deferred acceptance mechanism cannot improve the match at such a hospital or a region. This result offers a sense in which the flexible deferred acceptance mechanism avoids inflicting costs on underserved hospitals or regions.

6. DISCUSSION

This section provides several discussions. Subsection 6.1 studies mechanisms in other contexts and countries, and shows that those existing mechanisms suffer from problems similar to those we pointed out for Japanese medical residency match. Subsection 6.2 considers an alternative mechanism that is often suggested to us, and shows that it is not strategy-proof for doctors. In Subsection 6.3, we consider the rural hospital theorem of Roth (1986) and a related concept of the “match rate,” the ratio of the number of all matched doctors to the total number of doctors (matched plus unmatched). Subsection 6.4 studies the existence issue of a side-optimal stable matching, that is, a matching that is preferred by all doctors or by all hospitals. Subsection 6.5 generalizes stability and the flexible deferred acceptance mechanism. Subsection 6.6 examines the welfare effect of different choices of picking orders over hospitals, target capacities, and regional caps, and Subsection 6.7 considers “floor constraints” instead of “ceiling constraints” (regional caps).

6.1. Mechanisms in China and the United Kingdom. In the main sections we showed that the flexible deferred acceptance mechanism has desirable properties such as efficiency, stability, and strategy-proofness for doctors, and we observed that it outperforms the JRMP mechanism. In this section we analyze several other mechanisms used in

practice, by formulating them and pointing out their respective deficiencies. The analyses confirm the applicability of our mechanism in various markets around the world.

Specifically, in what follows we analyze the Chinese Graduate school admission and the medical matching problem in the United Kingdom. More comprehensive descriptions of the problems and the analyses can be found in Appendix A. Appendix A also discusses more examples mentioned in the Introduction such as college admission in Ukraine and matching of new teachers in Scotland.

6.1.1. *Chinese Graduate School Admission.* The first problem we study is the Chinese graduate school admission. As described briefly in the Introduction, master’s programs are categorized as either academic or professional, and the Chinese government is currently trying to reduce the number of academic master students. To achieve this goal, the government decided to reduce the available seats of each academic master’s program by about 25 percent by 2015.²⁸ In our framework, the policy goal of the Chinese government can be translated into imposing the “regional cap” on the set of all academic master’s programs, where the regional cap is about 75 percent of the sum of the true capacities across all the academic programs. Then the government sets the target capacity of each academic master’s program at about 75 percent of the true capacity.

Given the target capacities above, the main round of Chinese graduate admission runs as follows. Using an application website, each student applies to one graduate program. Given the set of applicants, each graduate school accepts its most preferred acceptable students up to its target capacity and rejects everyone else. All matches are final.²⁹

This mechanism suffers from several drawbacks. First, it is easy to see that this mechanism is not strategy-proof for students. Moreover, the mechanism may produce an unstable and inefficient matching. Here we focus on a particular source of inefficiency and instability that shares a certain feature with the JRMP mechanism: In the Chinese admission mechanism, the target capacities are used as rigid constraints, so the seats of a certain academic program beyond its target capacity *must* remain unfilled, even if some students prefer to be matched to these seats. This property holds true whatever mechanism the Chinese government uses as long as the target capacities are treated as

²⁸To achieve this goal gradually, the government plans to reduce the number of seats by about 5 percent every year until 2015.

²⁹Here we are describing the main round of the admission process. There is a “guaranteed assignment” in which especially high-achieving students are admitted before the main round begins, and there is also a supplementary round called an “adjustment process” which is largely decentralized. We do not discuss these processes in the main text because they are not directly relevant to the topic of distributional constraints. See instead Appendix A.2.1 for detail of these rounds.

rigid constraints, so the problem is orthogonal to the inefficiency coming from the fact that each student can list only one program. In fact, we show in Appendix A.2 that even a (complete information) pure-strategy subgame-perfect equilibrium outcome (of a game in which students submit their preference list first and then colleges admit students) can be unstable and inefficient, while we also show that the outcome is always stable and efficient if there is no binding regional cap (so that the target capacity of a program is equal to its physical capacity).

6.1.2. *United Kingdom.* As mentioned briefly in the Introduction, the mechanism used for medical match in the United Kingdom is based on a different idea from the mechanisms we have discussed so far. The process has two rounds, in which doctors are matched to a region first, and then to a program within their matched region. For a policy maker who desires to control the distribution of doctors across different regions, assigning doctors using such a two-round scheme may appear to be an appealing alternative to the JRMP mechanism or the Chinese mechanism. As we will see in Appendix A.4, however, this mechanism may also result in inefficiency and instability, and it entails incentive problems as well.

There have been several changes of the mechanism in recent years, and the matching mechanism in the second round varies across regions. To be specific, however, assume that both rounds use the serial dictatorship. Serial dictatorship is in use in the first round since 2012, and it is in use in the second round in Scotland since 2010. For simplicity we focus on this mechanism although the same points can be made in other mechanisms as well.³⁰

Suppose that a doctor whose first choice is a hospital h_1 in a region r_1 , the second choice is h_2 in another region r_2 , and her third choice is h_3 in region r_1 . Assume that, in the first round, the doctor lists r_1 as her first choice and is matched to it. However, it is possible that she is matched to h_3 in the second round, while h_2 in region r_2 prefers her to one of the doctors matched to it. Then this doctor and h_2 form a blocking pair that is not tolerated under our stability concept (or even the weak stability concept as defined in Kamada and Kojima (2014b)), implying that the resulting matching is not (weakly) stable.

Intuitively, instability can happen for the following reason: In the first round a student may apply to and is matched to a region where her preferred hospitals are located. But then in the second round she may end up being matched to a hospital that she prefers less to a hospital in another region which prefers her to one of its matched doctors. In other

³⁰This point will be shown more formally in Appendix A.4.2, in which another two-round mechanism is analyzed.

words, since the matching between doctors and regions are finalized before the ultimate matching to a hospital is decided, the resulting matching could result in instability. A similar example shows that the matching can be inefficient. Moreover, under this mechanism there does not generally exist a dominant strategy for doctors because a doctor's best report in the first round depends on which hospital in the region she will end up with. Concrete examples making these points are found in Appendix A.4.

Scotland's matching between new teachers and schools uses a similar two-round matching mechanism. As in the U.K. medical match, the matching clearinghouse first matches teachers to a local authority, which then assigns teachers matched to it to schools under its control. As such, this mechanism also suffers from similar instability, inefficiency, and strategic problems. See Appendix A.4 for detail.

6.2. The Iterated Deferred Acceptance Mechanism. As a solution to the efficiency and stability problems in the JRMP mechanism, we often encounter suggestions by government officials and matching theorists, saying that the **iterated deferred acceptance (iterated DA)** mechanism that uses the following algorithm may be useful: This algorithm consists of finite steps of rounds. In round 1, the deferred acceptance algorithm is run regarding the target capacities as the real capacities. If the resulting matching fills all the target capacities, then the algorithm stops. Otherwise, the algorithm proceeds to round 2 after the target capacities are modified as follows: hospitals set their new target capacities equal to their matched numbers of doctors if they have vacant seats relative to their target capacities; these vacant seats are reallocated to other hospitals in the same region according to a certain pre-specified rule. In round 2, the deferred acceptance algorithm is run with these modified target capacities. If the resulting matching fills all the new target capacities then the algorithm stops and otherwise it continues. We do the same in all other rounds, with a restriction that once a hospital has reduced its target capacity then it never increases (and require that the algorithm stop if no further reallocation is possible).

As one might expect, this mechanism produces a stable matching in Example 1. However, it turns out that this mechanism is not strategy-proof for doctors.

Example 3. Consider a market with two doctors, d_1 and d_2 , and two hospitals h_1 and h_2 in a single region with regional cap 2. Each doctor prefers h_1 to h_2 to being unmatched. Each hospital is associated with a capacity of 2 and a target capacity of 1, and prefers d_1 to d_2 to being unmatched. In this market, the iterated DA ends in one round, resulting

in the matching

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

Doctor d_2 has an incentive to misreport her preferences. For, if she reports that she prefers h_1 to being unmatched to h_2 , then the iterated DA proceeds to the second round with one seat moving from h_2 to h_1 , and in the second round the matching

$$\mu' = \begin{pmatrix} h_1 & h_2 \\ d_1, d_2 & \emptyset \end{pmatrix},$$

is realized and the algorithm stops. Since d_2 prefers μ'_{d_2} to μ_{d_2} , the iterated DA mechanism is not strategy-proof for doctors. \square

6.3. The Rural Hospital Theorem and The Match Rate. In this subsection, we show that the conclusion of the rural hospital theorem does not hold in our environment. Motivated by this finding, we study how the flexible deferred acceptance mechanism works in terms of the match rate, that is, the proportion of the number of all matched doctors to the total number of doctors (matched plus unmatched).

6.3.1. The Rural Hospital Theorem. The rural hospital theorem (Roth, 1986) states that, in a matching model without regional caps, any hospital that fails to fill all its positions in one stable matching is matched to an identical set of doctors in all stable matchings. It also states that the set of unmatched doctors is identical across all stable matchings.

The theorem is of particular interest when we consider allocating a sufficient number of doctors to rural areas. Although the rural hospital theorem might suggest that increasing the number of doctors in a particular set of hospitals is impossible, the conclusion of the theorem does not necessarily hold in our context with regional caps.³¹ The following example makes this point clear.

Example 4 (The conclusion of the rural hospital theorem does not hold). There is one region r with regional cap $q_r = 1$, in which two hospitals, h_1 and h_2 , reside. Each hospital h has a capacity of $q_h = 1$. Let the target capacities be arbitrary. Suppose that there are two doctors, d_1 and d_2 . We assume the following preferences:

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

³¹In fact, this example shows that the conclusion fails even with the most stringent concept of strong stability as studied by Kamada and Kojima (2014b).

It is straightforward to check that there are two stable matchings,

$$\mu = \begin{pmatrix} h_1 & h_2 & \emptyset \\ d_1 & \emptyset & d_2 \end{pmatrix},$$

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

Notice that hospital h_1 fills its capacity in matching μ while it does not do so in matching μ' . Also, d_1 is matched to a hospital in matching μ while unmatched in matching μ' . Hence both conclusions of the rural hospital theorem fail. \square

One might suspect that, although the rural hospital theorem does not apply, it might be the case that each region attracts the same number of doctors in any stable matchings. The following example shows that this is not true.

Example 5 (The number of doctors matched to hospitals in a rural region may be different in different stable matchings). We modify Example 4 by adding one more region r' , which we interpret here as a “rural region” for the sake of discussion. Region r' has the regional cap of $q_{r'} = 1$, and one hospital, h_3 , resides in it. Suppose that h_3 has a capacity of $q_{h_3} = 1$. The preferences are modified as follows:

$$\succ_{h_1}: d_1, \quad \succ_{h_2}: d_2, \quad \succ_{h_3}: d_1,$$

$$\succ_{d_1}: h_1, h_3, \quad \succ_{d_2}: h_2.$$

It is straightforward to check that there are two stable matchings for any target capacity profiles,

$$\mu = \begin{pmatrix} h_1 & h_2 & h_3 & \emptyset \\ d_1 & \emptyset & \emptyset & d_2 \end{pmatrix},$$

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

Thus the hospital in rural region r' does not attract any doctors in matching μ , while it attracts one doctor in matching μ' . \square

Hence, when the number of doctors matched to hospitals in rural regions matters, the choice of a mechanism *is* an important issue, in the presence of regional caps.

6.3.2. *The Match Rate.* Related to the rural hospital theorem is the notion of “match rate,” which is the ratio of the number of all matched doctors to the total number of doctors (matched plus unmatched). The match rate seems to be a measure that many

people care about. For example, match rates are listed on the annual reports published by the NRMP and the JRMP.³² This is perhaps because the match rate is an easy measure for participants to understand.³³

Although it would be desirable to select a matching that has the maximum match rate among the stable matchings, the following example shows that the flexible deferred acceptance mechanism fails to do so.

Example 6 (The flexible deferred acceptance mechanism does not necessarily select a matching with the highest match rate among stable matchings). Take the same example as in Example 5. Also, let the target profile be $(\bar{q}_{h_1}, \bar{q}_{h_2}, \bar{q}_{h_3}) = (1, 0, 1)$. Then, the flexible deferred acceptance mechanism always selects a matching μ defined in Example 5. But this has a match rate of $1/2$, while the other matching, namely μ' defined in Example 5, has a match rate of 1. \square

It is unfortunate that the flexible deferred acceptance mechanism does not necessarily maximize the match rate within stable matchings, but the following example shows that this is a necessary consequence of requiring strategy-proofness for doctors.

Example 7 (No mechanism that is strategy-proof for doctors can always select a matching with the highest match rate among stable matchings). Modify the environment in Example 5 as follows:

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

with everything else unchanged (thus hospitals h_1 and h_2 are in one region and h_3 is in the other, each region has a regional cap of one, and each hospital has capacity of one). Let $(\bar{q}_{h_1}, \bar{q}_{h_2}, \bar{q}_{h_3}) = (1, 0, 1)$. Notice that, given these preferences, there are two stable matchings, namely μ with $\mu_{d_1} = h_1$ and $\mu_{d_2} = h_3$, and μ' with $\mu'_{d_1} = h_3$ and $\mu'_{d_2} = h_2$. Take a mechanism that always selects a matching with the highest match rate among the stable matchings. We show that this mechanism cannot be strategy-proof. Since both μ and μ' have match rate of 1, both can potentially be chosen by the mechanism. Suppose that the mechanism chooses μ . Then, doctor d_2 has an incentive to misreport

³²For instance, see National Resident Matching Market (2010) and Japan Residency Matching Program (2009b).

³³The ease of understanding may not be a persuasive reason for economic theorists to care about the match rates, but it seems to be a crucial issue for market designers. For a mechanism to work well in practice, it is essential that people are willing to participate in the mechanism. To this end, providing information in an accessible manner, as in the form of the match rates, seems to be of great importance.

her preferences: If she reports that hospital h_2 is the only acceptable match, then given the new profile of the preferences, the only stable matching that maximizes the match rate among stable matchings is μ' . Since $\mu'_{d_2} \succ_{d_2} \mu_{d_2}$, doctor d_2 indeed has an incentive to misreport. A symmetric argument can be made for the case in which the mechanism chooses μ' given the true preference profile. Hence, there does not exist a mechanism that is strategy-proof for doctors and always selects a matching with the highest match rate among stable matchings. \square

Despite the above negative results, there are bounds on the match rates in the matchings produced by the flexible deferred acceptance mechanism. More specifically, the following comparison can be made with the JRMP mechanism as well as with the (unconstrained) deferred acceptance algorithm without regional caps:

Theorem 3. *For any preference profile,*

- (1) *Each doctor $d \in D$ weakly prefers a matching produced by the deferred acceptance mechanism to the one produced by the flexible deferred acceptance mechanism to the one produced by the JRMP mechanism.*
- (2) *If a doctor is unmatched in the deferred acceptance mechanism, she is unmatched in the flexible deferred acceptance mechanism. If a doctor is unmatched in the flexible deferred acceptance mechanism, she is unmatched in the JRMP mechanism.*

Notice that part (2) of the above result, which is a direct corollary of part (1), implies that the match rate is weakly higher in the deferred acceptance mechanism than in the flexible deferred acceptance mechanism, which in turn has a weakly higher match rate than the JRMP mechanism.³⁴

In Appendix B, we present simulation results to compare the performances of the deferred acceptance mechanism, the JRMP mechanism, and the flexible deferred acceptance mechanism. The effects highlighted by the above theorem are quite substantial. For instance, almost 600 additional doctors become unmatched under the JRMP mechanism compared to the unconstrained deferred acceptance mechanism (1396 versus 805), but this figure is reduced to about 200 (1010 versus 805) if the flexible deferred acceptance mechanism is used instead of the JRMP mechanism, while satisfying all the regional caps just like the JRMP mechanism does. Moreover, almost 1000 doctors—about 12 percent of the total number of doctors—prefer the outcomes of the flexible deferred acceptance

³⁴For an example in which the deferred acceptance mechanism and the flexible deferred acceptance mechanism differ in terms of match rates, see Example 4 (with an arbitrary target capacity profile). For the flexible deferred acceptance mechanism and the JRMP mechanism, see Example 1.

mechanism *strictly* to the JRMP outcomes, while every doctor prefers the former *weakly* to the latter, as predicted by Theorem 3. See Appendix B for detail of these results, as well as our simulation methods and other results.

Theorem 3 suggests that the flexible deferred acceptance mechanism matches reasonably many doctors. Characterizing stable mechanisms that achieve strategy-proofness for doctors and match “as many doctors as possible,” as well as studying their relationship with the flexible deferred acceptance mechanism, is an interesting open question.

6.4. Nonexistence of Side-Optimal Stable Matchings. There does not necessarily exist a doctor-optimal stable matching (a stable matching unanimously preferred to every stable matching by all doctors). Neither does there exist a hospital-optimal stable matching. To see this point, consider the environment presented in Example 4, and suppose that $(\bar{q}_{h_1}, \bar{q}_{h_2}) = (1, 0)$. There are two stable matchings, μ and μ' specified in Example 4, where only d_1 and h_1 are matched at μ while only d_2 and h_2 are matched at μ' . Clearly, d_1 and h_1 strictly prefer μ to μ' while d_2 and h_2 strictly prefer μ' to μ . Thus there exists neither a doctor-optimal stable matching nor a hospital-optimal stable matching. Moreover, this example shows that there exists neither a doctor-pessimal stable matching nor a hospital-pessimal stable matching in general.

6.5. Generalizations. As mentioned in Section 4, the notion of stability is based on the idea that if the result of a move of a doctor within a region does not equalize the excesses over the target capacities compared to the current matching, it is not deemed as a legitimate deviation. We argued that this is not the only reasonable definition as, for example, it may be natural to suppose that a hospital with capacity 20 is entitled to twice as many doctors (over the target) as a hospital with capacity 10. There may be other criteria, and Kamada and Kojima (2014a) explore the extent to which our analysis goes through. More specifically, they present a model in which each region is endowed with “regional preferences” over the set of distributions of doctors within the region. One special case of the regional preferences is when the region prefers to have more equal number of doctors in excess of targets. They define a stability concept that takes the regional preferences into consideration. They provide a condition on regional preferences under which a generalized version of the flexible deferred acceptance algorithm finds a stable matching as defined more generally and is group strategy-proof for doctors. The criteria mentioned so far satisfy their condition.

Kamada and Kojima (2014a) provide a further generalization: they consider the situation where there is a hierarchy of regional caps. They show that a generalization of the flexible deferred acceptance mechanism induces a stable matching appropriately defined.

This generalization not only has a theoretical appeal but also is practically important. For instance, one could consider a hierarchy of regional caps, say one cap for a prefecture and one for each district within the prefecture. Or the policy maker may desire to regulate the total number of doctors practicing in each specialty in each prefecture.

6.6. Welfare Effects of Picking Orders, Targets, and Regional Caps. The flexible deferred acceptance algorithm follows a certain picking order of hospitals in each region when there are some doctors remaining to be tentatively matched after hospitals have kept doctors up to their target capacities. One issue is how to decide the picking order. One natural conjecture may be that choosing earlier (that is, having an earlier order in the flexible deferred acceptance algorithm) benefits a hospital. As we have mentioned earlier, this would be a problematic property: If choosing earlier benefits a hospital, then how to order hospitals will be a sensitive policy issue to cope with because each hospital would have incentives to be granted an early picking order. Fortunately, the conjecture is not true, as shown in the following example. The example also shows that the different choices of orders result in different stable matchings, thus the choice of an order does matter for the algorithm's outcome.

Example 8 (Ordering a hospital earlier may make it worse off). Let there be hospitals h_1 , h_2 and h_3 in region r_1 , and h_4 in region r_2 . Suppose that $(q_{h_1}, q_{h_2}, q_{h_3}, q_{h_4}) = (2, 1, 1, 1)$ and $(\bar{q}_{h_1}, \bar{q}_{h_2}, \bar{q}_{h_3}, \bar{q}_{h_4}) = (1, 0, 1, 1)$. The regional cap of r_1 is 2 and that for r_2 is 1. Preferences are

$$\begin{aligned} \succ_{h_1}: d_1, d_4, d_2, \quad \succ_{h_2}: d_3, \quad \succ_{h_3}: \text{arbitrary}, \quad \succ_{h_4}: d_2, d_1, \\ \succ_{d_1}: h_4, h_1, \quad \succ_{d_2}: h_1, h_4, \quad \succ_{d_3}: h_2, \quad \succ_{d_4}: h_1. \end{aligned}$$

- (1) Assume that h_1 is ordered earlier than h_2 . In that case, in the flexible deferred acceptance mechanism, d_1 applies to h_4 , d_2 and d_4 apply to h_1 , and d_3 applies to h_2 . d_2 and d_4 are accepted while d_3 is rejected. The matching finalizes.
- (2) Assume that h_1 is ordered after h_2 . In that case, in the flexible deferred acceptance mechanism, d_1 applies to h_4 , d_2 and d_4 apply to h_1 , and d_3 applies to h_2 . But now d_2 is rejected while d_3 is accepted. Then d_2 applies to h_4 , displacing d_1 from h_4 . Then d_1 applies to h_1 . d_1 is accepted, displacing d_4 from h_1 . The matching finalizes.

First, notice that hospital h_2 is better off in case (2). Thus being ordered earlier helps h_2 in this example. However, if h_1 prefers $\{d_1\}$ to $\{d_2, d_4\}$ (which is consistent with the assumption that hospital preferences are responsive with capacities), then h_1 is also made

better off in case (2). Thus being ordered later helps h_1 if she prefers $\{d_1\}$ to $\{d_2, d_4\}$. Therefore, the effect of a picking order on hospitals' welfare is not monotone. \square

A related concern is about what could be called “target monotonicity.” That is, keeping everything else constant, does an increase of the target of a hospital make it better off under the flexible deferred acceptance mechanism? If so, then hospitals would have strong incentives to influence policy makers to give them large targets. The following example shows that target monotonicity is not necessarily true.

Example 9 (Target monotonicity may fail). Consider a market that is identical to the one in Example 8, except that the target of h_1 is now decreased to 0, with the order such that h_1 chooses before h_2 .³⁵ Then h_1 is matched to $\{d_1\}$ under the flexible deferred acceptance mechanism. Therefore, if h_1 prefers $\{d_1\}$ to $\{d_2, d_4\}$, then h_1 is made better off when its target capacity is smaller. \square

Note that both hospital and doctor preferences are heterogeneous in Examples 8 and 9. However, similar failures can occur even when hospitals or doctors have homogeneous preferences. Moreover, splitting or merging regions also has ambiguous welfare effects. These points are made by examples in Appendix C.

By contrast, there exist natural comparative statics results regarding welfare effects of the regional caps.

Proposition 3. *Fix a picking order in the flexible deferred acceptance mechanism. Let $(q_r)_{r \in R}$ and $(q'_r)_{r \in R}$ be regional caps such that $q'_r \leq q_r$ for each $r \in R$. Then the following statements hold.*

- (1) *Each doctor $d \in D$ weakly prefers a matching produced by the flexible deferred acceptance mechanism under regional caps $(q_r)_{r \in R}$ to the one under $(q'_r)_{r \in R}$.*
- (2) *For each region r such that $q_r = q'_r$, the number of doctors matched in r at a matching produced by the flexible deferred acceptance mechanism under regional caps $(q'_r)_{r \in R}$ is weakly larger than at the matching under $(q_r)_{r \in R}$.*

Thus all doctors are made weakly worse off when the regional caps become more stringent. Meanwhile, the number of doctors matched in a region whose regional cap is unchanged weakly increases when the regional caps of other regions become more stringent.

³⁵When the target capacity of h_1 is decreased, the sum of the target capacities becomes strictly smaller than the regional cap (note that such a situation is allowed in our model). If one wishes to keep the sum equal to the regional cap, the example can be modified by increasing the target capacity of h_3 by 1, and the conclusion of the example continues to hold.

This result highlights the tradeoff that a policy maker faces in using the flexible deferred acceptance mechanism: If the regional caps of urban regions are reduced, then the number of doctors matched to other regions weakly increases. However this change weakly decreases the welfare of doctors.

Remark 3. We obtain Proposition 3 as a corollary of a general comparative statics result that we prove in Kamada and Kojima (2014a). This result can be useful in analyzing matching with distributional constraints. For example, the following comparative statics about the JRMP mechanism can be shown using this result.

Proposition 4. *Let $(\bar{q}_h)_{h \in H}$ and $(\bar{q}'_h)_{h \in H}$ be target capacities such that $\bar{q}'_h \leq \bar{q}_h$ for each $h \in H$. Then the following statements hold.³⁶*

- (1) *Each doctor $d \in D$ weakly prefers a matching produced by the JRMP mechanism under target capacities $(\bar{q}_h)_{h \in H}$ to the one under $(\bar{q}'_h)_{h \in H}$.*
- (2) *Each hospital $h \in H$ such that $\bar{q}_h = \bar{q}'_h$ weakly prefers a matching produced by the JRMP mechanism under target capacities $(\bar{q}'_h)_{h \in H}$ to the one under $(\bar{q}_h)_{h \in H}$. Moreover, the number of doctors matched to any such h in the former matching is weakly larger than that in the latter.*

□

6.7. Floor Constraints. The present paper offers a practical solution for the Japanese resident matching problem with regional caps. However, the regional cap may not be an ultimate objective per se, but a means to allocate medical residents “evenly” to different areas. Setting a cap—a ceiling constraint on the number of residents in a region—is an obvious approach to this desideratum, but there may be other possible regulations. For example, one might wonder if setting floor constraints, as opposed to cap constraints, would be an easier and more direct solution. However, there are reasons that floor constraints may be difficult to use. First, even the existence of an individually rational matching that respects floor constraints is not guaranteed. For example, if no doctor finds any hospital in a certain region to be acceptable, then satisfying a positive floor constraint for the region results in an individually irrational matching (doctors matched with hospitals in the region would just reject taking the job). Second, even if an individually rational

³⁶ Since the JRMP mechanism is equivalent to the deferred acceptance mechanism with respect to the target capacities, this result can also be obtained by appealing to the “Capacity Lemma” by Konishi and Ünver (2006), although we obtain these results as corollaries of a more general result shown in Kamada and Kojima (2014a).

matching exists, it is not clear whether a stable matching exists. In fact, an appropriate definition of stability in the presence of floor constraints is unclear.³⁷

7. RELATED LITERATURE

In the one-to-one matching setting, McVitie and Wilson (1970) show that a doctor or a hospital that is unmatched at one stable matching is unmatched in every stable matching. This is the first statement of the rural hospital theorem to our knowledge, and its variants and extensions have been established in increasingly general settings by Gale and Sotomayor (1985a,b), Roth (1984, 1986), Martinez, Masso, Neme, and Oviedo (2000), and Hatfield and Milgrom (2005), among others. As recent results are quite general, it seems that placing more doctors in rural areas has been believed to be a difficult (if not impossible) task, and thus there are few studies offering solutions to this problem. The current paper explores possible ways to offer some positive results.

Roth (1991) points out that some hospitals in the United Kingdom prefer to hire no more than one female doctor while offering multiple positions. Similarly, some schools (or school districts) desire certain diversity characteristics of their incoming classes such as ethnicity and academic performance (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu, 2005; Ergin and Sönmez, 2006). Westkamp (2013) considers a college admission problem in which colleges have admission criteria based on trait-specific quotas. If one regards a region (instead of a hospital) as a single agent in our model, these models and ours appear similar in that an agent in both models has certain “preferences” over distributions more complex than responsive ones. However, the above models are different from ours. For instance, in our model, a distinction should be made between a matching of a doctor to one hospital in a region and a matching of the same doctor to a different hospital in the same region, but such a distinction cannot be even described in the former models. This distinction is essential in the context of residency matching because a doctor may have incentives to deviate by moving between hospitals within a single region. Thus results from these papers cannot be applied in this paper’s environment.

Despite the above-mentioned difficulty, there is a way to make an association between our model and an existing model, namely the model of matching with contracts as defined by Hatfield and Milgrom (2005).³⁸ More specifically, given a matching market with

³⁷Similar points are made in the context of school choice by Ehlers (2010), Ehlers, Hafalir, Yenmez, and Yildirim (2011), Fragiadakis, Iwasaki, Troyan, Ueda, and Yokoo (2012), and Fragiadakis and Troyan (2014).

³⁸Fleiner (2003) considers a framework that generalizes various mathematical results. A special case of his model corresponds to the model of Hatfield and Milgrom (2005), although not all results of the

regional caps, one can define an associated matching model with contracts such that a stable allocation in the latter model induces a stable matching in the former. This correspondence allows us to show some of our results by using properties of the matching with contracts model established by Hatfield and Milgrom (2005), Hatfield and Kojima (2009, 2010), and Hatfield and Kominers (2009, 2010).³⁹ This correspondence is explored by Kamada and Kojima (2014a). On the other hand, it is also worth noting that these models are still different. The reason is that certain types of blocks allowed in the matching model with contracts are considered infeasible in our context. Thus stable allocations in a matching model with contracts can induce only a subset of stable matchings in our model. For this reason, the structural properties of the set of stable matchings in our model are strikingly different from those in the matching model with contracts. For instance, a doctor-optimal stable allocation exists and the conclusion of the rural hospital theorem holds in their model but not in ours.⁴⁰

Abraham, Irving, and Manlove (2007) study allocation of students to projects where a lecturer may offer multiple projects. Both projects and lecturers have capacity constraints. Sönmez and Ünver (2006) analyze a related model in the context of school choice in which there may be multiple school programs in a school building. Motivated by the matching system for higher education in Hungary, Biró, Fleiner, Irving, and Manlove (2010) extend these models to cases in which capacity constraints are imposed on a nested system of sets. Their models are analogous to ours if we associate a lecturer and a project – and a school building and a school, respectively – in their models to a region and a hospital in our model, respectively. However, there are two notable differences. First, they assume that all projects provided by the same lecturer have identical preferences over acceptable students while such a restriction is not imposed in our model.⁴¹ Second, the stability

latter (e.g., those concerning incentives) are obtained in the former. See also Crawford and Knoer (1981) who observe that wages can represent general job descriptions in their model, given their assumption that firm preferences satisfy separability.

³⁹Note that residency matching and school choice with balance requirements mentioned in the last paragraph (Roth, 1991; Abdulkadiroğlu and Sönmez, 2003) can be modeled as special cases of this paper’s model. A related issue appears in the National Resident Matching Program where a hospital may have multiple types of residency positions (Roth and Peranson, 1999; Niederle, 2007).

⁴⁰More specifically, the former result holds under the property called the substitute condition, and the latter under the substitute condition and another property called the law of aggregate demand or size (or cardinal) monotonicity (Alkan, 2002; Alkan and Gale, 2003).

⁴¹In our context, it is important to allow different hospitals in the same region to have different preferences because two hospitals rarely have identical preferences in practice.

concepts employed in their models are different from ours, thus our results do not reduce to theirs even if all hospitals in the same region have identical preferences.

Milgrom (2009) and Budish, Che, Kojima, and Milgrom (2013) consider object allocation mechanisms with restrictions similar to the regional caps in our model. While their models are independent of ours (most notably, their analysis is primarily about object allocation, and stability is not studied), they share motivations with ours in that they consider flexible assignment in the face of complex constraints.

Two companion papers explore further issues in matching under distributional constraints. Kamada and Kojima (2014a) consider a model where welfare and distributional goals of the policy maker are given by “regional preferences.” They define stability in that framework and show that a generalization of the flexible deferred acceptance mechanism is strategy-proof for doctors and finds a stable matching. Furthermore, they show that the stability concept in the present paper is a special case, in which regions prefer to minimize a certain measure of imbalance of doctors across hospitals. Kamada and Kojima (2014b) define alternative stability concepts, strong and weak stabilities. They argue that strong stability is the most natural from a normative viewpoint, but they demonstrate that a strongly stable matching does not always exist. Motivated by the nonexistence of strongly stable matchings, they define a less demanding concept, weak stability. They show that weak stability can be characterized by efficiency and the absence of justified envy toward a matched agent, and that it is still strong enough to exclude some unappealing matchings such as those produced by JRMP. The stability concept in this paper can be shown to imply weak stability.

More broadly, this paper is part of a rapidly growing literature on matching market design. As advocated by Roth (2002), much of recent market design theory advanced by tackling problems arising in practical markets.⁴² For instance, practical considerations in designing school choice mechanisms in Boston and New York City are discussed by Abdulkadiroğlu, Pathak, and Roth (2005, 2009) and Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005, 2006). Abdulkadiroğlu, Che, and Yasuda (2014, 2011), Erdil and Ergin (2008), and Kesten (2009) analyze alternative mechanisms that may produce more efficient student placements than those that are currently used in New York City and Boston. Design issues motivated by an anti-trust lawsuit against the American medical resident matching clearinghouse are investigated by Bulow and Levin (2006), Kojima (2007), Konishi and Sapozhnikov (2008), Niederle (2007), and Niederle and Roth (2003).

⁴²Literature on auction market design also emphasizes the importance of solving practical problems (see Milgrom (2000, 2004) for instance).

A classical resource allocation problem with multi-unit demand has attracted renewed attention in the context of practical course allocation at business schools as studied by Sönmez and Ünver (2010), Budish and Cantillon (2010), and Budish (2010). Initiated by Roth, Sönmez, and Ünver (2004, 2005, 2007), even the organ transplantation problem has become a subject of market design research in recent years. See Roth and Sotomayor (1990) for a comprehensive survey of the matching literature in the first three decades, and Roth (2007a) and Sönmez and Ünver (2008) for discussion of more recent studies.

8. CONCLUSION

We showed that the current matching mechanisms used in various contexts around the world may result in avoidable inefficiency and instability even though some of them are similar to the celebrated deferred acceptance mechanism. We proposed a new mechanism, called the flexible deferred acceptance mechanism. This mechanism is (group) strategy-proof for doctors and generates a stable and (constrained) efficient matching.

With regional caps, defining stability is not a trivial task, and it seems that the right notion depends on the welfare and distributional goals that the policy maker wants to achieve. Hence there may not necessarily exist a unique choice of the mechanism, and there is room for the policy maker to select a particular stable matching based on such goals. Kamada and Kojima (2014a,b) study this issue further. We hope that the present paper serves as a basis for achieving such goals and, more broadly, that it contributes to the general agenda of matching/market design theory to address specific issues arising in practical problems.⁴³

We intentionally refrained from judging the merits of imposing regional caps itself (except for certain welfare results mentioned below). We took this approach because our model does not explicitly include patients or ethical concerns of the general populace, which may be underlying arguments for increasing doctors in rural areas. Similarly, we did not analyze other policies such as subsidies to incentivize residents to work in rural

⁴³Following the analysis of this paper, papers such as Goto, Iwasaki, Kawasaki, Yasuda, and Yokoo (2014); Goto, Hashimoto, Iwasaki, Kawasaki, Ueda, Yasuda, and Yokoo (2014) study matching problems with distributional constraints.

areas.⁴⁴ Instead, we took an approach in the new tradition of market design research, in which one regards constraints such as fairness and repugnance as requirements to be respected and offers solutions consistent with them.⁴⁵ That is, as regional caps seem to be a strong political reality, we believe that it is important to take them as given and provide a practical solution. To help the policy maker make informed judgements about the tradeoffs involved in imposing regional caps, we provided a number of comparative statics results.

The paper opens new avenues for further research topics. First, as mentioned before, strategy-proofness for every agent including hospitals is impossible even without regional caps if we also require stability. However, truth-telling is an approximately optimal strategy even for hospitals under the deferred acceptance mechanism in large markets under some assumptions (Roth and Peranson, 1999; Immorlica and Mahdian, 2005; Kojima and Pathak, 2009). Although such an analysis requires a much more specialized model structure than what this paper has and is outside the scope of this paper, approximate incentive compatibility similar to these papers may hold.

Second, studying more general constraint structures may be interesting. Kamada and Kojima (2014a) analyze the case in which there is a hierarchy of regional caps, and show that a stable matching can be found by a generalization of our flexible deferred acceptance mechanism. By contrast, they show by example that if the regions do not form a hierarchy, a stable matching does not necessarily exist.⁴⁶ A general recipe for defining a stability concept and finding a stable matching by an algorithm is an open question.

Third, it would be desirable to quantify the effect of using the flexible deferred acceptance mechanism instead of the existing mechanisms. As briefly discussed in Section 6.3, we conducted simulations to compare the outcome of the flexible deferred acceptance mechanism with that of the JRMP mechanism and the deferred acceptance mechanism. Since the data of submitted preferences have been unavailable to us so far, we used randomly generated preferences (while using publicly available data to mimic several aspects

⁴⁴This is not because subsidies are not important. In fact, subsidies are used to attract residents to rural areas in many countries such as the United States and Japan. However, there are political pressures to restrict the use of subsidies in the Japanese medical market. Beginning in 2011, for instance, the government will reduce subsidies to residency programs that pay annual salaries of more than 7,200,000 yen (about 85,000 U.S. dollars) to residents. In any case, our analysis is applicable given participants' preferences which reflect subsidies, thus our method can be employed on top of subsidies.

⁴⁵This approach is eloquently advocated by Roth (2007b).

⁴⁶Situations with non-hierarchies may be relevant if, for instance, a government would like to impose caps based on prefectures and specialities independently.

of the real market). However, a better prediction would be possible if we could simulate the performance of our mechanism based on actual data of preferences. In a new project joint with Jun Wako, we have started talking with the matching organizers to discuss such issues and put our mechanism in real use.

Finally, it would be nice to study markets that have similar structures to the ones in this paper. We explored a wide range of applications both in terms of geography (such as Japan, China, the United Kingdom, Hungary, and Ukraine) and in terms of the context (such as medical match, teacher allocation, college admission, and graduate school admission). We expect some general insights will carry over to other applications, while market-specific details may need to be carefully taken into account when we consider different markets in different political or cultural environments.

REFERENCES

- ABDULKADIROĞLU, A. (2005): “College admissions with affirmative action,” *International Journal of Game Theory*, 33(4), 535–549.
- ABDULKADIROĞLU, A., Y.-K. CHE, AND Y. YASUDA (2011): “Resolving Conflicting Preferences in School Choice: the Boston Mechanism Reconsidered,” *American Economic Review*, 101(1), 399–410.
- (2014): “Expanding “Choice” in School Choice,” *American Economic Journal: Microeconomics*, forthcoming.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2005): “The New York City High School Match,” *American Economic Review Papers and Proceedings*, 95, 364–367.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2009): “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the New York City High School Match,” *American Economic Review*, forthcoming.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2009): “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match,” *American Economic Review*, 99, 1954–1978.
- ABDULKADIROĞLU, A., P. A. PATHAK, A. E. ROTH, AND T. SÖNMEZ (2005): “The Boston Public School Match,” *American Economic Review Papers and Proceedings*, 95, 368–372.
- (2006): “Changing the Boston School Choice Mechanism: Strategy-proofness as Equal Access,” Harvard University and Boston College, unpublished mimeo.
- ABDULKADIROĞLU, A., AND T. SÖNMEZ (2003): “School Choice: A Mechanism Design Approach,” *American Economic Review*, 93, 729–747.

- ABRAHAM, D. J., R. IRVING, AND D. MANLOVE (2007): “Two algorithms for the student-project allocation problem,” *Journal of Discrete Algorithms*, 5, 73–90.
- ALCOBA, N. (2009): “Hitting the road to alleviate India’s rural doctor shortage,” *the Canadian Medical Association Journal*,
<http://www.cmaj.ca/cgi/reprint/180/10/E34.pdf>.
- ALKAN, A. (2002): “A class of multipartner matching markets with a strong lattice structure,” *Economic Theory*, 19(4), 737–746.
- ALKAN, A., AND D. GALE (2003): “Stable Schedule Matching Under Revealed Preferences,” *Journal of Economic Theory*, 84, 73–94.
- ASSOCIATION OF MEDICAL STUDENTS (2009): “Public comment on the proposed revision to the residency match (in Japanese),” <http://students.umin.jp/>.
- BIRÓ, P., T. FLEINER, R. W. IRVING, AND D. F. MANLOVE (2010): “The College Admissions problem with lower and common quotas,” *Theoretical Computer Science*, 411, 3136–3153.
- BUDISH, E. (2010): “The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes,” *Harvard University, Unpublished mimeo*.
- BUDISH, E., AND E. CANTILLON (2010): “The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard,” Discussion paper, mimeo.
- BUDISH, E., Y.-K. CHE, F. KOJIMA, AND P. R. MILGROM (2013): “Designing Random Allocation Mechanisms: Theory and Applications,” *American Economic Review*, 103(2), 585–623.
- BULOW, J., AND J. LEVIN (2006): “Matching and Price Competition,” *American Economic Review*, 96, 652–668.
- CRAWFORD, V., AND E. M. KNOER (1981): “Job Matching with Heterogeneous Firms and Workers,” *Econometrica*, 49, 437–450.
- DUBINS, L. E., AND D. A. FREEDMAN (1981): “Machiavelli and the Gale-Shapley algorithm,” *American Mathematical Monthly*, 88, 485–494.
- ECHENIQUE, F., AND J. OVIEDO (2006): “A Theory of Stability in Many-to-Many Matching,” *Theoretical Economics*, 1, 233–273.
- EHLERS, L. (2010): “Controlled School Choice,” Unpublished mimeo.
- EHLERS, L., I. HAFALIR, M. YENMEZ, AND M. YILDIRIM (2011): “School Choice with Controlled Choice Constraints: Hard Bounds versus Soft Bounds,” *mimeo*.
- ERDIL, A., AND H. ERGIN (2008): “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98, 669–689.

- ERGIN, H. (2002): “Efficient Resource Allocation on the Basis of Priorities,” *Econometrica*, 70, 2489–2498.
- ERGIN, H., AND T. SÖNMEZ (2006): “Games of School Choice under the Boston Mechanism,” *Journal of Public Economics*, 90, 215–237.
- FLEINER, T. (2003): “A Fixed-Point Approach to Stable Matchings and Some Applications,” *Mathematics of Operations Research*, 28, 103–126.
- FRAGIADAKIS, D., A. IWASAKI, P. TROYAN, S. UEDA, AND M. YOKOO (2012): “Strategyproof Matching with Minimum Quotas,” *mimeo*.
- FRAGIADAKIS, D., AND P. TROYAN (2014): “Market Design under Distributional Constraints: Diversity in School Choice and Other Applications,” .
- GALE, D., AND L. S. SHAPLEY (1962): “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 69, 9–15.
- GALE, D., AND M. A. O. SOTOMAYOR (1985a): “Ms Machiavelli and the Stable matching problem,” *American Mathematical Monthly*, 92, 261–268.
- (1985b): “Some Remarks on the Stable Matching Problem,” *Discrete Applied Mathematics*, 11, 223–232.
- GOTO, M., N. HASHIMOTO, A. IWASAKI, Y. KAWASAKI, S. UEDA, Y. YASUDA, AND M. YOKOO (2014): “Strategy-proof matching with regional minimum quotas,” in *AAMAS2014*.
- GOTO, M., A. IWASAKI, Y. KAWASAKI, Y. YASUDA, AND M. YOKOO (2014): “Improving Fairness and Efficiency in Matching Markets with Regional Caps: Priority-list based Deferred Acceptance Mechanism,” *mimeo* (the latest version is available at <http://mpra.ub.uni-muenchen.de/53409/>).
- HAERINGER, G., AND V. IEHLE (2010): “Knocking out the supply and sorting in centralized job markets,” *mimeo*.
- HATFIELD, J. W., AND F. KOJIMA (2009): “Group Incentive Compatibility for Matching with Contracts,” *Games and Economic Behavior*, pp. 745–749.
- (2010): “Substitutes and Stability for Matching with Contracts,” *Journal of Economic Theory*, pp. 1704–1723.
- HATFIELD, J. W., AND S. D. KOMINERS (2009): “Contract Design and Stability in Matching Markets,” *mimeo*.
- (2010): “Matching in Networks with Bilateral Contracts,” *mimeo*.
- HATFIELD, J. W., AND P. MILGROM (2005): “Matching with Contracts,” *American Economic Review*, 95, 913–935.

- IMMORLICA, N., AND M. MAHDIAN (2005): “Marriage, Honesty, and Stability,” *SODA*, pp. 53–62.
- IRVING, R., AND D. MANLOVE (2009): “Finding large stable matchings,” *ACM Journal of Experimental Algorithmics*, 14, Section 1 article no. 2.
- JAPAN RESIDENCY MATCHING PROGRAM (2009a): “Result of the Residency Match,” <http://www.jrmp.jp/toukei/2009/2009toukei-5.pdf>.
- (2009b): “Result of the Residency Match,” <http://www.jrmp.jp/koho/2009/2009press.htm>.
- (2010): “Rules for Matching Participation (residency hospitals and university hospitals),” <http://www.jrmp.jp/kiyaku-byouin.htm>.
- KAMADA, Y., AND F. KOJIMA (2014a): “General Theory of Matching under Distributional Constraints,” mimeo.
- (2014b): “Stability Concepts in Matching with Distributional Constraints,” mimeo.
- KESTEN, O. (2009): “School Choice with Consent,” *Quarterly Journal of Economics*, forthcoming.
- KISELGOFF, S. (2012): “Matching Practices for Universities—Ukraine,” <http://www.matching-in-practice.eu/index.php/matching-in-practice/universities/ukraine>.
- KOJIMA, F. (2007): “Matching and price competition: comment,” *The American Economic Review*, 97(3), 1027–1031.
- KOJIMA, F., AND P. PATHAK (2009): “Incentives and stability in large two-sided matching markets,” *American Economic Review*, 99(3), 608–627.
- KONISHI, H., AND M. SAPOZHNIKOV (2008): “Decentralized matching markets with endogenous salaries,” *Games and Economic Behavior*, 64(1), 193–218.
- KONISHI, H., AND M. U. ÜNVER (2006): “Games of Capacity Manipulation in the Hospital-Intern Market,” *Social Choice and Welfare*, 27, 3–24.
- KOZU, T. (2006): “Medical education in Japan,” *Academic Medicine*, 81(12), 1069–1075.
- LIN, H. (2011): “Conference Minutes for Ministry of Education,” <http://www.mpacc.net.cn/ReadNews.asp?NewsID=1689&BigClassName=%CA%D4%B5%E3%D4%BA%D0%A3&SmallClassName=%D4%BA%D0%A3%B6%AF%CC%AC&SpecialID=34>.
- MARTINEZ, R., J. MASSO, A. NEME, AND J. OVIEDO (2000): “Single agents and the set of many-to-one stable matchings,” *Journal of Economic Theory*, 91, 91–105.
- (2004): “On group strategy-proof mechanisms for a many-to-one matching model,” *International Journal of Game Theory*, 33, 115–128.

- MCVITIE, D. G., AND L. WILSON (1970): “Stable Marriage Assignments for Unequal Sets,” *BIT*, 10, 295–309.
- MILGROM, P. (2000): “Putting Auction Theory to Work: The Simultaneous Ascending Auction,” *Journal of Political Economy*, 108(2), 245–272.
- (2004): *Putting Auction Theory to Work*. Cambridge University Press, Cambridge.
- (2009): “Assignment messages and exchanges,” *American Economic Journal: Microeconomics*, 1(2), 95–113.
- MINISTRY OF EDUCATION OF CHINA (2010): “Notification about Master’s Program Enrollment of the year 2010 (in Chinese),” http://www.moe.gov.cn/publicfiles/business/htmlfiles/moe/s180/201007/xxgk_91621.html.
- MINISTRY OF HEALTH, LABOUR AND WELFARE (2005a): “Proposed survey on no-doctor districts (in Japanese),” <http://www.mhlw.go.jp/shingi/2005/01/s0124-11c.html>.
- (2005b): “Report on no-doctor districts (in Japanese),” <http://www.mhlw.go.jp/toukei/list/76-16.html>.
- (2009a): “On opinions expressed in the public comment procedure (in Japanese),” <http://www.mhlw.go.jp/shingi/2009/04/dl/s0423-8a.pdf>.
- (2009b): “On the outcome of the residency match (in Japanese),” <http://www.mhlw.go.jp/topics/bukyoku/isei/rinsyo/matching/dl/091030-1a.pdf>.
- NAMBIAR, M., AND J. BAVAS (2010): “Rudd plan won’t fix rural doctor shortage: RDA,” <http://www.abc.net.au/news/stories/2010/03/16/2846827.htm>.
- NATIONAL RESIDENT MATCHING MARKET (2010): “Results and Data, 2010 Main Residency Match,” <http://www.nrmp.org/data/resultsanddata2010.pdf>.
- NICHOLSON, S. (2003): “Barriers to Entering Medical Specialties,” NBER Working Paper 9649.
- NIEDERLE, M. (2007): “Competitive wages in a match with ordered contracts,” *The American Economic Review*, 97(5), 1957–1969.
- NIEDERLE, M., AND A. ROTH (2003): “Relationship between wages and presence of a match in medical fellowships,” *JAMA*, 290(9), 1153.
- ONISHI, H., AND I. YOSHIDA (2004): “Rapid change in Japanese medical education,” *Medical Teacher*, 26(5), 403–408.
- PATHAK, P., AND T. SONMEZ (2008): “Leveling the playing field: Sincere and sophisticated players in the Boston mechanism,” *The American Economic Review*, 98(4), 1636–1652.

- PEOPLE'S REPUBLIC OF CHINA (2010): "National Long-Term Plan for Education Reform and Development (2010-2020) (in Chinese)," http://www.gov.cn/jrzg/2010-07/29/content_1667143.htm.
- PRONOVOST, P., D. ANGUS, T. DORMAN, K. ROBINSON, T. DREMSIZOV, AND T. YOUNG (2002): "Physician staffing patterns and clinical outcomes in critically ill patients: a systematic review," *JAMA, the Journal of the American Medical Association*, 288(17), 2151.
- ROTH, A. E. (1982): "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–628.
- (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.
- (1985): "The College Admission Problem is not Equivalent to the Marriage Problem," *Journal of Economic Theory*, 36, 277–288.
- (1986): "On the allocation of residents to rural hospitals: a general property of two-sided matching markets," *Econometrica*, 54(2), 425–427.
- (1991): "A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.," *American Economic Review*, 81, 415–440.
- (2002): "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics, Fisher-Schultz Lecture," *Econometrica*, 70, 1341–1378.
- (2007a): "Deferred Acceptance Algorithms: History, Theory, Practice and Open Questions," forthcoming, *International Journal of Game Theory*.
- (2007b): "Repugnance as a Constraint on Markets," *The Journal of Economic Perspectives*, pp. 37–58.
- ROTH, A. E., AND E. PERANSON (1999): "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89, 748–780.
- ROTH, A. E., T. SÖNMEZ, AND U. ÜNVER (2004): "Kidney Exchange," *Quarterly Journal of Economics*, 119, 457–488.
- (2005): "Pairwise Kidney Exchange," *Journal of Economic Theory*, 125, 151–188.
- (2007): "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences," *American Economic Review*, 97, 828–851.

- ROTH, A. E., AND M. A. O. SOTOMAYOR (1990): *Two-sided Matching: a Study in Game-theoretic Modeling and Analysis*. Econometric Society monographs, Cambridge.
- ROTH, A. E., AND X. XING (1994): “Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions,” *American Economic Review*, 84, 992–1044.
- (1997): “Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists,” *Journal of Political Economy*, 105, 284–329.
- SHALLCROSS, T. (2005): “General Medicine in a Remote & Rural Hospital: Problems and Proposals – a personal view,” <http://www.rcpe.ac.uk/publications/rarm/shalcross.pdf>.
- SÖNMEZ, T., AND M. ÜNVER (2006): “School Choice with Aggregate Capacity Constraints,” mimeo.
- (2008): “Matching, Allocation, and Exchange of Discrete Resources,” *Handbook of Social Economics, Elsevier, forthcoming*.
- (2010): “Course bidding at business schools,” *International Economic Review*, 51(1), 99–123.
- SOTOMAYOR, M. A. O. (2004): “Implementation in the many-to-many matching market,” *Games and Economic Behavior*, 46(1), 199–212.
- TALBOTT, C. (2007): “Shortage of Doctors Affects Rural U.S.,” *the Washington Post*, <http://www.washingtonpost.com/wp-dyn/content/article/2007/07/21/AR2007072100432.html>.
- TEO, A. (2007): “The current state of medical education in Japan: a system under reform,” *Medical education*, 41(3), 302–308.
- TOTTORI PREFECTURE (2009): “Emergency request to the proposed revision to the residency match by four prefectural governors (in Japanese),” <http://www.pref.tottori.lg.jp/dd.aspx?menuid=102291>.
- TOYABE, S. (2009): “Trend in geographic distribution of physicians in Japan,” *International Journal for Equity in Health*, 8(1), 5.
- WESTKAMP, A. (2013): “An Analysis of the German University Admissions System,” *Economic Theory*, 53(3), 561–589.
- WONGRUANG, P. (2010): “Diagnosis bleak for rural doctors,” *the Bangkok Post*, <http://www.bangkokpost.com/news/investigation/34445/diagnosis-bleak-for-rural-doctors>.
- ZHAI, Y. (2011): “Development of Professional Education is Left Behind,” *Chinese News of Science*, <http://edu.qq.com/a/20111117/000077.htm>.

APPENDIX A. MATCHING MARKETS IN PRACTICE

A.1. Residency Matching in Japan. In Japan, about 8,000 doctors and 1,500 residency programs participate in the matching process each year. This section describes how this process has evolved and how it has affected the debate on the geographical distribution of residents. For further details of Japanese medical education written in English, see Teo (2007) and Kozu (2006). Also, information about the matching program written in Japanese is available on the websites of the government ministry and the matching organizer.⁴⁷

The Japanese residency matching started in 2003 as part of a comprehensive reform of the medical residency program. Prior to the reform, clinical departments in university hospitals, called *ikyoku*, had de facto authority to allocate doctors. The system was criticized because it was seen to have given clinical departments too much power and resulted in opaque, inefficient, and unfair allocations of doctors against their will.⁴⁸ Describing the situation, Onishi and Yoshida (2004) write “This clinical-department-centred system was often compared to the feudal hierarchy.”

To cope with the above problem a new system, the Japan Residency Matching Program (JRMP), introduced a centralized matching procedure using the (doctor-proposing) deferred acceptance algorithm by Gale and Shapley (1962). Unlike its U.S. counterpart, the National Resident Matching Program (NRMP), the system has no “match variation” (Roth and Peranson, 1999) such as married couples, which would cause many of the good properties of the deferred acceptance algorithm to fail.

Although the matching system was welcomed by many, it has also received a lot of criticisms. This is because some hospitals, especially university hospitals in rural areas, felt that they attracted fewer residents under the new matching mechanism. They argued that the new system provided too much opportunity for doctors to work for urban hospitals rather than rural hospitals, resulting in severe doctor shortages in rural areas. While there is no conclusive evidence supporting their claim, an empirical study by Toyabe (2009) finds that the geographical imbalance of doctors has increased in recent years according to several measures (the Gini coefficient, Atkinson index, and Theil index of the per-capita number of doctors across regions). By contrast, he also finds that the

⁴⁷See the websites of the Ministry of Health, Labor and Welfare (<http://www.mhlw.go.jp/topics/bukyoku/isei/rinsyo/>) and the Japan Residency Matching Program (<http://www.jrmp.jp/>).

⁴⁸The criticism appears to have some justification. For instance, Niederle and Roth (2003) offer empirical evidence that a system without a centralized matching procedure reduces mobility and efficiency of resident allocation in the context of the U.S. gastroenterologist match.

imbalance is lower when residents are excluded from the calculation. Based on these findings, he suggests that the matching system introduced in 2003 may have contributed to the widening regional imbalance of doctors.

To put such criticisms into context, we note that the regional imbalance of doctors has been a long-standing and serious problem in Japan. As of 2004, there were over 160,000 people living in the so-called *mui-chiku*, which means “districts with no doctors” (Ministry of Health, Labour and Welfare, 2005b)⁴⁹ and many more who were allegedly underserved. One government official told one of the authors (personal communication) that the regional imbalance is one of the most important problems in the government’s health care policy, together with financing health care cost. Popular media regularly report stories of doctor shortages, often in a very sensational tone.⁵⁰ There is evidence that the sufficient staffing of doctors in hospitals is positively correlated with the quality of medical care such as lower mortality (see Pronovost, Angus, Dorman, Robinson, Dremsizov, and Young (2002) for instance); thus the doctor shortage in rural areas may lead to bad medical care.

In response to the criticisms against the matching mechanism, the Japanese government introduced a new system with regional caps beginning with the matching conducted in 2009. More specifically, a regional cap was imposed on the number of residents in each of the 47 prefectures that partition the country. If the sum of the hospital capacities in a region exceeds its regional cap, then the capacity of each hospital is reduced to equalize the total capacity with the regional cap.⁵¹ Then the deferred acceptance algorithm is implemented under the reduced capacities. We call this mechanism the Japan Residency Matching Program (JRMP) mechanism. The basic intuition behind this policy is that if

⁴⁹A *mui-chiku* is defined by various criteria such as the ease of access to hospitals, the population, the regularity of clinic openings, and so forth (Ministry of Health, Labour and Welfare, 2005a).

⁵⁰For instance, the Yomiuri Shimbun newspaper, with circulation of over 10,000,000, recently provoked a controversy by its article about the only doctor in Kamikoani-mura village, where 2,800 people live (Yomiuri Shimbun newspaper, 03/19/2010). Although the doctor, aged 65, took only 18 days off a year, she was persistently criticized by some “unreasonable demanding” patients. When she announced that she wanted to quit (which means that the village will be left with no doctor) because she was “exhausted,” 600 signatures were collected in only 10 days, to change her mind.

⁵¹The capacity of a hospital is reduced proportionately to its original capacity in principle (subject to integrality constraints) although there are a number of fine adjustments and exceptions. Although this rule might suggest that hospitals have incentives to misreport their true capacities, the Japanese government regulates how many positions each hospital can offer so that the capacity can be considered exogenous. More specifically, the government decides the physical capacity of a hospital based on verifiable information such as the number of beds in it.

residents are denied from urban hospitals because of the reduced capacities, then some of them will work for rural hospitals.




FIGURE 1. For each prefecture, the total capacity is the sum of advertised positions in hospitals located in the prefecture in 2008. The regional caps are based on the government's plan in 2008 (Ministry of Health, Labour and Welfare, 2009a). Negative values of total capacities in some prefectures indicate the excess amount of regional caps beyond the advertised positions.

The magnitude of the regional caps is illustrated in Figure 1. Relatively large reductions are imposed on urban areas. For instance, hospitals in Tokyo and Osaka advertised 1,582 and 860 positions in 2008, respectively, but the government set the regional caps of 1,287 and 533, the largest reductions in the number of positions. The largest reduction in proportion is imposed on Kyoto, which offered 353 positions in 2008 but the number is dropped to 190, a reduction of about 46 percent. Indeed, the projected changes were so large that the government provided a temporary measure that limits per-year reductions within a certain bound in the first years of operation, though the plan is to reach the planned regional cap eventually. In total, 34 out of 47 prefectures are given regional caps smaller than the numbers of advertised positions in 2008.

The new JRMP mechanism with regional caps was used in 2009 for the first time. The government claims that the change alleviated the regional imbalance of residents:

It reports that the proportion of residents matched to hospitals in rural areas has risen to 52.3 percent, an increase of one percentage point from the previous year (Ministry of Health, Labour and Welfare, 2009b).⁵² However, there is mounting criticism to the JRMP mechanism as well. For instance, a number of governors of rural prefectures (see Tottori Prefecture (2009) for instance) and a student group (Association of Medical Students, 2009) have demanded that the government modify or abolish the JRMP mechanism with regional caps.⁵³ Among other things, a commonly expressed concern is that the current system with regional caps causes efficiency losses, for instance by preventing residents from learning their desired skills for practicing medical treatments.

In the main text we formalized the JRMP mechanism (Section 3), explored its properties (Example 1, Remark 1, and Proposition 4), and compared it to the flexible deferred acceptance mechanism (Theorem 3). Our analysis suggests that the current JRMP mechanism needs to be changed to the flexible deferred acceptance mechanism.

A.2. Chinese Graduate Admission. This section describes the Chinese graduate admission in detail, and formally shows that the mechanism may result in an unstable and inefficient matching.⁵⁴

A.2.1. Institutional Background. Chinese society is changing rapidly, and it is widely believed that there is need for more workers with professional master’s degrees.⁵⁵ However, professional master’s degrees have traditionally been regarded as inferior to academic master’s degrees by many.⁵⁶ And there are not as many students in professional master’s programs as the government aims.⁵⁷

⁵²Ministry of Health, Labour and Welfare (2009b) defines “rural areas” as all prefectures except for 6 prefectures, Tokyo, Kyoto, Osaka, Kanagawa, Aichi, and Fukuoka, which have large cities.

⁵³Interestingly, even regional governments in rural areas such as Tokushima and Tottori opposed to the JRMP mechanism. They were worried that since the system reduces capacities of *individual hospitals* in the region and some of them could hire more residents, it can reduce the number of residents allocated in the regions even further. This feature - inflexibility of the way capacities are reduced - is one of the problems of the current JRMP mechanism that we try to remedy by our alternative mechanism.

⁵⁴We greatly benefited from discussing Chinese graduate school admission with Jin Chen.

⁵⁵Ministry of Education of China (2010) states that “[the education authority and graduate schools] should put emphasis on the promotion of education for advanced professionals, especially full-time professional master’s degree.”

⁵⁶Chinese government has been emphasizing that the only differences between professional and academic master’s programs are in the types and goals of education, and not in standards, but the reputation of a professional master’s degree is still not as good as an academic one (Zhai, 2011).

⁵⁷China’s master-level education system traditionally emphasized academic training, rather than professional training. However, most graduates from master’s programs pursue a professional career instead

To address this issue, Chinese government started a new regulation to increase enrollment in professional master’s programs in 2010 (People’s Republic of China, 2010). More specifically, the government began to impose constraints on the total number of academic master students, while increasing the number of professional master students. To achieve this goal, the government decided to reduce the available seats of each academic master’s program by about 25 percent by 2015, while increasing capacities for professional programs.⁵⁸

Although Chinese graduate school admission is different from Japanese residency match in many ways, there is a clear isomorphism between the structures of the problems that these markets are faced with. Just as there is demand for increasing resident allocation in rural Japan, there is demand for increasing professional master students in Chinese graduate education. Moreover, in both Japanese and Chinese cases, the feasibility requirements are placed on the total numbers of allocations for a subset of institutions (hospitals in each prefecture in the Japanese case, and the academic master’s programs in the Chinese case). Lastly, when implementing the requirement, both governments place rigid restrictions on the allowed seats of each institution (each hospital in Japan, and each graduate school in China).

Remark 4. As mentioned in Section 6.1.1 of the main text, the matching mechanism for Chinese graduate school admission has additional rounds. First, there is a recommendation-based admission for high achievers before the main round. Second, there is a round called an “adjustment process” for those who have not been matched by the end of the main round. We do not formally analyze these rounds because they are not directly related to the issue of distributional constraints, and processes similar to them and the associated problems have been analyzed by other works.⁵⁹ Nevertheless, we describe these rounds for completeness.

of an academic one: in 2009, for instance, enrollment for master’s programs was around 415,000 while that for PhD programs was around 60,000 (in China, the master’s degree is similar to M.Phil. in countries like the U.K. in that a student seeking PhD first attend a master’s program). Between 2008 and 2011, the proportion of professional masters’ enrollment has increased from just 7 percent to 30 percent of the total enrollment. Ministry of Education aims at 50 percent by 2015 (Lin, 2011).

⁵⁸To achieve this goal gradually, the government plans to reduce the number of seats by about 5 percent every year until 2015 (to our knowledge, the government has not disclosed whether it will continue imposing the reduction beyond 2015).

⁵⁹Abdulkadiroğlu, Pathak, and Roth (2005); Abdulkadiroğlu, Pathak, and Roth (2009) study New York City’s high school match. In NYC, top 2 percent students are automatically admitted to certain schools if they prefer, similarly to top performers who can be admitted to some schools in China’s recommendation-based admission. Roth and Xing (1994, 1997) study labor markets that proceed in real

In the recommendation-based admission, students are recommended to graduate schools directly even before the graduate entrance examination that everyone else should take for admission. This round works as a shortcut for excellent students to enter graduate schools. Most recommended students are admitted to and attend the same university's graduate school where he or she attended college, and few of the recommended students transfer to another graduate school. This round is decentralized, and how to evaluate and admit students is largely up to each school, and thus the admission policy varies from school to school.

Students who were unmatched in the recommendation-based admission and the main round, as well as graduate schools which were not full in these rounds, enter the adjustment process. This round proceeds in real time, and each student maintains an application list which consists of at most two schools at any moment during this process. More specifically, at the beginning of this round, each student enters at most two schools into an online system.⁶⁰ Each school sees students who listed it, decides whether to give each applicant the chance for interview or not, and sends out interview invitations to them. Upon receiving an interview invitation, a student chooses whether to accept it or not. If a student fails to receive an interview invitation in 48 hours or is rejected after the interview from a school, the student can remove the school and add another school into her online application list. Once changed, a student must keep the new school in the list for at least 48 hours unless it interviews her. On the interview day, a graduate school admits or rejects students. If the capacity becomes full, the graduate school completes the admission. Each student can confirm admission from at most one school. Once she confirms, she exits the matching process.

The adjustment process happens in real time: In 2012, for instance, the process was in session from April 1st to May 5th. Given this time constraint, it is widely believed that students and schools contact each other ahead of the official adjustment process.⁶¹ \square

A.2.2. *Formal Analysis.* Let us use the same notation as in the main text, although now we call h a program instead of a hospital, and d a student instead of a doctor. Further assume that the set of all programs is partitioned into the set of all academic programs r

time, and highlight the time constraint and associated strategic behavior and inefficiency of the resulting matching.

⁶⁰The application is maintained at <http://yz.chsi.com.cn/>.

⁶¹See "Experts: 2012 graduate school entrance examination, advice and strategies for adjustment" from *Kuakao Education*, a counseling agency in China for students applying for graduate schools, available at <http://yz.chsi.com.cn/kyzx/fstj/201203/20120322/293594478.html>.

and the set of all professional programs r' . Throughout, assume $q_{r'} > \sum_{h \in H_{r'}} q_h$ so that the cap for professional programs is not binding.

We define the main round of the Chinese graduate admission formally.⁶² As described in the main text, given the cap q_r , the main round of Chinese graduate admission runs as follows (we describe the mechanism for a general value of q_r , although in China q_r is an integer close to 75 percent of the sum of academic program capacities). Set $\bar{q}_h \leq q_h$ for each h in such a way that $\sum_{h \in H_r} \bar{q}_h \leq q_r$ (in Chinese graduate admission, \bar{q}_h is an integer that is at most 75 percent of q_h for each $h \in H_r$). Each student applies to at most one program. Given the set of applicants, each program h accepts its most preferred students up to its target capacity \bar{q}_h and rejects everyone else. All matchings are final.

We model the behavior in this mechanism by considering the following two-stage extensive-form game. In the first stage, students simultaneously apply to programs, one for each student. Then in the second stage, each program admits students from those who applied to it up to its target capacity. In this game, the following result holds:

Result 1. *Suppose that $q_r > \sum_{h \in H_r} q_h$ and $\bar{q}_h = q_h$ for all h . Then the set of the pure-strategy subgame-perfect equilibrium outcomes in the game induced by the Chinese graduate admission coincides with the set of stable matchings.*

Proof. When $q_r > \sum_{h \in H_r} q_h$ and $\bar{q}_h = q_h$ for all h , the stability concept of this paper is equivalent to the standard stability concept (as in Roth and Sotomayor (1990) for example). By Sotomayor (2004) and Echenique and Oviedo (2006), the set of subgame perfect equilibrium outcomes of this game is equivalent to the set of stable matchings in the standard sense. These two observations complete the proof. \square

Thus if there is no binding cap on academic programs, then the equilibrium outcomes are stable and hence efficient. When the cap on academic programs is binding as in the Chinese admission mechanism, however, neither of these good properties hold even for equilibrium outcomes. The following example, which is an adaptation of Example 1, is such a case.

Example 10 (Equilibrium Outcomes under the Chinese Mechanism Can Be Unstable and Inefficient). The “regional cap” for academic programs r is $q_r = 10$. There are two academic programs h_1 and h_2 and no professional program. Each program h has a capacity

⁶²The description is based on the website of the “National Graduate Admissions Information Network” (<http://yz.chsi.com.cn/>), which provides information on graduate admission and host online applications for graduate schools.

of $q_h = 10$. Let the target capacities be $\bar{q}_{h_1} = \bar{q}_{h_2} = 5$. There are 10 students, d_1, \dots, d_{10} . Preference profile \succ is as follows:

$$\begin{aligned} \succ_{h_i}: d_1, d_2, \dots, d_{10} \quad \text{for } i = 1, 2, \\ \succ_{d_j}: h_1 \text{ if } j \leq 3 \quad \text{and} \quad \succ_{d_j}: h_2 \text{ if } j \geq 4. \end{aligned}$$

It is easy to see that the only pure-strategy subgame perfect equilibrium outcome is

$$\mu = \begin{pmatrix} h_1 & h_2 & \emptyset \\ d_1, d_2, d_3 & d_4, d_5, d_6, d_7, d_8 & d_9, d_{10} \end{pmatrix}.$$

Consider a matching μ' defined by,

$$\mu' = \begin{pmatrix} h_1 & h_2 \\ d_1, d_2, d_3 & d_4, d_5, d_6, d_7, d_8, d_9, d_{10} \end{pmatrix}.$$

Since the cap for academic programs is still respected, μ' is feasible. Moreover, every student is weakly better off with students d_9 and d_{10} being strictly better off than at μ . Hence we conclude that the Chinese mechanism can result in an inefficient matching. We also note that μ is not stable: For example, program h_2 and student d_9 constitute a blocking pair while the cap for r is not binding. \square

A.3. College Admission in Ukraine. A problem similar to Japanese residency match and Chinese graduate school admission is found in college admission in Ukraine as well. In Ukraine, some of the seats are financed by the state, while other “open-enrollment” seats require that students pay tuition (Kiselgof, 2012).⁶³ There is a cap on the number of state-financed seats, apparently as there is a limit on the budget that can be used to finance college study. The government implements the cap on the number of state-financed seats by imposing a cap on each program as in Japanese residency match and the Chinese graduate admission. Although the specific mechanism the Ukrainian college admission system uses is different from JRMP and Chinese graduate school admissions (see Kiselgof (2012) for detail), instability and inefficiency because of the constraints similarly result.

A.4. Medical Matching in the United Kingdom.

⁶³A similar cap on the number of state-financed college seats exists in Hungarian college admission as well. The situation is somewhat different here, however, as ranking by colleges are based on a common exam and hence is common for different university programs on the same subject. Biró, Fleiner, Irving, and Manlove (2010) propose an elegant matching mechanism in such an environment.

A.4.1. *Institutional Backgrounds.* In recent years, how to organize medical training has been a contentious topic in the U.K., and the system has undergone a number of drastic changes. This section describes the current system, whose basic structure was set up in 2005.⁶⁴

In order to practice medicine in the U.K., graduates from medical schools must undertake two years of training. The arrangement is called the Foundation Programme, and places about 7,000 medical school graduates to training programs every year.⁶⁵ In the first round of the matching scheme of the Foundation Programme, applicants are matched to one of 25 “foundation schools” by a national matching process. A foundation school is a consortium made of medical schools and other organizations, and each foundation school largely corresponds to a region of the country. Upon being matched to a foundation school, students are matched to individual training programs within that foundation school in the second round of the matching process. In these processes, applicants are assigned a numerical score, which may result in ties. Until 2011, the Boston mechanism (also known as the “first-choice-first” mechanism in the U.K.) was in use, based on the numerical score and a random tie-breaking.⁶⁶ Beginning in 2012, a serial dictatorship algorithm based on the score with random tie-breaking is used for allocation to foundation schools.⁶⁷ It is up to individual foundation schools as to how they match their assigned applicants to programs in their region. In Scotland, for example, a stable mechanism was in use to allocate students to programs within the region until serial dictatorship based on applicant scores and random tie-breaking replaced it in 2010.⁶⁸

For our purposes, an especially interesting point is that the mechanism used in U.K. medical matching has two rounds, in which students are assigned to a region first, and then to a program within their assigned region. Although the specific mechanisms used in the second round vary from region to region, the United Kingdom as a whole uses a two-round mechanism.

⁶⁴We are grateful to Peter Biró, Rob Irving, and David Manlove for answering our questions about medical match in the U.K.

⁶⁵Some institutional details and the statistics reported here can be found in the Foundation Programme’s website, especially in its annual reports: see for example its 2011 annual report at http://www.foundationprogramme.nhs.uk/download.asp?file=Foundation_Programme_Annual_Report_Nov11_FINAL.pdf.

⁶⁶See Abdulkadiroğlu and Sönmez (2003) who study the Boston mechanism in the school choice context.

⁶⁷The algorithm is described at <http://www.foundationprogramme.nhs.uk/pages/medical-students/faqs#answer39>

⁶⁸There are applicants who participate as couples, and the algorithms handle these couples in certain manners. See Irving and Manlove (2009) for details.

A.4.2. *Formal Analysis.* As indicated above, mechanisms used in U.K. medical match (and Scottish teacher matching as mentioned in Appendix A.5) have many variations, but their basic structure is common in the sense that applicants are matched by a two-round procedure. Formally, we consider a mechanism in which applicants are matched to a region (up to its regional cap) in the first round, and then they are matched to a hospital within the assigned region in the second round. For concreteness we focus on the mechanism in which serial dictatorship is used in both rounds, but the main conclusions can be obtained for other mechanisms as well (we describe the details later in this section).

The first example shows that the outcome of this two-round mechanism may be unstable.

Example 11. There are two regions r_1 and r_2 with regional caps $q_{r_1} = q_{r_2} = 2$. There are two hospitals h_1 and h_3 in r_1 while there is one hospital h_2 in r_2 . Hospital capacities are $q_{h_1} = 1$, $q_{h_2} = 2$, and $q_{h_3} = 1$. Suppose that there are 3 doctors, d_1 , d_2 , and d_3 . Preference profile \succ is as follows:

$$\succ_{h_i}: d_1, d_2, d_3 \text{ for all } i,$$

$$\succ_{d_j}: h_1, h_2, h_3 \text{ for all } j.$$

And let us assume that, in both rounds, the serial dictatorship is used with respect to the ordering d_1 , d_2 , and d_3 . That is, d_1 is matched to her most preferred region (or hospital), d_2 is matched to the most preferred region (or hospital) that are still available, and so on. Note that we assume that hospital preferences are common and coincide with the applicant ordering in the serial dictatorship. This assumption is meant to make stability as easy to obtain as possible, because if serial order and hospital preferences are different, it is almost trivial to obtain unstable matchings (indeed, under the original serial dictatorship, the resulting matching is stable under this assumption, but not otherwise). Assume that each doctor prefers r_1 most and r_2 second.⁶⁹ Let target capacities be arbitrary.

In the first round of this mechanism, d_1 and d_2 are matched to r_1 , while d_3 is matched to r_2 . In the second round, d_1 is matched to her first choice h_1 , d_2 is matched to h_3 , which is the only remaining hospital in region r_1 , and d_3 is matched to hospital h_2 , resulting in

$$\mu = \begin{pmatrix} h_1 & h_2 & h_3 \\ d_1 & d_3 & d_2 \end{pmatrix}.$$

This matching μ is unstable, because d_2 and h_2 form a legitimate blocking pair: d_2 prefers h_2 to its match h_3 and h_2 prefers d_2 to its match d_3 . \square

⁶⁹Such reported preferences may arise if, for instance, she believes that there is nonzero probability to be matched with h_1 and her cardinal utility from h_1 is sufficiently high.

Although we phrased the above example in the context of a mechanism both of whose rounds employ the serial dictatorship, the same point can be made for other two-round mechanisms. Consider, for instance, the case in which the first-round procedure is a Boston mechanism (as was the case in the U.K. until 2011) based on the above ordering, while the second round is a serial dictatorship.⁷⁰ In the above market, under this procedure d_1 and d_2 are still matched to region r_1 in the first round, and then d_2 is matched to h_3 in the second round, leading to the same matching μ of Example 11, thus to instability. This observation shows that the problem of instability is not restricted to the detail of the current mechanism using serial dictatorship in both rounds, but rather a general feature of two-round systems that have been the basic framework of the U.K. medical match.

The following example shows that the matching resulting from the U.K. medical match can be inefficient.

Example 12. There are two regions r_1 and r_2 with regional caps $q_{r_1} = 2$ and $q_{r_2} = 1$. There are two hospitals h_1 and h_3 in r_1 while there is one hospital h_2 in r_2 . Each hospital h has a capacity of $q_h = 1$. Suppose that there are 3 doctors, d_1, d_2, d_3 . Preference profile \succ is as follows:

$$\begin{aligned} \succ_{h_i} &: d_1, d_2, d_3 \text{ for all } i, \\ \succ_{d_1} &: h_1, h_2, h_3, \\ \succ_{d_2} &: h_1, h_2, \\ \succ_{d_3} &: h_1, h_3, h_2. \end{aligned}$$

As before, let us assume that the serial dictatorship with ordering d_1, d_2 , and d_3 is used in both rounds. Assume further that doctors' preferences over the regions are induced in the manner specified in Example 11.

At the first round of this mechanism, d_1 and d_2 are matched to r_1 , while d_3 is matched to r_2 . In the second round, d_1 is matched to her first choice h_1 , d_2 is unmatched, and d_3 is matched to hospital h_2 , resulting in matching

$$\mu = \begin{pmatrix} h_1 & h_2 & h_3 & \emptyset \\ d_1 & d_3 & \emptyset & d_2 \end{pmatrix}.$$

⁷⁰Recall that the first round algorithm was changed from the Boston mechanism to the serial dictatorship only beginning in 2012 in the U.K. medical match, while serial dictatorship was already in use in Scotland in 2009.

Consider a matching μ' defined by,

$$\mu' = \begin{pmatrix} h_1 & h_2 & h_3 \\ d_1 & d_2 & d_3 \end{pmatrix}.$$

The latter matching satisfies all regional caps and Pareto dominates the former matching μ : d_1 and h_1 are indifferent between μ and μ' , while every other agent is made strictly better off at μ' than at μ . Therefore the matching μ is inefficient. \square

The last drawback of the two-round mechanism we point out involves incentives. Serial dictatorship is strategy-proof, and this property is often regarded as one of the main advantages of this mechanism. In a two-round mechanism, however, there exists no dominant strategy even if both rounds employ serial dictatorship.

Example 13. Consider the market defined in Example 11. Note that it is a weakly dominant strategy to report true preferences in any subgame of the second round, i.e., once doctors are matched to regions, so assume that all doctors report true preferences in the second round. Suppose that doctors report preferences over regions as in Example 11. Then doctor d_2 is assigned to her third choice hospital h_3 . However, if d_2 reports region r_2 to be her most preferred region while no other doctor changes his reported preference, then she is matched to h_2 , which is the optimal matching possible for any of her reported preferences. In other words, reporting r_2 as the most preferred region is a best response while reporting r_1 is not. Next, consider a report of d_1 that reports r_2 to be his most preferred region. Then d_2 is matched to h_1 if she reports r_1 to be her most preferred region while she is matched to her less preferred hospital h_2 if she reports r_2 to be her most preferred region. In other words, reporting r_1 as the most preferred region is a best response while reporting r_2 is not. Therefore there is no dominant strategy.⁷¹ \square

A.5. Probationary Teacher Matching in Scotland. Another problem of interest is the matching of new teachers (called probationary teachers) to schools. Teachers in Scotland need to get a training as probationers for one year. The General Teaching Council for Scotland (GTCS) runs a procedure called the Teacher Induction Scheme, which allocates probationary teachers to training posts in Scottish schools.⁷² Scotland has 32 local authorities, and probationary teachers and these local authorities are matched in the first round of the mechanism. Information about the algorithm used is unavailable to our knowledge, but some documents suggest that a slight variant of the random serial

⁷¹It is trivial, and hence omitted, to show that reporting no region to be acceptable is weakly dominated, so the above argument is enough to establish the claim.

⁷²See <http://www.gtcs.org.uk/home/students/teacher-induction-scheme-faq.aspx>.

dictatorship is used.⁷³ Then each local authority decides which probationers matched to it are sent to which schools under its control, and that round occurs subsequently to the first round. The mechanism that local authorities use in this round is up to each local authority, and appears to vary widely from one local authority to another.

This scheme has a lot in common with the previous example. As in the U.K. medical match, the matching clearinghouse first assigns teachers to a local authority, who then assigns them to schools under its control.

⁷³“Teacher Induction Scheme 2008/2009,” <http://www.scotland.gov.uk/Resource/Doc/200891/0053701.pdf> states that “A computer system will match and allocate students to local authorities using each local authority’s vacancy list and student’s preference list. You will be chosen at random and matched against your five preferences, beginning with your first preference. Where an appropriate vacancy is unavailable, you will be matched against your second preference, and so on until an appropriate match is found.” As indicated above, a probationary teacher is asked to only rank 5 local authorities, unlike the exact random serial dictatorship. Another complication is that a student can alternatively tick a preference waiver box indicating that they are happy to work anywhere in Scotland. Those who choose the preference waiver option are paid additional compensation.

ONLINE SUPPLEMENTARY APPENDIX

APPENDIX B. SIMULATION METHODS AND RESULTS FOR THE CASE OF THE JAPANESE RESIDENCY MATCHING PROGRAM

In this appendix we provide results of simulations using the data on Japanese medical residency match. As we have discussed in the main text of the paper (and in the appendix), there are a number of instances around the world where distributional constraints are imposed. We chose the Japanese case as our principal target for simulation for various reasons: First, all the data we used in our simulation are available online for free, so one can replicate our simulation results easily. Second, the Japanese medical match is highly centralized, so the conclusions from the simulation results are more meaningful than in some other applications where a (sometimes nontrivial) part of the matching procedure is not explicitly specified. Third, compared to the current practice (the JRMP mechanism), our proposal (the flexible deferred acceptance (FDA) mechanism) has an advantage in terms of efficiency and stability, while its effect on the regional balance of doctors is ambiguous from the theoretical perspective. Thus simulations are useful. Finally, we have been talking to the Japanese government officials about the FDA mechanism, and by quantifying the trade-off that we have just mentioned, we have a better chance of persuading them to use the FDA mechanism. Since the practicality of the theory is one of our main goals, we view simulations as a useful analysis to implement.

As we reviewed in the paper, the JRMP mechanism was introduced in 2009, and after that, hospitals have been asked to gradually decrease their respective capacities, to eventually match the total capacities to the planned regional cap. Since the government publicizes only the reduced capacities, we use the hospital capacities in the data just before 2009. More specifically, we use the data from 2007, because the government specified the regional cap based on that year's data. For consistency, we also use other parts of information from the data of the same year. All the data used here can be obtained at the webpage of Japanese Medical Residency Program (<http://www.jrmp.jp>). The data are in Japanese.

B.1. Simulation Method. We obtain the data of regional caps of all 47 prefectures, the capacity and the region (prefecture) of each hospital.

Using these data, we set the “target capacity” for each hospital following the description in footnote 13 of the paper. That is, if the sum of the advertised positions in the hospital's region is no more than the regional cap, then the hospital's target capacity is equal to the number of its advertised positions; Otherwise, the target is given by the advertised

number times the fraction of the regional cap over the total number of the advertised positions in the region.⁷⁴

The market size. In the simulation, the numbers of doctors and hospitals are 8,291 and 1,357, respectively, which are the number of doctors and hospitals that actually submitted preference lists in the Japanese residency match in 2007.

Preference lists. We do not have the actual data on submitted or true preferences of the doctors and hospitals, so given the above information we generated preferences and ran the simulation. Fortunately, various public data enabled us to set parameters that mimic the Japanese case, as we explain below.

- (1) **Doctors.** We obtain the data on the distribution of the length of preference lists of doctors (i.e., the number of hospitals listed in the submitted preference list of each doctor) up to length 8. In the data, the number of doctors who listed k hospitals is not available for $k \geq 9$, while the total number of doctors who listed 9 or more hospitals is available. We also obtain from the data the average number of hospitals listed, which is 3.48.

For the doctors who list 9 or more hospitals in their preference list, we used the truncated exponential distribution such that (i) the number of doctors who list k hospitals is 0.6 times the number of doctors who list $k - 1$ doctors, for $k = 10, 11$, (modulo integer constraints) (ii) the number of doctors at length 9 is adjusted so that the average number of listed hospitals is 3.48, the average from the data, and (iii) the maximum length is 15. With this specification the number of doctors at length 9 is smaller than the number for length 8.

The data describe, for each hospital, the number of doctors who listed that hospital in their preference lists. Using these data, for each hospital h , we define p_h to be the number of doctors who listed it in their preference list divided by the sum of all those numbers across all hospitals (so that it becomes probability, i.e., the numbers sum up to one). Then each doctor with preference list length k independently draws hospitals based on this distribution $(p_h)_{h \in H}$, repeatedly k times without replacement, listing her first pick as the first choice, her second pick as the second choice, and so on.⁷⁵

⁷⁴If the resulting number is not an integer, then we round the numbers to one of the adjacent integers in such a way that the sum of the target capacities in the region is equal to the regional cap.

⁷⁵This manner of preference generation is used in a number of matching papers, such as Kojima and Pathak (2009).

- (2) **Hospitals.** Each hospital ranks doctors uniform randomly, viewing every doctor acceptable.⁷⁶

Remark 5. There could be alternatives for this method. For example we could have each hospital always rank a doctor with a shorter preference-list length higher than the one with a longer length and those who have the same length are ranked uniformly randomly: this might as well be closer to the data, because those doctors who rank only a small number of hospitals may be doing so because they are confident that hospitals rank them high. However, without better data on the doctor preference, we did not have better foundation for conducting such biased data generations, and hence stuck to the uniformly-random data generation. We hope that the Japanese government disclose anonymous data of preference lists.

B.2. Simulation Results. Medical matching has two sides, namely doctors and hospitals. In the context of the Japanese medical match, another important issue is the distributional balance of doctors in different regions. Therefore, in the following we discuss the simulation results pertinent to welfare of doctors and hospitals, and then discuss the distributional consequences of the mechanisms across regions.

B.2.1. *Doctors and Hospitals.*

- (1) **The number of matched doctors.** Figure 2 shows that the cost of using the JRMP mechanism is quite significant: almost 600 doctors who are matched in the unconstrained deferred acceptance (DA) mechanism become unmatched in JRMP (1396 versus 805). However, much of the negative effects can be alleviated if we switch to FDA: the number of additional doctors who become unmatched compared to the DA is about 205 (1010 versus 805), which is only about one third of 600, the corresponding number for the JRMP mechanism. Importantly, FDA achieves this improvement while satisfying all the regional caps just like the JRMP does, while DA does not respect regional caps.

⁷⁶In our simulation code, each hospital actually orders only doctors who find the hospital acceptable. This is without any consequence because none of the algorithms we consider in this paper is affected by whether a doctor who finds a hospital unacceptable is acceptable to the hospital.




FIGURE 2. The numbers of matched doctors under different mechanisms

- (2) **The number of doctors who are made strictly better off from JRMP to FDA, from FDA to DA, and from JRMP to DA.**

| From/To | DA | FDA | JRMP |
|---------|---------------|--------------|------|
| DA | 0 | 0 | 0 |
| FDA | 606 (7.3 %) | 0 | 0 |
| JRMP | 1547 (18.7 %) | 996 (12.0 %) | 0 |

TABLE 1. The number of doctors who are made strictly better off from JRMP to FDA, from FDA to DA, and from JRMP to DA.

Table 1 demonstrates the numbers of doctors who become strictly better off by changing the mechanism from the one in the row to the one in the column. For example, 996 doctors become strictly better off by changing the mechanism from the JRMP mechanism to the FDA mechanism. As Theorem 3(1) predicts, there are no doctors who become strictly better off from FDA to JRMP, or from DA to the other two mechanisms, and this prediction is confirmed by the zeros in Table 1.

The theorem predicts that doctors are *weakly* better off by the change from JRMP to the other two mechanisms and from FDA to DA, but it does not pin

down how many doctors become *strictly* better off, and in general it is hard to obtain an analytical result on strict improvement without making additional assumptions on the preference distributions. This is one of the main motivations for our simulations. Since DA is unconstrained with respect to the number of doctors that can be matched to each region, large magnitudes of improvement from the other two mechanisms to DA is expected. Even so, the simulation result shows that the improvement from FDA to DA is moderate (7.3%). The result also shows that the effect of the change from JRMP to FDA—both of which are constrained by the regional cap—is large (12.0%). In view of the fact that the DA gives a (loose) upper bound of what FDA can possibly achieve, the simulation result demonstrates FDA’s surprisingly large improvements in doctor welfare upon the current JRMP mechanism.

(3) **Cumulative number of doctors matched to their k -th or better choices**




FIGURE 3. Cumulative number of doctors matched to their k -th or better choices.

In the above graph, the horizontal axis describes the ranking, and the vertical axis describes the number of doctors. For each of the mechanisms, we plot the

cumulative number of doctors who are matched to their k -th or better choices, for each value k in the horizontal axis.⁷⁷

The graph confirms our prediction in Theorem 3(1). That is, the doctors are better off under DA than under FDA and under FDA than under JRMP. As we discussed in (2), the theorem does not predict magnitudes of the improvement, and this motivates simulations. Although regional caps certainly result in worse outcomes for doctors, more than a half of the loss caused by JRMP compared to the unconstrained DA can be avoided once FDA is used, even though regional caps are satisfied in FDA just as in JRMP. The effect is large: for example, about 500 more doctors are matched to their first choices under FDA compared to JRMP.

(4) **The number of hospitals that are matched to more doctors in DA, FDA, and JRMP.**

| From\To | DA | FDA | JRMP |
|---------|--------------|--------------|--------------|
| DA | 0 | 138 (10.2%) | 222 (16.4%) |
| FDA | 104 (7.7 %) | 0 | 158 (11.6 %) |
| JRMP | 366 (27.0 %) | 376 (27.7 %) | 0 |

TABLE 2. The number of hospitals that are matched to more doctors in DA, FDA, and JRMP.

The above table describes the number of the hospitals that gained more doctors in one mechanism than another. For example, 376 hospitals (27.7 % of the total of 1,357 hospital programs) are matched with more doctors in FDA than in JRMP. Unlike the corresponding table for the doctors, (2), some hospitals receives additional doctors while others lose doctors in any transition between the 3 mechanisms. But overall, the number of hospitals that receive more doctors is larger than the number of hospitals that lose doctors in transitions from JRMP to FDA or DA, which is not surprising given that more doctors are matched in FDA and DA than in JRMP (although this does not necessarily imply that more hospitals are matched under FDA or DA than under JRMP).

B.2.2. *Regions.* It is not surprising that improvement happens for some prefectures from JRMP or FDA to DA. Also, more than 3/4 of the prefectures are assigned more doctors

⁷⁷In this figure, we plot only the doctors who are matched with some hospital. This is because it is the information that JRMP provides in their reports (and that appears to be reasonable statistics which most people care about).

under FDA than under DA. This indicates that, due to the introduction of regional caps, many prefectures became better off in terms of the number of doctors. The more relevant question is the comparison of the improvements from JRMP to FDA and from FDA to JRMP. The exact comparison of the numbers may not make too much sense, but the numbers indicate that the introduction of FDA does not create a situation where “most regions get worse off.” This is one finding that the theory did not tell us.

| From\To | DA | FDA | JRMP |
|---------|-------------|-------------|-------------|
| DA | 0 | 37 (78.7%) | 30 (63.8%) |
| FDA | 6 (12.8 %) | 0 | 25 (53.2 %) |
| JRMP | 16 (34.0 %) | 20 (42.6 %) | 0 |

TABLE 3. The number of regions that are assigned strictly more doctors from JRMP to FDA, from FDA to DA, and from JRMP to DA.

An issue of interest that is suppressed in the above table is the magnitude of improvements and decline in different regions.

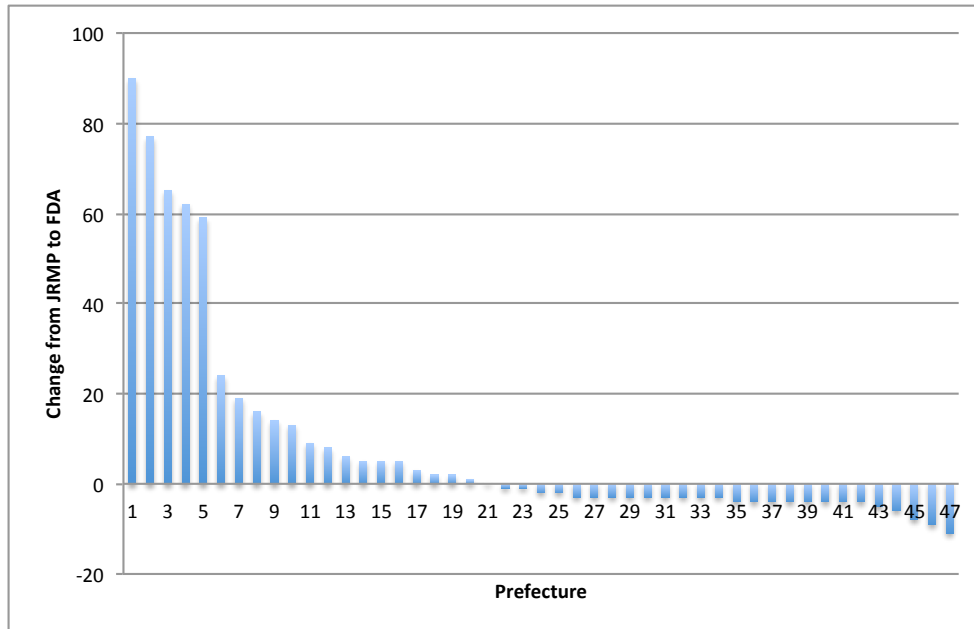


FIGURE 4. The magnitude of improvements and decline in different regions.

To better see such magnitude, in the above figure we plot the change in the number of doctors in each region, in a descending order. The graph shows that the magnitude of improvement from JRMP to FDA is large while that of decline is small. To take some

numbers, the maximum improvement is 90 while the maximum decline is 11. The area above the positive region is larger than the one in the negative region, which is consistent with our overall finding that FDA assigns about 400 more doctors in hospitals than JRMP does.

This graph is, however, silent about the distributional consequences across regions. This issue appears to be one of the main concerns in Japan. To study this issue, the figure below plots which regions become better off and which become worse off in the transition from JRMP to FDA.




FIGURE 5. Distributional consequence of the change from JRMP to FDA.

The horizontal axis measures the number of doctors (not limited to residents) per 100,000 population from 2006, which is the latest data before 2007 that is available to us, and the vertical axis describes the proportional increase of doctor assignment in the region caused by the change of the mechanism from JRMP to FDA. The motivation for this figure is to use the number of doctors per capita as a proxy for how popular each prefecture is among doctors, and study whether there is any redistribution of doctors between popular and unpopular areas.

This figure suggests that there is virtually no adverse distributional consequence against rural areas. The linear regression suggests only a slight amount of positive relation between popularity of the prefectures and improvement/decline of doctor assignment, and indeed the R^2 value is as low as 0.00327, suggesting that there is virtually no statistical correlation between the improvement/decline of doctor assignment and how popular the area is.

APPENDIX C. ADDITIONAL EXAMPLES

In this section we present three additional examples that present various comparative statics.

The first two examples strengthen the examples on comparative statics regarding regional preferences in the main text by showing that they hold under stronger assumptions on hospital preferences.

Example 14 (Ordering a hospital earlier may make it worse off even under homogenous hospital preferences). Let there be hospitals h_1 and h_2 in region r_1 , and h_3 and h_4 in region r_2 . Suppose that $(q_{h_1}, q_{h_2}, q_{h_3}, q_{h_4}) = (2, 2, 2, 2)$ and $(\bar{q}_{h_1}, \bar{q}_{h_2}, \bar{q}_{h_3}, \bar{q}_{h_4}) = (1, 0, 0, 0)$. The regional cap of r_1 is 2 and that for r_2 is 1. Preferences are

$$\begin{aligned} &\succ_{h_i}: d_1, d_2, d_3, d_4 \quad \text{for all } i = 1, \dots, 4, \\ &\succ_{d_1}: h_4, h_1, \quad \succ_{d_2}: h_1, \quad \succ_{d_3}: h_2, \quad \succ_{d_4}: h_1, h_3. \end{aligned}$$

We assume that h_3 is ordered earlier than h_4 .

- (1) Assume that h_1 is ordered earlier than h_2 . In that case, in the flexible deferred acceptance mechanism, d_1 applies to h_4 , d_2 and d_4 apply to h_1 , and d_3 applies to h_2 . d_1 , d_2 , and d_4 are accepted while d_3 is rejected. The matching finalizes with:

$$\mu = \begin{pmatrix} h_1 & h_2 & h_3 & h_4 & \emptyset \\ d_2, d_4 & \emptyset & \emptyset & d_1 & d_3 \end{pmatrix}.$$

- (2) Assume that h_1 is ordered after h_2 . In that case, in the flexible deferred acceptance mechanism, d_1 applies to h_4 , d_2 and d_4 apply to h_1 , and d_3 applies to h_2 . d_1 , d_2 , and d_3 are accepted while d_4 is rejected. d_4 applies to h_3 next, and d_1 is rejected. d_1 then applies to h_1 , which now rejects d_2 . The matching finalizes with:

$$\mu' = \begin{pmatrix} h_1 & h_2 & h_3 & h_4 & \emptyset \\ d_1 & d_3 & d_4 & \emptyset & d_2 \end{pmatrix}.$$

First, notice that hospital h_2 is better off in case (2). Thus being ordered earlier helps h_2 in this example. However, if h_1 prefers $\{d_1\}$ to $\{d_2, d_4\}$ (which is consistent with the

assumption that hospital preferences are responsive with capacities), then h_1 is also made better off in case (2). Therefore, the effect of a picking order on hospitals' welfare is not monotone. \square

Example 15 (Target monotonicity may fail even under homogenous hospital preferences). Consider a market that is identical to the one in Example 8, except that the target of h_1 is now decreased to 0, with the order such that h_1 chooses before h_2 . Then h_1 is matched to $\{d_1\}$ under the flexible deferred acceptance mechanism. Therefore, if h_1 prefers $\{d_1\}$ to $\{d_2, d_4\}$, then h_1 is made better off when its target capacity is smaller. \square

In these examples, it is hospitals that have homogeneous preferences. However, these examples can be modified so that doctors have homogeneous preferences. To do so, modify preferences to

$$\begin{aligned} \succ_{h_1}: d_1, d_2, d_4, \quad \succ_{h_2}: d_3, \quad \succ_{h_3}: d_4, \quad \succ_{h_4}: d_1, \\ \succ_{d_i}: h_4, h_1, h_3, h_2 \quad \text{for all } i = 1, \dots, 4. \end{aligned}$$

That is, hospital h finds doctor d acceptable if and only if d finds h acceptable in the previous examples, while all doctors find all hospitals acceptable and the ranking between two hospitals are consistent with the rankings between two acceptable hospitals in the previous examples. By construction, the matchings produced by the flexible deferred acceptance algorithm in this market are identical to those in the previous examples.

The next example studies comparative statics. Consider splitting a region into a number of smaller regions that partition the original region, and dividing the original regional cap among the new smaller regions. One might suspect that doing so makes doctors weakly worse off because the new set of constraints based on smaller regions may appear more stringent. The following example shows that this conjecture is incorrect. In fact, splitting regions can make some doctors and hospitals strictly better off, while making other doctors and hospitals strictly worse off.

Example 16 (Splitting regions has ambiguous welfare effects). Let there be three hospitals, h_i for $i = 1, 2, 3$ in the grand region r with regional cap of 1. The capacity of each hospital is 1. There are three doctors in the market, d_i for $i = 1, 2, 3$. Suppose that the regional preferences are such that $(1, 0, 0) \succ_r (0, 1, 0) \succ_r (0, 0, 1)$.

We examine the effect of splitting region r into two smaller regions, $r' = \{h_1, h_3\}$ and $r'' = \{h_2\}$. The splitting needs some rule of allocating the regional cap to the smaller regions, which in this example corresponds to allocating the cap 1 of r either to r' or to r''

(while allocating the regional cap of zero to the other region).⁷⁸ In what follows we show that in either case, there exists a preference profile such that the welfare effect of splitting is ambiguous (i.e., under such a preference profile it is not the case that every agent of one side of the market becomes weakly better/worse off) under the flexible deferred acceptance mechanism.

Suppose first that the cap 1 of r is allocated to r' . Then, suppose

$$\succ_{d_i}: h_i, \quad \succ_{h_i}: d_i$$

for $i = 2, 3$, and d_1 and h_1 regard no one as acceptable. The flexible deferred acceptance mechanism produces a matching μ such that $\mu_{d_2} = h_2$ before splitting, while it produces a matching μ' such that $\mu'_{d_3} = h_3$ after splitting (no other doctors are matched in either matching). Thus, splitting the region r makes d_2 and h_2 strictly worse off, while making d_3 and h_3 strictly better off.

Suppose second that the cap 1 of r is allocated to r'' . Then, suppose

$$\succ_{d_i}: h_i, \quad \succ_{h_i}: d_i$$

for $i = 1, 2$, and d_3 and h_3 regard no one as acceptable. The flexible deferred acceptance mechanism produces a matching μ such that $\mu_{d_1} = h_1$ before splitting, while it produces a matching μ' such that $\mu'_{d_2} = h_2$ after splitting (no other doctors are matched in either matching). Thus, splitting the region r makes d_1 and h_1 strictly worse off, while making d_2 and h_2 strictly better off. \square

Note that an analogous example can be easily constructed to show that the effect of splitting on the welfare of the hospitals outside the split region is also ambiguous. Finally, also note that the conclusion holds regardless of how we define regional preferences after splitting the grand region r .

⁷⁸It is only for simplicity that we use an example in which the regional cap of the grand region is one, and thus one of the smaller regions has a regional cap of zero. Our conclusion does not depend on this (perhaps unrealistic) assumption: The same point can be made in examples in which a region with regional cap larger than one is split.