Welfare and Strategic Externalities in Matching Markets with Interviews

Erling Skancke

This version: November 14, 2021, First version: November 10, 2021

For the latest version, click here

Abstract

Recent debate in the medical literature has brought attention to issues with the pre-match interview process for residency and fellowship positions at hospitals. However, little is known about the economics of this decentralized process. In this paper, I build a game-theoretic model in which hospitals simultaneously decide on which doctors to interview, in order to learn their preferences over doctors. I show that increased interview activity by any hospital imposes an unambiguous negative welfare externality on all other hospitals. In equilibrium, both hospitals and doctors may be better off by a coordinated reduction in interview activity. The strategic externality is more subtle, and conditions are derived under which the game exhibits either strategic complementarities or substitutes. Moreover, an increase in market size may exacerbate the interview externalities, preventing agents from reaping the thick market benefits that would arise in the absence of the costly interviews. This effect increases participants' incentives to match outside of the centralized clearinghouse as markets become thicker, jeopardizing the long-term viability of the clearinghouse. The model also provides new insights into several market design interventions that have recently been proposed.

Keywords: Matching with Interviews, Market Design, Strategic Complementarities, Negative Welfare Externalities, Inefficient equilibrium, Residency match, NRMP, Market Thickness

[*]For useful commments and advice I thank Alvin E. Roth, Doug Bernheim, Mira Frick, Ryota Iijima, Nadia Kotova, Negar Matoorian Pour, Paul Milgrom, Thayer Morrill, Muriel Niederle, Akhil Vohra, Daniel Walton, Anthony Lee Zhang, David Zuckerman and seminar participants at Stanford University. I gratefully acknowledge financial support from the Leonard W. Ely and Shirley R. Ely Graduate Student Fellowship.

[†]Stanford University, Department of Economics: erling.skancke@gmail.com

1 Introduction

Many matching markets involve various forms of costly information acquisition through which market participants learn about their preferences. Examples include students attending visit days at schools, colleges or universities before deciding where to embark on their education; people dating before finding a partner they want to settle down with; and home buyers visiting a number of different homes before deciding to place an offer. A similar costly process is also evident in many labor markets, in which both firms and workers spend significant time and resources conducting either in-person or virtual job interviews, in addition to the sending and processing of job applications. Despite the prevalence of this costly search activity, little is known about the economics of this process: What incentives do firms and workers face when making their interview decisions? How do participants' interview decisions affect the welfare and interview decisions of others? How does agents' interview activity respond to changing market conditions, such as a change in the interview costs/technology, or an increase in market size? Finally, how does the costly search activity affect agents' incentives to participate in a centralized clearinghouse?

This paper considers the above questions in the context of the entry- and intermediate-level markets for doctors in the US, through the analysis of a game-theoretic model of matching with interviews. These are interesting markets in their own right, in part due to their size: In 2021, participation in the matches organized by the National Residency Matching Program (NRMP) reached record numbers, with a total of more than 48,700 applicants and 38,100 residency positions participating in the Main Residency Match, and more than 12,900 applicants and 11,700 fellowship positions participating in additional matches across 68 subspecialties. These markets are also some of the best known examples of labor markets that make use of a centralized clearinghouse for matching workers to firms. Indeed, the NRMP employs a variant of the celebrated deferred-acceptance algorithm developed by Gale and Shapley (1962) and Roth and Peranson (1999), which has been shown to have many desirable properties (Roth and Sotomayor (1992), Roth (1991)). The use of a stable matching mechanism makes these markets particularly convenient to study, as it simplifies the link between market participants' preferences/rank-order lists submitted to the clearinghouse and the match outcome. On the other hand, contrary to the centralized final matching of doctors to positions at hospitals, participants' pre-match job search and interview activity is a largely decentralized process. The interview process has recently been the subject of debate in the medical literature: Some raise concerns about the high level of costs (Watson et al.

(2017), Agarwal et al. (2017)), while others point to excessive interview activity as both sides of the market seek to ensure a successful match (Strand and Sonn (2018)), a situation which may result in interview hoarding (Bernstein (2021)). Despite being some of the most widely studied markets in the matching literature, the pre-match search and interview process has largely been ignored.

In this paper, I build a game-theoretic model in which hospitals simultaneously invite doctors to conduct costly interviews, in order for hospitals to determine doctor compatibility with their program. I show that an increase in hospital interview activity imposes unambiguous negative welfare externalities on all other hospitals. When hosptials' interview costs are sufficiently low, the strategic externalities cause hospitals to increase their number of interviews whenever their competitors do, leading to a game of strategic complementarities and equilibria characterized by inefficiently many interviews. When doctors' interview costs are sufficiently high, they too may be worse off as hospitals decide to interview more candidates. As a result, there may be equilibria in which both sides of the market can be made better off by a coordinated reduction in interview activity. I also explore the effect of changing market conditions that parallel those seen in the entry-level market for doctors in the US; such as the apparent reduction in interview costs during the 2020-2021 match season when most candidates and programs moved to virtual interviews. I show that the welfare effects of this change are generally ambiguous. I also investigate the effect of increasing the number of both available positions and participating candidates, and find that hospital equilibrium welfare may fall as markets thicken. While not unique, this result is in stark contrast to the many papers that emphasize the positive effects of thick markets. Moreover, as market participants' equilibrium welfare decreases, participants' incentives to match *outside* the centralized clearinghouse increase. This effect could jeopardize the long-term viability of clearing-houses such as the NRMP in the absence of an appropriate market design intervention. Finally, the model allows us to assess several market design interventions that have been proposed in the medical literature, but whose theoretical properties remain unknown.

In the model, hospitals have partial pre-interview information about their match values which rank doctors uniformly and independently, but must conduct interviews to learn their full preferences. Doctors are assumed to accept all interview offers they receive, an assumption meant to represent the empirical fact that most of the interview selection happens on the hospital side. Indeed, most doctors tend to accept the vast majority of the interviews they are offered (National Residency Matching Program and Committee (2017)). When doctors' interview costs are sufficiently low, I show that it is individually rational for them to accept all interviews. After the interviews are conducted, and participants learn their preferences, hospitals and doctors are matched according to a deferred-acceptance algorithm in which doctors and hospitals only rank those that they previously interviewed with. My focus is on *Bayesian interview equilibria*: Hospitals optimally choose to interview doctors until the marginal cost of an additional interview exceeds the marginal benefit, given other hospitals' interview strategies. My focus is on equilibria in *anonymous strategies*, strategies which do not depend on doctors' identities, such that hospitals cannot condition their interview decisions on the identity of the doctors interviewed by their competitors.

Having set up the model, I show that a hospital is always worse off as its competitors increase their interview activity. I also analyze the effect of the competitors' increased interview activity on the hospital's benefit of an additional interview. I show that the effect can be decomposed into two terms: The *competition effect* reduces the value of any interview that the hospital conducts. As hospital h's competitors interview more, it increases the probability that any of the candidates interviewed by h is also interviewed by its competitors, and potentially match with them. This reduces h's incentives to interview any doctor. The competition effect is partially offset by the *probability reallocation effect*: During the deferred-accpetance algorithm, if hospital h must go far down on its rank-order list to find a match, it is likely because some of the higher ranked doctors were matched with other hospitals, and thus more of the hospital's competitors already filled their positions. As hospital h is less likely to compete for its lower ranked doctors, this increases the marginal benefit of having a longer rank-order list, and thus interviewing more doctors. Thus, when the hospital's interview costs are sufficiently low and its competitors increase their interview activity, the hospital's best-response is to increase its own number of interviews.

From the doctors' perspective the situation is more complicated. On the one hand, I show that doctors' expected match values are monotonically increasing in the number of interviews hospitals conduct. On the other hand, I also assume doctors incur interview costs (even though in the game they do not have the option to reject any interviews). The more interviews hospitals conduct, the more likely a doctor is to receive any interview, and the more likely they are to get matched during the deferred-acceptance algorithm. However, doctors' expected number of interviews always increases faster than their expected match rate. As a result, when doctors' interview costs are sufficiently high, their marginal increase in expected interview costs exceed their marginal increase in expected match quality, as hospitals increase their interview activity. Interestingly, if doctors' interview costs are not too large, it may be individually rational for them to accept all interviews even though they would be better off with hospitals reducing their interview activity.

Combining the above, my analysis shows that when interview costs are sufficiently low, and the game exhibits strategic complementarities, the set of pure equilibria in the interview game forms a lattice ordered according to how many interviews each hospital conducts. As the equilibrium number of interviews increases, all hospitals are worse off. Moreover, if doctors' interview costs are sufficiently high, doctors' welfare is also decreasing in the equilibrium number of interviews. Thus, welfare on both sides of the market is decreasing in the equilibrium number of interviews. This may be surprising, as classical results in the theory of stable matching state that if one side of the market is worse off, then all agents on the other side of the market are better off. Once market participants' search and interviewing costs are taken into account, this classical result no longer holds. As a result, participants on both sides of the market could benefit from reduced interview activity, a fact that becomes relevant for my later discussion of market design interventions.

After characterizing the model's set of equilibria, I next turn to comparative statics: First, as the market for residency and fellowship positions largely moved to virtual interviews during the 2020-2021 job market season, a natural question is whether this relatively low-cost interview method improved participants' welfare compared to more traditional and high-cost in-person interviews. Intuitively, lower per-interview costs make everyone better off through the direct positive effect on each interview conducted. However, it also provides incentives to hospitals to increase their interview activity, exposing other hospitals to increased negative welfare externalities, as described above. The overall effect is ambiguous, and equilibrium welfare may not be monotonic in interview costs, because of the counter-acting forces. Second, the past decades have seen increasing trends in the participation of both hospitals and doctors in the NRMP. Intuition suggests that hospitals benefit from an increase in the number of doctors, but are worse off as the number of hospitals increases, due to increased competition. What happens to welfare and the equilibrium number of interviews when both sides of the market increase proportionately is ambiguous. If interview costs are low, hospitals will be better off as they expand their search for candidates. As markets become sufficiently thick, the benefits of improved match utilities fade against the negative welfare externalities and increased interview costs, making hospitals worse off. Under the assumption of no pre-interview information, I prove the existence of a symmetric equilibrium in anonymous strategies that maximizes hospital welfare, and I construct an algorithm to find this equilibrium. Through numerical simulations, I compare the welfare-maximizing equilibrium for different levels of market thickness and interview costs, which confirms the above intuition. That welfare decreases as markets become thicker is in stark contrast to the results in many existing papers which highlight the benefits of thick markets. However, these papers typically ignore the search costs incurred by agents. In the presence of search costs, market participants may not be able to reap the full benefits of market thickness. I illustrate this point by comparing the equilibrium welfare with interview costs to the welfare that would result in the case with zero interview costs. As markets become sufficiently thick, the two measures of welfare indeed diverge.

These results also cast doubt on the long-term viability of the centralized clearinghouse. Both hospitals and doctors may have access to matching opportunities outside of those offered by the match (such as the hiring of medical students who previously interned at hospitals). Such opportunities will not pursued as long as agents expect to obtain an even better outcome through the centralized match. If the expected utility from participating in the match decreases as markets thicken, then the incentives to pursue opportunities outside of the match increase. Eventually, this may lower participation in the clearinghouse, causing agents to forego the benefits of the centralized match. Taken to its extreme, the decrease in equilibrium welfare may trigger markets to unravel.

I finally turn to analyze several market design interventions that have recently been proposed in the medical literature, such as imposing a limit on the number of applications doctors can send (Burbano et al. (2019)). Major advantages of this intervention is that it could be easy to enforce and inexpensive to implement, while it may also provide large benefits if hospitals' costs of processing applications are high. That said, this intervention does not address the fundamental issue of welfare- and strategic externalities in the interview process. Other proposed market interventions that impose restrictions on the number of interviews agents conduct (e.g. Wapnir et al. (2021)) address these issues more directly. By way of example, I illustrate that a policy that limits the number of applications doctors are allowed to send indeed improves upon the unconstrained equilibrium, but is less efficient than an appropriate two-sided restriction on interview activity.

The rest of the paper is organized as follows: Section 2 discusses the related literature. Section 3 introduces the model and the interview game, introduces the solution concept used throughout, and establishes key intermediary results on hospitals' optimal strategies and the existence of equilibrium. In Section 4, I consider the externalities that result from hospitals' interview decisions. I analyze the welfare consequences from the perspective of both hospitals and doctors, and derive conditions under which the interview game exhibits either strategic complementarities or substitutes. Section 5 explores comparative statics: I first show that equilibrium welfare is generally non-monotonic in interview costs, before investigating the effect of market thickness. In Section 6, I discuss several market design interventions that have been recently been proposed. Section 7 concludes.

2 Related Literature

While there is a large literature on post-interview matching, relatively few papers have studied the pre-match application and interview processes. In the computer science literature, some papers have looked at "matching with partial information", although the research agenda is typically far removed from that of my paper: For instance, Rastegari et al. (2013) derive the minimum number of interviews needed to ensure a stable matching with respect to the underlying true preferences, noting that the problem is NP-hard, while Rastegari et al. (2014) show that one can decide in polynomial time whether such a stable matching exists or not. While this analysis may prove crucial to the design of appropriate market design interventions, it does not provide too many insights into the nature of any market failure that may result in inefficiencies. The primary focus in my paper is on the equilibrium interview assignments that result from the strategic interaction between hospitals and doctors, and the nature and direction of the externalities to which hospitals and doctors are exposed.

In the economics literature, early work focused on models with correlated preferences, i.e. when doctors agree on the ranking of hospitals, and hospitals agree on who the top candidates are. These papers drew attention to candidates who "fall through the cracks", leading to non-assortative matches. This may result from hospitals deciding to spread their interviews between the very top candidates and lower ranked "safety candidates", creating a vacuum for medium quality candidates (Lien (2009), Kadam (2015)), or by firms simply making mistakes in noisy environments (Das and Li (2014)). Contrary to my focus on excessively high interview activity, a main concern in these papers has been that some agents fail to attract *enough* interviews. I abstract from the difficulties of correlated preferences, and assume that all agents have independent preferences. While a simplification, the assumption seems more relevant after the United States Medical Licensing Examination moved from reporting a three-digit numeric score to reporting only pass/fail (USMLE (2020)), providing less basis for hospitals to agree on a common ranking of candidates.

A recent paper by Manjunath and Morrill (2021) introduces welfare results that, on the surface, echo some of those presented in this project. In an intriguing analysis that is solely based on the properties of the deferred-acceptance algorithm, they show that all agents on the proposing side are made worse off when their side decides to increase their interview capacities. While compelling, their analysis ignores agents' strategic considerations in their interview decisions, and the interview allocations they consider generally do not form equilibria of any underlying game. That increasing interview capacities is not unambiguously positive has already been shown in Kadam (2015), who shows that high quality candidates may benefit from interviewing more, while lower quality candidates are made worse off. The overall match rate may also fall from increased interview capacities. However, Kadam (2015) only considers changes in the interview capacities of candidates, who are on the receiving side in the associated interview game. Specifically, he does not consider changes in firms' interview capacities. My analysis is fundamentally different, in that the focal point is on equilibrium interview decisions when firms must trade off the marginal benefit of an additional interview against its marginal cost.

The model in my paper is most closely related to the one analyzed in Lee and Schwarz (2017). The structure of their interview game is almost identical, but they consider a case where hospitals have no pre-interview information about their preferences, hence their game is not one with private information. My model generalizes their setup. More importantly, the purpose of the analysis in Lee and Schwarz (2017) is very different: Their focus is on the properties of particular interview assignments, and they restrict attention to deriving conditions on the market environment (interview costs) under which these assignments form equilibria. In contrast, I fix the market conditions and then characterize the properties of hospitals' best-response correspondences, the qualitiative and quantitative nature of the welfare externalities emerging from hospitals' interview decisions, and in some cases, the properties of the set of equilibria under a particular equilibrium selection criterion. My approach allows me both to compare the different equilibra for a given set of market conditions, and more importantly, to compare the set of equilibria as market conditions change.

The concept of interviewing is also closely related to that of applications, which has received a lot of attention in the college admission literature (see e.g. Chade et al. (2014)). As with interviews, matching games with applications typically exhibit inefficient equilibria. While considering different contexts, both Arnosti et al. (2021) and Beyhaghi and Tardos (2021) show that lower caps on applications may lead to welfare improvements. In the medical literature, Burbano et al. (2019) cite Arnosti et al. (2021) to claim that the introduction of a cap on applications is "the most reasonable approach to address the issue [of application and interview costs]". On the contrary, my project helps to highlight the differences between applications and interviews. While placing a cap on doctors' applications may have positive effects on the market performance, it far from solves the issue of negative welfare externalities from excessive interview activity.

3 Model

3.1 Notation, main assumptions, timing of the game and agents' strategies

Consider a two-sided matching market consisting of a set of hospitals, $\mathcal{H} = \{h_1, ..., h_{|H|}\}$, and a set of doctors, $\mathcal{D} = \{d_1, ..., d_{|\mathcal{D}|}\}$. Each hospital has one position they seek to fill, and each doctor is looking for at most one position. Agents will be matched using the hospital-proposing deferred-acceptance algorithm (HPDA), with outcome denoted by μ.

Hospitals are endowed with pre-interview information about their match values with doctors, contained in a real-valued vector $\theta_h = (\theta_{hd_1}, ..., \theta_{hd_{|D|}})$, referred to as the hospital's type. θ_h is drawn from a distribution G with support Θ. Conditional on θ_{hd}, hospital h's value of matching with doctor d, denoted v_{hd}, is drawn according to a distribution $F_{\theta_{hd}}$. Let $\beta := \mathbb{P}\big(v_{hd} > 0\big) > 0$ denote the probability that h finds d acceptable, assumed the same for all hospital-doctor pairs.[1] Denote by v^+_{hd} the positive part of v_{hd}, and F^+_θ the positive part of F_θ. From Section 4 onward, we will also make use of the following notation: For any $S \subset \mathcal{D}$, define $v^+_{h(k,S)}$ to be the k-th highest values of $(v^+_{hd})_{d\in S}$, conditional on all the doctors in S being found acceptable. We refer to $\mathbb{E}\big[v^+_{h(k,S)}\big]$ as the k-th *rank-order statistic* of the set S. For $k > |S|$, define $v^+_{h(k,S)} \equiv 0$. We assume the following:

Assumptions

(A1) $\mathbb{E}[v^+_{hd} - y | v^+_{hd} > y]$ is decreasing in y for every θ_{hd}

(A2) there exists bounds $0 < \underline{v}_{\mathcal{H}} \leq \overline{v}_{\mathcal{H}} < \infty$ such that $\underline{v}_{\mathcal{H}} \leq \mathbb{E}[v^+_{hd}] \leq \overline{v}_{\mathcal{H}}$ for every θ_{hd}.

(A3) if $\theta_{hd} \geq \theta_{hd}$ then v_{hd} first-order stochastically dominates $v_{hd'}$

(A4) v_{hd} independent of both $v_{hd'}$ and $v_{h'd''}$ for $d' \neq d$, $h' \neq h$, and $d'' \in \mathcal{D}$, conditional on θ_h.

(A5) hospitals' pre-interview information θ_h are draws from a distribution G such that:

(a) for any hospital h, conditional on θ_h, all other hospitals look the same: $\mathbb{P}_{G,h'}\big(\theta|\theta_h\big) = \mathbb{P}_{G,h''}\big(\theta|\theta_h\big)$ for all θ, all $h', h'' \neq h$, and

[1] It is possible to relax the assumption of a common probability of acceptability. Since, however, this complicates the algebra without adding much economic intuition, we have made the simplifying assumption of a common β for all doctors. In the appendix, we briefly discuss how the model can be extended to account for differences in the probability of acceptability.

(b) for every pair (h, h'), conditional on h's pre-interview information θ_h, for every set $D \subset \mathcal{D}$ and h' type $\theta_{h'}$ such that $\theta_{h'd} > 0$ for all $d \in D$, for every permutation π^D over D, we have $\mathbb{P}\big(\pi^D \theta_{h'} \,|\, \theta_h\big) = \mathbb{P}\big(\theta_{h'} \,|\, \theta_h\big)$.

(A6) doctors find all hospitals acceptable, their match utilities $v_d = (v_{dh})_{h \in \mathcal{H}}$ are private information and are independent draws from a distribution $F_{\mathcal{D}}$ with $supp(F_{\mathcal{D}}) = [\underline{v_{\mathcal{D}}}, \overline{v_{\mathcal{D}}}]^{|\mathcal{H}|} > 0$.

Assumptions[for Sections 4.2 and 4.3]

(A5*) Hospitals' pre-interview information θ_h are *independent* draws from a distribution G such that for every permutation $\pi^{\mathcal{D}}$ over $\mathcal{D}$, we have $\mathbb{P}(\pi^{\mathcal{D}}\theta) = \mathbb{P}(\theta)$.

(A7) for any doctor $d' \in \mathcal{D}$ and any set $S \subset \mathcal{D} \setminus \{d'\}$, if $\theta_{hd} \geq \theta_{hd'}\ \forall\, d \in S$, then $\mathbb{E}[v^+_{h(k,S\cup\{d'\})} - v^+_{h(k,S)}] \leq \mathbb{E}[v^+_{h(k+1,S\cup\{d'\})} - v^+_{h(k+1,S)}]$ for any $k \leq |S|$.

(A8) there exists $\varepsilon \in (0,1]$ such that for any set $S \subset D$ and any pair of doctors $d', d'' \notin S$, if $\theta_{hd'} \geq \theta_{hd''}$ and $\theta_{hd} \geq \theta_{hd''}$ for all $d \in S$, then
$\mathbb{E}\big[v^+_{h(|S|+2,S\cup\{d',d''\})} - \big(v^+_{h(|S|+1,S\cup\{d',d''\})} - v^+_{h(|S|+1,S\cup\{d'\})}\big)\big] \geq \varepsilon\, \mathbb{E}\big[v^+_{h(|S|+1,S\cup\{d''\})}\big]$

While we have alluded to many of the elements of the interview game, we here formalize the timing of the game and explicitly describe agents' strategy spaces:

Timing of the Interview Game:

1. Hospitals learn their pre-interview information θ_h.

2. Hospitals simultaneously choose which doctors to interview. Doctors accept all interviews.

3. Hospitals and doctors learn their match utilities, v_h and v_d, only for those agents with whom they interviewed.[2]

4. Hospitals and doctors all submit rank-order lists. Doctors rank hospitals according to their ex ante match utilities. Hospitals rank according to their post-interview preferences all doctors they both (i) interviewed and (ii) found acceptable, and do not rank the remaining doctors.

[2] As noted by Lee and Schwarz (2017), the exact timing of when doctors learn their match utilities is not of significance for any of the results: All results would go through if doctors knew their match utilities from the very beginning of the game.

The tuple $(\mathcal{H}, \mathcal{D}, c_{\mathcal{H}}, G, (F)_{\theta\in\Theta})$ describes an *interview game* of incomplete information between the hospitals in $\mathcal{H}$ (doctors are passive players). Hospitals' strategies consist of choosing which doctors to interview: For each realization of ex ante preferences, θ_h, hospital h's strategy specifies a subset (or a mixing over subsets) of $\mathcal{D}$:

Definition 3.1 *A* ***mixed Bayesian interview equilibrium*** *is a tuple* $\big(\sigma_h(\theta_h)\big)_{\theta_h\in\Theta,\, h\in\mathcal{H}}$ *such that for each* $h \in \mathcal{H}$ *and each* $\theta_h \in \Theta$

$$\sigma_h(\theta_h) \in \Delta \arg\max_{S\subseteq D} \left\{ \sum_{d\in S} \mathbb{P}\Big(\mu(h_j) = d \big| \sigma_{-h}\Big) \mathbb{E}\big[v_{hd} \,\big|\, \mu(h) = d,\, \sigma_{-h}\big] - c_{\mathcal{H}}|S| \right\} \tag{1}$$

where $V_{\theta_h}(S, \sigma_{-h}) := \sum_{d\in S} \mathbb{P}\Big(\mu(h_j) = d\big|\sigma_{-h}\Big) E\big[v_{hd} \,\big|\, \mu(h) = d,\, \sigma_{-h}\big]$ *denotes h's expected match value from interviewing the set* S*, conditional on* θ_h *and conditional on the other hospitals' strategies* σ_{-h}*.*

For the rest of the paper, we will restrict attention to equilibria in anonymous strategies:

Definition 3.2 *A strategy* σ *is called* ***anonymous*** *if for any* $\theta \in supp(G)$ *and for any permutation* π*,* $\sigma(\pi\theta) = \pi\sigma(\theta)$*.*

Note that in the case of no pre-interview preferences ($\theta_{hd} \equiv \theta$ for all $d \in \mathcal{D}$), an anonymous strategy implies that h interviews each doctor with the same probability.[3]

Before proceeding with our analysis and results, we will discuss some of the main assumptions we have made: Assumption (A1) ensures that the benefit of an additional interview decreases the more interviews a hospital is conducting, allowing us to employ standard tools from convex optimization when studying hospitals' optimization problem. (A2) ensures that hospitals' optimization problems are well-defined, and that a hospital would be willing to interview any doctor, as long as the interview cost is sufficiently small. As we will show, (A3) implies that hospitals find it optimal to interview the doctors in descending order of the pre-interview preferences, while (A4)-(A5) imposes no correlation (neither conditional nor unconditional) in ' preferences both prior and post

[3]It is worth contrasting the above definitions to those used in Lee and Schwarz (2017). They refer to "mixed-strategy equilibrium" as an equilibrium in which hospitals interview x doctors at random. In our terminology, their "mixed-strategy equilibrium" corresponds to an equilibrium in anonymous strategies in which all hospitals conduct x interviews each (and have no pre-interview information). As will become apparent later, I allow for an additional kind of mixing, e.g. one where hospitals mix between conducting x and $x + 1$ interviews. As I show in Section 3.2, this allows us to state general existence result without imposing additional conditions on interview costs, as do Lee and Schwarz (2017).

interviews, on both sides of the market. (A6) states that doctors find all hospitals acceptable. This is without loss of generality: If doctors found hospitals unacceptable with positive probability (and this probability was independent and identical across both doctors and hospitals), then doctors' probability of unacceptability could be subsumed in the hospital probability of unacceptability, β. (A6) also states that doctors' preferences over hospitals are independent draws from a bounded distribution. The independence assumption plays an important role for most of the results of the paper, while the boundedness is only used to evaluate doctor welfare in Sections 4.1 and 4.3.

The model and its assumptions are either fully or partially consistent with much of the earlier work on matching with interviews. First, (A1)-(A4) and (A6) satisfy the no learning setup studied in Manjunath and Morrill (2021) in which $|supp(F_\theta)| = 1$ for all θ. From an individual hospital's perspective, Kadam (2015) imposes all of assumptions (A1)-(A4), although he considers a case with correlated preferences, violating (A5). Note that our assumption (A6) encompasses the case in which doctors fully agree on the ranking of hospitals when $|supp(F_{\mathcal{D}})| = 1$. Our model is perhaps most closely related to the one in Lee and Schwarz (2017), who explicitly or implicitly assume (A1)-(A4) and (A6). However, they do not consider pre-interview preferences, meaning that (A3) and (A5) are automatically satisfied. Moreover, Lee and Schwarz (2017) consider the case in which all hospitals are treated the same by doctors, while our assumption (A6) is more general, and allows for doctors to agree on their ranking of hospitals.

It is also worth noting that different authors address the rationale for the existence of interviews in different ways: For instance, Manjunath and Morrill (2021), who consider a context with no learning, simply take the presence of the interview step as given, assuming interviews are allocated according to a many-to-many stable mechanism with interview capacities exogenously given, before agents are eventually matched in a second-stage in which participants can only rank those with whom they interviewed. An immediate objection to their model is that the interview stage serves no purpose. Indeed, since all agents know their own preferences prior to the interviews, one could immediately run the second-stage final matching. Lee and Schwarz (2017) take the antithetical route, imposing the strict assumption that firms' expected match value with any worker is negative (meaning they would be unacceptable), implying that no firm would ever want to be matched with any worker without first having interviewed them, which indeed justifies the presence of the interview stage. In this paper, we will remain mute about the exact justification for the interview stage. As long as the probability a hospital finds a doctor unacceptable is positive ($\beta < 1$), the model is consistent with the justification in Lee and Schwarz (2017). However, it is also worth pointing out that it may be in both hospitals' and doctors' best interest to conduct interviews even if they believe all doctors to be acceptable ($\beta = 1$), as long as the interviews provide a sufficient amount of new information about their post-interview preferences. Finally, in many of the examples I present in subsequent sections, I will often consider cases in which the interview stage may seem superfluous, although this is done in order to present the simplest example possible to illustrate the point at hand.

Related to the justification of the presence of an interview stage, a discussion of the exact process through which interviews are allocated is also warranted. Again, authors differ widely in their assumptions. As mentioned above, Manjunath and Morrill (2021) assume that interviews are allocated according to the hospital-proposing deferred acceptance algorithm, citing it as "an approximation of the decentralized process by which hospitals invite doctors in rounds, extending invitations to further doctors when invitations are declined". An immediate issue with this assumption is that it is in neither hospitals' nor doctors' best interest to report their preferences truthfully, as can easily be verified formally, and is also perfectly illustrated by many of the results Manjunath and Morrill (2021) present. As such, the assumption that interviews are allocated according to an ordinal stable mechanism is not immediately consistent with optimal behavior by agents. Kadam (2015) also studies an equilibrium model of interview allocation, considering a two-step game in which firms first send out interview invitations followed by a round of acceptances by workers. Importantly, contrary to Manjunath and Morrill (2021), Kadam (2015) does not allow firms to send new interview requests after some of the original ones may have been declined. While he does not completely characterize agents' optimal strategies, due to the assumption that all workers agree on the ranking of firms, and firms are aware of this ranking, the game can be solved by iterated elimination of dominated strategies. In equilibrium, doctors accept the interview requests they receive in order of their preferences, up to their (exogenously given) interview capacities, while firms take into account the strategies of all higher ranked firms, and need not find it optimal to invite doctors according to their pre-interview preferences. Note that this solution concept hinges on the assumption of the perfect and commonly known agreement of the ranking of firms. If this assumption were to be relaxed, the assumption that firms only need one round of interview requests to fill their capacities would likely no longer hold. Lee and Schwarz (2017) choose an approach that is similar to the one in this paper. In particular, they assume that workers accept all interview requests, which they justify on the basis of their assumption that workers have zero interview costs, and that any rejection at the interview stage could equivalently be accomplished by failing to list the relevant hospitals at the final matching stage. They further impose assumptions on the distribution of doctors' match values to ensure that doctors will always want to truthfully report their preferences, thus effectively circumventing the issue. In contrast, while I maintain the assumption that doctors accept all interviews, I also allow for doctors to incur interview costs. This immediately raises the concern that doctors could be better off rejecting some of their interviews. On the other hand, empirical evidence suggests that doctors tend to accept the overwhelming majority of the interviews they are offered (National Residency Matching Program and Committee (2017)), in line with the view that most of the interview selection happens on the hospital side. Moreover, under certain assumptions (see Section 4 and Theorem 4.6) it is indeed in the best interest of doctors to accept all interviews they are offered, despite the fact that they may be worse off as hospitals increase their interview activity. On a more technical note, constructing a model in which hospitals and doctors reach a "stable" interview assignment, one in which agents on both sides choose the optimal set of interviews among those offered, and all mutually beneficial interview opportunities are pursued, may not be a fruitful avenue: Indeed, due to the strategic externalities, the core in matching markets with interviews may be empty, hence the existence of a stable interview assignment cannot be guaranteed. While efforts to better understand the incentives faced by agents on both sides of the market is a promising avenue for future research, a model in which doctors play an entirely passive role at the interviewing stage has the benefit of both being tractable while also shedding light on the many issues raised in, for instance, the medical literature.

The assumption of uncorrelated preferences on the hospital side of the market is the most substantive of the assumptions above. Perhaps not surprisingly, some empirical evidence suggests a certain degree of correlation in agents' preferences (see, for instance, Agarwal (2015)). The reason for focusing on the case of uncorrelated preferences is two-fold: First, while based on strong assumptions, this simple model still allows us to understand many of the important phenomena that are at play in matching markets with interviews. It is indeed instructive to study the effect of preference formation on market outcomes in the simplest and most tractable setting possible. Second, previous work on matching with interviews has shown that the characteriation of equilibria can be untractable once one allows for correlation in preferences. Moreover, previous work also makes clear that certain conclusions are very sensitive to the assumptions on agents' preference structure. To illustrate, suppose all doctors agree on the ranking of hospitals, and that this ranking is common knowledge among all market participants. Then the most preferred hospital, let's denote it by h^*, will always be matched to its post-interview top choice, regardless of other hospitals' interview decisions. Although h^* imposes an externality on other hospitals, the actions of other hospitals will neither impose a welfare externality nor a strategic externality on h^*. As a result, absent any transfers, no equilibrium will be Pareto inefficient, since h^* can never be made better off. While this conclusion is valid, it is likely not illustrative of the current state of the entry- and intermediate-level markets for doctors in the US. As will become apparent in later sections, our model based on uncorrelated preferences will be well-equipped to speak to many of the issues raised in the medical literature, while also formalizing the discussion in a tractable model. Third, in private conversations, individuals involved in interview decisions at Stanford Hospital, both for residency and fellowship positions, have revealed that once one disregards the very top and the bottom of the spectrum of doctors in the market, one is left with a large set of doctors who either may be difficult to differentiate, or for whom it may be difficult to infer the assessment other hospitals may have of said candidates. Finally, starting in 2022, the United States Medical Licensing Examination Step 1 score will transition from a three-digit numeric score to pass/fail outcomes only, providing less basis for hospitals to agree on a common ranking of doctors (USMLE (2020)).

3.2 Optimal Strategies and Equilibrium Existence

Having set up the model, the rest of this section establishes the existence of a Bayesian interview equilibrium and provides a general characterization of hospitals' optimal strategies. These result will prove useful for our analysis in subsequent sections.

Lemma 3.3 *A Bayesian interview equilibrium in anonymous strategies always exists.*

The proof is in the appendix, and is similar to standard proofs of equilibrium existence in games of incomplete information, with the additional argument that if an interview set S maximizes a hospital's utility at θ_h, then πS maximizes the hospital's utility at $\pi\theta_h$, for any permutation π. In essence, the hospital's problem looks "identical" from the perspective of every permutation of their pre-interview information. This shows that whenever all other hospitals are playing anonymous strategies, a hospital always has an optimal best response that is anonymous.

It is also worth contrasting the above existence result with those stated in Lee and Schwarz (2017): Their focus is on symmetric equilibria in which all hospitals conduct $x \in \{0, .., |\mathcal{D}|\}$ interviews each. They then show that there exist interview costs such that this indeed forms an equilibrium. Since we allow for hospitals to mix between multiple optimal sets of doctors, our existence result is more general, and an equilibrium (not necessarily symmetric) exists for all possible interview costs.

Lee and Schwarz (2017) focus on the case where each hospital conducts an exact number of interviews, and argue that this is a result of hospitals' decreasing benefits of adding doctors to their interview lists (a consequence of Assumption (A1)). This logic extends to our setting, and it turns out that for any optimal strategy (holding other hospitals' strategies fixed), there exists a k_h such that hospital h either finds it strictly optimal to conduct exactly k_h interviews, or is indifferent between conducting k_h and $k_h - 1$ interviews, as the following result shows:

Lemma 3.4 *Suppose all other hospitals are playing anonymous strategies. Given θ_h, any optimal pure strategy by hospital h consists of adding the doctors to its interview list in descending order of pre-interview signals, until the benefit of an additional interview falls short of the additional interview cost. That is, any optimal pure strategy consists of choosing $S = \{d_{(1)}, ..., d_{(|S|)}\} \subset \mathcal{D}$ such that $\theta_{hd_{(1)}} \geq ... \geq \theta_{hd_{(|S|)}} \geq \theta_{hd}$ for $d \notin S$, such that*

(i) $V_{\theta_h}(S, \sigma_{-h}) - V_{\theta_h}(S \setminus \{d_{(|S|)}\}, \sigma_{-h}) \geq c_{\mathcal{H}}$

(ii) $V_{\theta_h}(S \cup \{d\}, \sigma_{-h}) - V_{\theta_h}(S, \sigma_{-h}) \leq c_{\mathcal{H}}$ *for all* $d \notin S$

Remark: In the case where pre-interview preferences are strict, there are at most two sets, S and $S \cup \{d'\}$, that satisfy the above two conditions. This happens whenever h is indifferent between adding the last doctor d' to its interview list or not. Similarly, in the case where pre-interview preferences are not strict, the size of the optimal sets take at most two values, i.e. there exists $k \leq |\mathcal{D}|$ such that for every optimal set S we have $|S| \in \{k-1, k\}$. This is illustrated in Figure 1.

The above result is a corollary of the following two lemmas; one stating that the benefit of adding a doctor to an interview lists is smaller the bigger is the set of doctors already being interviewed, and the other stating that it's always better to add a doctor with higher pre-interview information to any interview set. The proofs of all three results are in the appendix.

Lemma 3.5 *Suppose all other hospitals are playing anonymous strategies. Hospital h's benefit of adding a doctor to a small interview list is larger than adding the doctor to a large interview list: Let $S \subset S' \subseteq \mathcal{D}$ and $d^* \in \mathcal{D}$, then for any θ_h*

$$V_{\theta_h}(S \cup \{d^*\}, \sigma_{-h}) \geq V_{\theta_h}(S' \cup \{d^*\}, \sigma_{-h})$$

Lemma 3.6 *Suppose all other hospitals are playing anonymous strategies. Let $S \subset \mathcal{D}$, and $d, d' \in \mathcal{D} \setminus S$. Suppose $\theta_{hd} \geq \theta_{hd'}$, then the benefit of adding d to the interview list S is greater than the benefit of adding d': $V(S \cup \{d\}, \sigma_{-h}) \geq V(S \cup \{d'\}, \sigma_{-h})$.*

Figure 1: **Optimal interview strategy given θ_h and other hospitals' strategies σ_{-h}**
An optimal strategy interviews doctors in descending order according to θ_h until the benefit of an additional interview no longer exceeds the interview costs $c_{\mathcal{H}}$. The marginal benefit curve is decreasing in the number of total interviews. In the case where the benefit of the last interview exactly equals $c_{\mathcal{H}}$, the hospital is indifferent between adding the last doctor to its interview list or not.

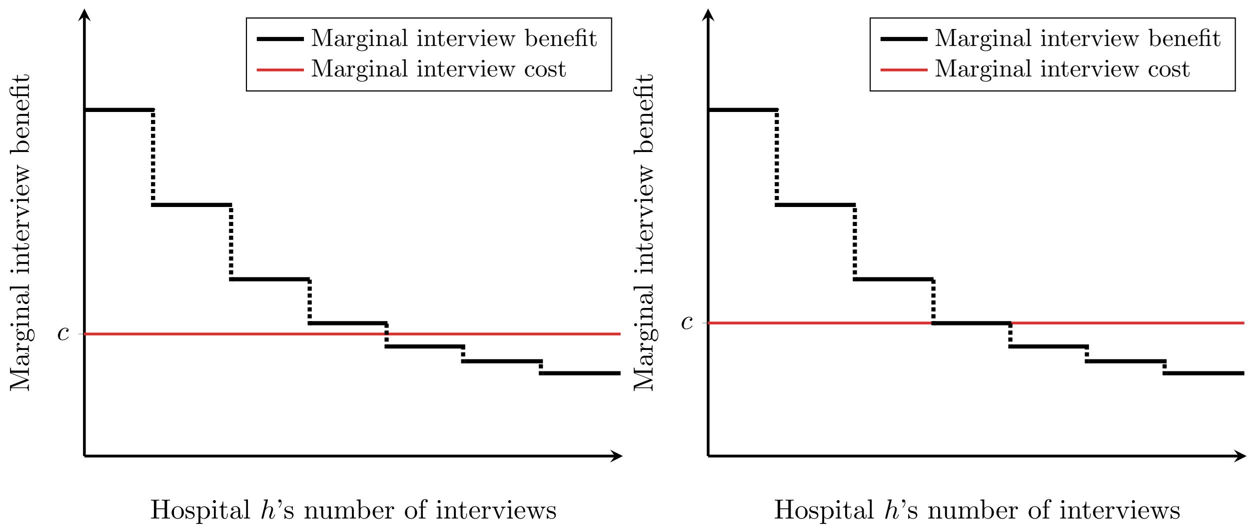

Note that Lemma 3.4 is a generalization of Lemma 2 in the appendix of Lee and Schwarz (2017). First, they indeed focus on the case in which a hospital's competitors play the same strategies, while our results are stated for all anonymous strategies (which may differ across hospitals). Second, we also generalize the result by allowing for pre-interview preferences.[4]

4 Interview Externalities

In this section, we explore the externalities that arise as hospitals increase their interview activity. We start by formally defining what we mean by "increased interview activity": First, conditional on the pre-interview information θ_h, a (mixed-) strategy σ_h specifies a probability distribution $\mathbb{P}_{\sigma_h}(\cdot|\theta_h)$ over subsets of $\mathcal{D}$ (with underlying probability space Ω). We write $\sigma'_h(\theta_h) \geq \sigma_h(\theta_h)$ if

[4] The proof of Lemma 3.5 also fixes a minor mistake in the proof of Lemma 2 in Lee and Schwarz (2017): Indeed, they claim that the probability that a hospital h is matched to the same doctor d after adding d_k to its interview list equals the sum of the probability h prefers d after interviewing d_k and the probability that h prefers d_k, but is rejected by d. However, one can show that even if h and d_k block the matching that would prevail if h did not interview d_k, it is still possible that h is matched to d after interviewing and proposing to d_k. Our proof circumvents this issue, and is based on an appropriate partitioning of the probability space and the fact that the hospital-proposing deferred-acceptance algorithm is strategy-proof for the hospitals. See the appendix for details.

there exists a partition Ψ of Ω such that for every $\psi \in \Psi$ we have $\sigma'_h(\theta_h, \psi) \supseteq \sigma_h(\theta_h, \psi)$. For any two anonymous strategies σ'_h and σ_h we write $\sigma'_h \geq \sigma_h$ if $\sigma'_h(\theta_h) \geq \sigma_h(\theta)$ for all $\theta_h \in \Theta$.[5] We write $\sigma'_{-h} \geq \sigma_{-h}$ if $\sigma'_{h'} \geq \sigma_{h'}$ for all $h' \neq h$ and $\sigma' \geq \sigma$ if $\sigma'_h \geq \sigma_h$ for all $h \in \mathcal{H}$.

We will first explore welfare externalities in Section 4.1, by first considering how a hospital's welfare is impacted if all its competitors increase their interview activity, and then looking at the welfare impact for doctors as hospitals all increase their interview activity. In Section 4.2, we reinterpret the welfare externalities imposed on hospitals in terms of the probability a hospital will match with its most preferred doctors, and we use this to further study how increased interview activity by a hospital's competitors impact the hospitals incentives to conduct more or fewer interviews. We finally combine the results from these two Sections in Section 4.3 where we explore the possibility of strategic complementarities combined with negative welfare externalities on both sides of the market.

4.1 Welfare Externalities

We start by exploring welfare externalities as they pertain to agents' expected match utilities. We then later explore the consequences of including interview costs, both for hospitals and doctors.

Theorem 4.1 *An increase in competitors' interview activity reduces the expected match utility of a hospital: For any $h \in \mathcal{H}$ and any set of doctors $S \subseteq \mathcal{D}$ interviewed by h, for any two anonymous strategies profiles $\sigma_{-h}, \sigma'_{-h}$ such that $\sigma'_{-h} \geq \sigma_{-h}$ we have $V_{\theta_h}(S, \sigma'_{-h}) \leq V_{\theta_h}(S, \sigma_{-h})$.*

Moreover, as all hospitals conduct more interviews, the doctors' expected match utilities increase: If σ and σ' are two anonymous strategies such that $\sigma' \geq \sigma$, then $V_d(\sigma') \geq V_d(\sigma)$ for all $d \in \mathcal{D}$.

The proof is in the appendix. The idea behind the theorem is that as hospitals conduct more interviews, they are more likely to arrive at the final matching stage with longer rank-order lists (ROLs). Using Assumption (A5) and the fact that we're considering anonymous strategies, by an appropriate partitioning of the probability space, I show that an increase in interview activity can be interpreted as an increase in the probability that a hospital *appends* previously unacceptable doctors to the end of its ROL at the final matching stage. Using a result by Gale and Sotomayor

[5] To illustrate, consider the case with no pre-interview information ($\Theta = \{\theta\}$) in which σ_h mixes between interviewing $\{d_1\}$ and $\{d_2\}$, each with probability 1/2, while σ'_h is the strategy that mixes between $\{d_1\}$, $\{d_2\}$, each with probability 1/4, and $\{d_1, d_2\}$ with probability 1/2. Letting $\Psi = \{\psi_1, \psi_2, \psi_3, \psi_4\}$, we can define $\sigma_h(\psi_1) = \sigma_h(\psi_2) = \{d_1\}$, $\sigma_h(\psi_3) = \sigma_h(\psi_4) = \{d_2\}$ and $\sigma'_h(\psi_1) = \{d_1\}$, $\sigma'_h(\psi_3) = \{d_2\}$ and $\sigma'_h(\psi_2) = \sigma'_h(\psi_4) = \{d_1, d_2\}$ such that $\sigma'_h(\psi) \supseteq \sigma_h(\psi)$ for all $\psi \in \Psi$.

(1985)[6], it follows that at every element of the partition of the probability space, hospitals are worse off and doctors are better off.

A consequence of the above result is that, from the hospitals' perspective, the negative welfare persist when we also consider hospitals' interview costs and we allow hospitals to best respond to an increase in their competitors' interview activity. As a result, we obtain a partial ordering of the set of equilibria in anonymous strategies:

Corollary 4.2 *Suppose hospital h's competitors increase their interview activity. Even as h best responds to its competitors' strategies, h is worse off. That is, let σ'_{-h} and σ_{-h} be two anonymous strategy profiles such that $\sigma'_{-h} \geq \sigma_{-h}$, and let $\sigma_h \in B(\sigma_{-h})$, $\sigma'_h \in B(\sigma'_{-h})$. Then*

$$V_{\theta_h}(\sigma'_h(\theta_h), \sigma'_{-h}) - c_{\mathcal{H}}|\sigma'_h(\theta_h)| \leq V_{\theta_h}(\sigma_h(\theta_h), \sigma_{-h}) - c_{\mathcal{H}}|\sigma_h(\theta_h)|$$

As a consequence, if σ, σ' are two equilibria in anonymous strategies with $\sigma' \geq \sigma$, then all hospitals prefer σ to σ'.

Proof: Using the above lemma and the definition of the best response we get

$$\begin{aligned} V_{\theta_h}(\sigma'_h(\theta_h), \sigma'_{-h}) - c_{\mathcal{H}}|\sigma'_h(\theta_h)| &\leq V_{\theta_h}(\sigma'_h(\theta_h), \sigma_{-h}) - c_{\mathcal{H}}|\sigma'_h(\theta_h)| \\ &\leq V_{\theta_h}(\sigma_h(\theta_h), \sigma_{-h}) - c_{\mathcal{H}}|\sigma_h(\theta_h)| \end{aligned}$$

■

Example 1 *Consider $\mathcal{H} = \{h_1, h_2\}$, and $|\mathcal{D}| = 2$, with $v_{hd}, v_{dh} \overset{iid}{\sim} U(0,1)$, with $\beta = 1$. Assume $c_{\mathcal{H}} \in (1/12, 1/2)$. Suppose first h_2 does not conduct any interviews; $\sigma_{h_2} = \emptyset$. If h_1 interviews one doctor at random, its expected utility is $1/2 - c_{\mathcal{H}} > 0$, while with 2 interviews h_1 gets $2/3 - 2c_{\mathcal{H}}$. Then h_1 prefers to conduct 1 interview as long as $c_{\mathcal{H}} > 1/12$. Suppose now that h_2 decides to interview one of the doctors at random; $|\sigma'_{h_2}| = 1$. If h_1 continues to conduct one interview, the two hospitals interview the same doctor 50% of the time, and each hospital hires the doctor 25% of the time. h_1's expected match utility decreases to $\frac{3}{4}1/2$. If h_1 rather interviews both doctors, h_1 achieves an expected utility of $\frac{3}{4}2/3 + \frac{1}{4}1/3 - 2c_{\mathcal{H}}$. It's in the best interest for h_1 to interview both doctors as long as $c_{\mathcal{H}} < 5/24$. In either case, h_1 is worse off as h_2 increases its interview activity.*

When only h_1 conducts one interview, doctors each match with h_1 50% of the time, with expected match utility of $1/2$. If h_2 also decides to conduct an interview, then doctors' match probability increases to 75%. Moreover, 25% of the time, a doctor will receive two proposals at the final matching stage, and achieve a match utility of $2/3$.

[6]See also Theorem 2.25 in Roth and Sotomayor (1992).

We next further explore the implications for doctors of increased interview activity by hospitals. Recall that our model allows for doctors to incur interview costs, $c_{\mathcal{D}}$, despite the fact that they are otherwise assumed to be passive players in the game. Our interest is in doctors' trade-off between improved match quality and increased interview costs. As we established in Theorem 4.1, doctors' expected match values increase the more interviews hospitals conduct. In a series of lemmas, however, we will show that the increase in doctors' expected match utility (as hospitals interview more) slows down the higher are hospitals' probabilities of getting matched. On the other hand, doctors' expected interview costs increase linearly with the number of interviews hospitals conduct: Indeed, given an anonymous strategy profile σ, all doctors are equally likely to get matched, and, for each hospital conducting a positive amount of interviews, any doctor is equally likely to receive one of the hospital's interviews.

Lemma 4.3 *Consider two anonymous strategy profiles σ and σ' such that $\sigma' \geq \sigma$. The upper bound on the increase in doctors' expected match rate is negatively related to hospitals' match rates under σ. That is,*

$$\mathbb{E}\big[|\mu(\sigma')| - |\mu(\sigma)|\big] \leq \beta \textstyle\sum_{h\in\mathcal{H}} \mathbb{E}\Big[\big(1-\mu_h(\sigma)\big)\big(|\sigma'_h| - |\sigma_h|\big)\Big] \leq \beta \displaystyle\sum_{h\in\mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big]$$

where $\underline{\mu_h}(\sigma) := \min_\theta \mathbb{E}\big[\mu_h(\sigma_h(\theta), \sigma_{-h})\big]$ is the minimum match probability hospital h achieves.

The proof is in the appendix. The easiest way to illustrate the point of the above lemma is to consider the case in which there's a single hospital, as we explore in the following example:

Example 2 *Consider a market with a single hospital, $\mathcal{H} = \{h\}$, and multiple doctors, $|\mathcal{D}| > 1$, with $\beta = 1$. If h interviews multiple doctors, the hospital will match with its most preferred doctor at the final matching stage. Hence it may be in the best interest of the hospital to conduct multiple interviews. Suppose h randomly interviews k doctors. Then each doctor will be interviewed with probability $k/|\mathcal{D}|$. In order to match with h, a doctor must also be the highest ranked doctor according to h's post-interview preferences. This happens with probability $1/k$. Therefore, the expected utility of a doctor, as a function of the number of interviews, k, that h decides to conduct, is given by*

$$\frac{v_{dh}}{|\mathcal{D}|} - \frac{k}{|\mathcal{D}|} c_{\mathcal{D}}$$

Note that the expected match utility of the doctors does not increase in the number of interviews h conducts (as long as $k \geq 1$), while doctors' expected interview costs do. Moreover, as long as

$c_{\mathcal{D}} < \underline{v_{\mathcal{D}}}/k$, it's individually rational for the doctors to accept the interview, despite the fact that they would be better off if h conducted fewer interviews.

As we observed in Example 1, the effect on doctors of increased interview activity by hospitals can be divided into two components: Doctors probability of being matched, and doctors' probability of receiving multiple proposals during the final matching stage. The following lemma states the probability of the latter is small when hospitals' match probabilities are already high:

Lemma 4.4 *Suppose σ' and σ are anonymous, with $\sigma' \geq \sigma$. As hospitals increase their interview activity, for any $d \in \mathcal{D}$, conditional on being matched under σ, the upper bound on the increase in d's expected match utility is negatively related to hospitals' match rates under σ. That is,*

$$\begin{aligned}\mathbb{E}\big[v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)} \,|\, \mu_d(\sigma) \neq \emptyset\big] &\leq \bar{v}_{\mathcal{D}}\beta\mathbb{E}\big[\sum_{h\in\mathcal{H}} \big(1 - |\mu_h(\sigma)|\big)\big(|\sigma'_h| - |\sigma_h|\big)\big] \\ &\leq \bar{v}_{\mathcal{D}}\beta \sum_{h\in\mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big]\end{aligned}$$

where $|\underline{\mu_h}(\sigma)| := \min_\theta \mathbb{E}\big[|\mu_h(\sigma_h(\theta), \sigma_{-h})|\big]$.

The proof is in the appendix. The intuition behind the proofs of both Lemma 4.3 and Lemma A.2 follows the logic behind Theorem 4.1: As hospitals increase their interview activity, they are more likely to show up at the final matching stage with longer rank-order lists. By Assumption (A5) and the fact that we're only considering anonymous strategy profiles, the increased interview activity can be interpreted as the appending of doctors to the bottom of hospitals rank-order lists. However, a hospital appending doctors to the end of its rank-order lists only has a chance of changing the match outcome if the hospital would otherwise be unmatched: Indeed, if all hospitals are already able to match, then the deferred acceptance algorithm ends at exactly the same step as before even after hospitals append doctors to the end of their ROLs. It turns out that the extent to which hospitals are more likely to receive multiple offers at the final matchin stage depends on the market tightness, i.e. the fraction of doctors relative to hospitals, as the following example illustrates:

Example 3 *Consider again the setup in Example 1, with $\mathcal{D} = \{d_1, d_2\}$, $v_{hd}, v_{dh} \overset{iid}{\sim} U(0,1)$ and $\beta = 1$, but assume now that there are more hospitals than doctors, e.g. $|\mathcal{H}| = 3$. If all hospitals conduct 1 interview each, then a doctor will receive two proposals at the final matching stage with probability $\frac{1}{2}\binom{3}{2}\left(\frac{1}{2}\right)^3 = \frac{3}{16}$ and will receive three proposals with probability $\frac{1}{16}$. Note that hospitals'* *probability of getting matched is less than 2/3. If instead hospitals all conduct 2 interviews each, both doctors will always receive at least two proposals at the final matching stage: Even if some of the hospitals first propose to d_1, only one of them can be accepted, and hence any hospital rejected by d_1 will then make a proposal to d_2.*

Having established respective upper bounds on doctors' match rates and match utility from increased hospital interview activity, we can combine the two lemmas to establish an upper bound for the increase in doctors' unconditional match utility:

Corollary 4.5 *As hospitals increase their interview activity, the upper bound on the increase in doctors' expected match utilities is negatively related to hospitals' match rates: Consider two anonymous strategy profiles σ' and σ such that $\sigma' \geq \sigma$. Then*

$$\mathbb{E}\big[v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)}\big] \leq \frac{\beta}{|\mathcal{D}|}\Big[\overline{v}_{\mathcal{D}} + (\overline{v}_{\mathcal{D}} - \underline{v}_{\mathcal{D}})\mathbb{E}[|\mu(\sigma)|]\Big] \sum_{h\in\mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big]$$

where $|\underline{\mu_h}(\sigma)| := \min_\theta \mathbb{E}\big[|\mu_h(\sigma_h(\theta), \sigma_{-h})|\big]$.

As we pointed out in Example 2, doctors' expected interview costs increase linearly in hospitals' interview activity, while the previous results established that the increase in doctors' expected match utility is negatively related to hospitals' match probabilities. Example 2 also illustrated that it's possible for doctors to be worse off as hospitals increase their interview activity, while it still being individually rational for them to accept all interviews, a result we generalize below:

Proposition 4.6 *Consider two anonymous strategy profiles σ and σ' such that $\sigma' \geq \sigma$. Let $|\underline{\mu_{\mathcal{H}}}|$ denote the lowest match probability for any hospital-type under either σ or σ'. If $|\underline{\mu_{\mathcal{H}}}| > \frac{\beta|\mathcal{D}|}{\beta|\mathcal{D}|+1}$, then there exist constants γ and $\underline{c_{\mathcal{D}}} < \overline{c_{\mathcal{D}}}$ such that if $\overline{v_{\mathcal{D}}} > \gamma(\overline{v_{\mathcal{D}}} - \underline{v_{\mathcal{D}}})$*

1. *if $c_{\mathcal{D}} < \overline{c_{\mathcal{D}}}$, it is individually rational for doctors to accept all interviews under both σ and σ'.*

2. *if $c_{\mathcal{D}} > \underline{c_{\mathcal{D}}}$, doctors are worse off under σ' than under σ.*

The proof is in the appendix. Intuitively, when γ is large, the difference $\overline{v_{\mathcal{D}}} - \underline{v_{\mathcal{D}}}$ is small, and doctors only care about the probability of getting matched. If all hospitals match with very high probability under both σ and σ', then the benefits to doctors from the increased interview activity under σ' is small, while their interview costs increase proportionally to $\mathbb{E}[|\sigma'| - |\sigma|]$.

4.2 Strategic externalities

As we saw in Example 1, an increase in the interview activity by a hospital's competitors affected the probability that the hospital matched with its first and second highest ranked doctor. We also illustrated that, for certain interview costs, the increase in competitor interview activity changed a hospital's optimal interview strategy In this section, we further generalize these observations and draw the connection between the welfare externalities and a hospital's match probabilities and hence the incentives to change its interview strategy. We begin by introducing the following notation:

Definition 4.7 *For any h, θ_h, for any set $S \subseteq \mathcal{D}$ of doctors interviewed by h, for any set $\overline{S} \subseteq S$ of doctors found acceptable by h after interviewing S, and any ranking P over $\overline{S}$, let $q_{h,j}(\sigma, S, \overline{S}, P)$ denote the probability that h matches to its j-th highest ranked doctor in $\overline{S}$, conditional on σ_{-h}:*

$$q^{(j)}_{h,\theta_h}(\sigma_{-h}, S, \overline{S}, P) := \mathbb{P}\Big(\mu(h) = P^{(j)} \,\big|\, v_{hd} > 0\,\forall\, d \in \overline{S},\, v_{hd} \leq 0\,\forall\, d \in S \setminus \overline{S},\, P,\, \sigma_{-h}\Big)$$

Drawing once again on Example 1, we noticed that when h_2 conducted no interviews, then h_1 would match with its post-interview highest ranked doctor with probability 1, regardless of the identity of this doctor. Similarly, when h_2 conducted one interview at random, the probability that h_1 would match with its post-interview highest ranked doctor fell to $\frac{3}{4}$, again irrespective of the identity of this doctor. It turns out that when hospitals play anonymous strategies, this property holds more generally:

Lemma 4.8 *Assume σ_{-h} is anonymous. Then $q_j(\sigma_{-h}, \cdot, \cdot)$ does not depend on the set of doctors interviewed by h, nor on the set of doctors found acceptable by h, nor h's ranking P of the acceptable doctors: For any $S, S' \subseteq \mathcal{D}$, for any $\overline{S} \subseteq S$, $\overline{S}' \subseteq S'$, and any rankings P and P' over $\overline{S}$ and $\overline{S}'$ respectively, for every $j \leq \min(|\overline{S}|, |\overline{S}'|)$ we have $q^{(j)}_{h,\theta_h}(\sigma_{-h}, S, \overline{S}, P) = q^{(j)}_{h,\theta_h}(\sigma_{-h}, S', \overline{S}', P')$*

Essentially, the Lemma says that the probability a hospital matches to its j-th choice does not depend on the identity of this doctor, nor the way in which the hospital ranks all of its acceptable doctors. Of course, the probability that a hospital matches to its j-th choice if it only found $k < j$ doctors acceptable is zero. Keeping this in mind, with some abuse of notation, we can drop the dependence of $q^{(j)}_{h,\theta_h}$ on S, $\overline{S}$ and P, and simply write it as a function of the competitors' strategies; $q^{(j)}_{h,\theta_h}(\sigma_{-h})$. In the following, it will be useful to also specify the conditional probabilities, i.e. the probability that a hospital matches with its j-th choice, conditional on not matching with any of its $j-1$ higher ranked choices:

$$p^{(j)}_{h,\theta_h}(\sigma_{-h}) := \frac{q^{(j)}_{h,\theta_h}(\sigma_{-h})}{1-\sum_{i<j} q^{(i)}_{h,\theta_h}(\sigma_{-h})} \quad \Rightarrow \quad q^{(j)}_{h,\theta_h}(\sigma_{-h}) = \prod_{i<j}\Big(1-p^{(i)}_{h,\theta_h}(\sigma_{-h})\Big)p^{(j)}_{h,\theta_h}(\sigma_{-h})$$

Recall that we denote by $v_{h(i,\underline{S})}$ the i-th highest match value among the doctors in the set $\underline{S}$ (all assumed acceptable). Combining this with the notation above, we can now rewrite hospital h's expected match utility, given h's interview set S, and given that all other hospitals play an anonymous strategy σ_{-h}, as

$$\begin{aligned} V_{\theta_h}(S,\sigma_{-h}) &= \sum_{\underline{S}\subset S}\prod_{\overline{d}\notin\underline{S}}(1-\beta_{\theta_{h\underline{d}}})\prod_{\underline{d}\in\underline{S}}\beta_{\theta_{h\overline{d}}}\sum_{i=1}^{|\underline{S}|} q^{(i)}_{h,\theta_h}(\sigma_{-h})\mathbb{E}\big[v_{h(i,\underline{S})}\big] \\ &= \sum_{\underline{S}\subset S}\prod_{\overline{d}\notin\underline{S}}(1-\beta_{\theta_{h\underline{d}}})\prod_{\underline{d}\in\underline{S}}\beta_{\theta_{h\overline{d}}}\sum_{i=1}^{|\underline{S}|}\prod_{j<i}\Big(1-p^{(j)}_{h,\theta_h}(\sigma_{-h})\Big)p^{(i)}_{h,\theta_h}(\sigma_{-h})\mathbb{E}\big[v_{h(i,\underline{S})}\big] \\ &= \sum_{\underline{S}\subset S}\prod_{\overline{d}\notin\underline{S}}(1-\beta_{\theta_{h\underline{d}}})\prod_{\underline{d}\in\underline{S}}\beta_{\theta_{h\overline{d}}}\mathbb{E}\big[v_{h(1,\underline{S})}\big] + \sum_{i=1}^{|S|}\prod_{j\leq i}\big(1-p^{(j)}_{h,\theta_h}(\sigma_{-h})\big)\Bigg(\\ &\sum_{\substack{\underline{S}\subset S\\ |\underline{S}|>i}}\prod_{\overline{d}\notin\underline{S}}(1-\beta_{\theta_{h\underline{d}}})\prod_{\underline{d}\in\underline{S}}\beta_{\theta_{h\overline{d}}}\mathbb{E}\big[v_{h(i+1,\underline{S})}-v_{h(i,\underline{S})}\big] - \sum_{\substack{\underline{S}\subset S\\ |\underline{S}|=i}}\prod_{\underline{d}\in\underline{S}}(1-\beta_{\theta_{h\underline{d}}})\prod_{\overline{d}\notin\underline{S}}\beta_{\theta_{h\overline{d}}}\mathbb{E}\big[v_{h(i,\underline{S})}\big]\Bigg) \end{aligned}$$

Having introduced this notation, we can now reinterpret the welfare externalities analyzed in the previous section. Indeed, the statement "h is worse off when its competitors increase their interview activity" can be reinterpreted in terms of h's match probabilities:

Lemma 4.9 *Assume σ_{-h} and σ'_{-h} are anonymous, with $\sigma'_{-h}\geq\sigma_{-h}$. Then for all k, the probability h matches with any of its k highest choices under σ'_{-h} is lower than under σ_{-h}. That is, $\sum_{j\leq k} q^{(j)}_{h,\theta_h}(\sigma'_{-h})\leq\sum_{j\leq k} q^{(j)}_{h,\theta_h}(\sigma_{-h})$ and $\prod_{j\leq k}\big(1-p^{(j)}_{h,\theta_h}(\sigma'_{-h})\big)\leq\prod_{j\leq k}\big(1-p^{(j)}_{h,\theta_h}(\sigma_{-h})\big)$ for all k.*

Proof: In the proof of Theorem 4.1, we show that σ'_{-h} results in a distribution of rank-order lists that can be seen as appending (at the end of the list) previously unacceptable doctors to the rank-order lists resulting from σ_{-h}. By an appropriate partitioning of the probability space (over agents' strategies and preferences), we then argue that h is worse off under σ'_{-h} than under σ_{-h} at every element of that partition. Specifically, h is matched to a doctor further down on their rank-order list. Therefore, for all elements of the partition in which h is not matched to any of its k highest ranked acceptable doctors under σ_{-h}, h is not matched to any of its k highest ranked acceptable doctors under σ'_{-h}. Since this holds for all k, the result follows. ∎

Our interest is in understanding how h's benefit of adding a doctor d to an interview set S is affected by an increase in its competitors' interview activity, i.e. evaluating

$$V_{\theta_h}\big(S\cup\{d\},\sigma'_{-h}\big)-V_{\theta_h}\big(S,\sigma'_{-h}\big)-\Big(V_{\theta_h}\big(S\cup\{d\},\sigma_{-h}\big)-V_{\theta_h}\big(S,\sigma_{-h}\big)\Big)$$

Keeping in mind that the benefit of adding d to the interview list is only relevant in the cases where d is found acceptable by h, some algebra reveals that for $S\subset\mathcal{D}\setminus\{d\}$ we have

$$\begin{aligned}
&V_{\theta_h}\big(S\cup\{d\},\sigma_{-h}\big)-V_{\theta_h}\big(S,\sigma_{-h}\big)\\
&=\beta_{\theta_{hd}}\Bigg\{\sum_{\underline{S}\subset S}\prod_{\underline{d}\in\underline{S}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}}\beta_{\theta_{h\overline{d}}}\mathbb{E}\big[v_{h(1,\underline{S}\cup\{d\})}-v_{h(1,\underline{S})}\big]+\sum_{i=1}^{|\mathcal{H}|-1}\prod_{j\le i}\Big(1-p^{(j)}_{h,\theta_h}(\sigma_{-h})\Big)\Big(\\
&\qquad\sum_{\substack{\underline{S}\subset S\\|\underline{S}|>i}}\prod_{\underline{d}\in\underline{S}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}}\beta_{\theta_{h\overline{d}}}\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{d\})}-v_{h(i+1,\underline{S})}-\big(v_{h(i,\underline{S}\cup\{d\})}-v_{h(i,\underline{S})}\big)\big]\\
&\qquad+\sum_{\substack{\underline{S}\subset S\\|\underline{S}|=i}}\prod_{\underline{d}\in\underline{S}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}}\beta_{\theta_{h\overline{d}}}\,\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{d\})}-\big(v_{h(i,\underline{S}\cup\{d\})}-v_{h(i,\underline{S})}\big)\big]\\
&\qquad-\sum_{\substack{\underline{S}\subset S\\|\underline{S}|=i-1}}\prod_{\underline{d}\in\underline{S}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}}\beta_{\theta_{h\overline{d}}}\,\mathbb{E}\big[v_{h(i,\underline{S}\cup\{d\})}\big]\Big)\Bigg\}
\end{aligned}$$

To further study the strategic externalities, we will impose Assumption (A5*) and (A7)-(A8) stated in Section 3. Neither assumption is implied by the previously stated assumptions. However, Lee and Schwarz (2017) mention that log-concavity of the distribution F is a sufficient condition for their assumption (A1). It turns out that log-concavity also implies (A7), as is shown in the appendix.

Note that by Assumptions (A7)-(A8), the terms $\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{d\})}-v_{h(i+1,\underline{S})}-\big(v_{h(i,\underline{S}\cup\{d\})}-v_{h(i,\underline{S})}\big)\big]$ and $\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{d\})}-\big(v_{h(i,\underline{S}\cup\{d\})}-v_{h(i,\underline{S})}\big)\big]$ are all positive.

Using the above lemma and the additional distributional assumptions, the strategic effect of an increase in the interview activity by h's competitors can now be decomposed into two parts; a negative **competition effect**

$$-\beta_{\theta_{hd}}\sum_{\substack{\underline{S}^{i-1}\subset S\\|\underline{S}^{i-1}|=i-1}}\prod_{\underline{d}\in\underline{S}^{i-1}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}^{i-1}}\beta_{\theta_{h\overline{d}}}\,\mathbb{E}\big[v_{h(i,\underline{S}^{i-1}\cup\{d\})}\big]$$

and a positive **probability reallocation effect**

$$\beta_{\theta_{hd}}\Bigg(\sum_{\substack{\underline{S}^{>i}\subset S\\|\underline{S}^{>i}|>i}}\prod_{\underline{d}\in\underline{S}^{>i}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}^{>i}}\beta_{\theta_{h\overline{d}}}\mathbb{E}\big[v_{h(i+1,\underline{S}^{>i}\cup\{d\})}-v_{h(i+1,\underline{S}^{>i})}-\big(v_{h(i,\underline{S}^{>i}\cup\{d\})}-v_{h(i,\underline{S}^{>i})}\big)\big]$$

$$+\sum_{\substack{\underline{S}^{i}\subset S\\|\underline{S}^{i}|=i}}\prod_{\underline{d}\in\underline{S}^{i}}\big(1-\beta_{\theta_{h\underline{d}}}\big)\prod_{\overline{d}\notin\underline{S}^{i}}\beta_{\theta_{h\overline{d}}}\,\mathbb{E}\big[v_{h(i+1,\underline{S}^{i}\cup\{d\})}-\big(v_{h(i,\underline{S}^{i}\cup\{d\})}-v_{h(i,\underline{S}^{i})}\big)\big]\Bigg)$$

The competition effect captures the idea that when the other hospitals increase their interview activity, h is less likely to be matched to any of its k-th highest ranked doctors, for $k<|\mathcal{H}|$. This unambiguously reduces h's value of interviewing doctors who are likely to be ranked among its k highest ranked after the interviews are conducted. The probability reallocation effect represents the increased benefit of submitting longer rank-order lists when the competing hospitals conduct more interviews. Indeed, the more h's competitors interview, the more likely h is to match with a doctor far down on its rank-order list, which increases h's incentives to improve the expected match value of its lower ranked doctors. The two effects are illustrated both separately and combined in Figure 2 below. Note that the probability reallocation effect is always zero when the hospital is conducting only one interview.

Further algebra shows that the marginal interview benefit can be rewritten as

$$V_{\theta_h}\big(S\cup\{d\},\sigma_{-h}\big)-V_{\theta_h}\big(S,\sigma_{-h}\big)$$

$$=\beta\Bigg\{\sum_{\underline{S}\subset S}(1-\beta)^{|S|-|\underline{S}|}\beta^{|\underline{S}|}\,\mathbb{E}\big[v_{h(1,\underline{S}\cup\{d\})}-v_{h(1,\underline{S})}\big]+\sum_{i=1}^{|\mathcal{H}|-1}\prod_{j\le i}\Big(1-p^{(j)}_{h,\theta_h}(\sigma_{-h})\Big)\Bigg($$

$$\sum_{\substack{\underline{S}\subset S\\|\underline{S}|>i}}(1-\beta)^{|S|-|\underline{S}|}\beta^{|\underline{S}|}\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{d\})}-v_{h(i+1,\underline{S})}-\big(v_{h(i,\underline{S}\cup\{d\})}-v_{h(i,\underline{S})}\big)\big]$$

$$+\frac{(1-\beta)^{|S|-i}\beta^{i-1}}{i}\sum_{\substack{\underline{S}\subset S\\|\underline{S}|=i-1}}\Bigg[$$

$$\beta\sum_{\underline{d}\in S\setminus\underline{S}}\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{\underline{d},d\})}-\big(v_{h(i,\underline{S}\cup\{\underline{d},d\})}-v_{h(i,\underline{S}\cup\{\underline{d}\})}\big)\big]-(1-\beta)i\mathbb{E}\big[v_{h(i,\underline{S}\cup\{d\})}\big]\Bigg]\Bigg)\Bigg\}$$

By Assumption (A8), for each $i\le|\mathcal{H}|-1$, the last term (in which $|\underline{S}|=i-1$) can be bounded

Figure 2: **The competition and probability reallocation effect, with two hospitals**
The competition effect reduces the benefit of any interview, and can be illustrated as a negative "shift" of the marginal benefit curve. The probability reallocation effect increases the benefit of long rank-order lists, and hence many interviews. The net effect (right panel) is negative for the "first" interviews but is positive for doctors ranked low according to pre-interview information θ.

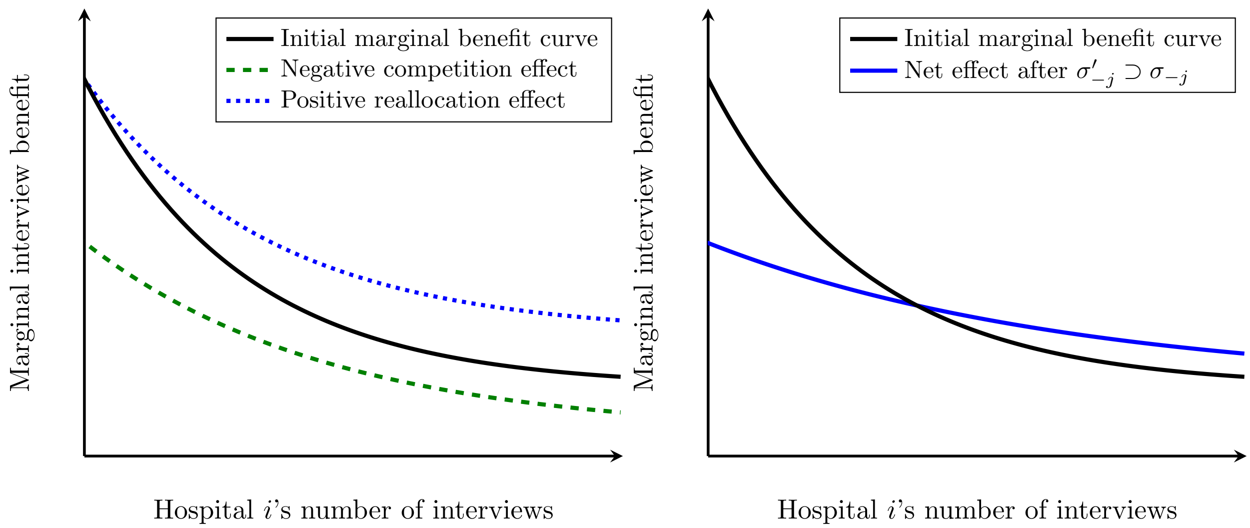

from below as

$$\beta \sum_{\underline{d} \in S \setminus \underline{S}} \mathbb{E}\left[v_{h(i+1, \underline{S} \cup \{\underline{d}, d\})} - \left(v_{h(i, \underline{S} \cup \{\underline{d}, d\})} - v_{h(i, \underline{S} \cup \{\underline{d}\})}\right)\right] - (1-\beta) i \mathbb{E}\left[v_{h(i, \underline{S} \cup \{d\})}\right]$$
$$\geq \left[\beta\left(|S| + 1 - i\right)\varepsilon - (1-\beta) i\right] \mathbb{E}\left[v_{h(i, \underline{S} \cup \{d\})}\right]$$

We can now combine Lemma 4.9 with the above expressions to determine how a hospital's incentives to either increase or decrease its own number of interviews change as its competitors increase their interview activity. Specifically, our notion of strategic externalities will be based on the properties of h's best-response correspondences: In the following, we will consider the strong set order induced by the set inclusion order (where meet is intersection and join is union).

Definition 4.10 *A best-response correspondence $B(\cdot)$ exhibits* strategic complementarities *(*strategic substitutes*) if $B(\cdot)$ is increasing (decreasing) in the strong set order. A game exhibits* best-response strategic complementarities *(*best-response strategic substitutes*) if all players' best-response correspondences exhibit* strategic complementarities *(*strategic substitutes*).*

A consequence of the above lemma, combined with the above expression for the benefit of a marginal interview is the following:

Proposition 4.11 *Restrict attention to anonymous strategies, and suppose Assumptions (A5*), (A7)-(A8) hold. Then there exists a constant $0 < \bar{c}$ such that*

1. *if interview costs satisfy $c_{\mathcal{H}} > \bar{c}$, the game exhibits best-response strategic substitutes.*
2. *if $\beta\varepsilon|\mathcal{D}| \geq (1-\beta+\varepsilon\beta)|\mathcal{H}|$, then there exists a constant $\underline{c} \leq \bar{c}$ such that if interview costs satisfy $c_{\mathcal{H}} < \underline{c}$, then the game exhibits best-response strategic complementarities.*

The proof is in the appendix. The idea behind Part 2 of the above Proposition is that when the marginal interview costs are low, hospitals will always want to conduct many interviews. Moreover, when the number of interviews are sufficiently high, then the positive probability reallocation effect always outweighs the negative competition effect, leading to the strategic complementarity in interviews. To see this, note that for each $S \subset \mathcal{D}$, there are $\binom{|S|}{i-1}$ sets of size $i-1$, while there are $\binom{|S|}{i} = \binom{|S|}{i-1}\frac{|S|+1-i}{i}$ sets of size i. As $|S|$ grows, i.e. the more interviews h conducts, the ratio of subsets of S of size i relative to subsets of size $i-1$ increases. By Assumptions (A7)-(A8), the probability reallocation effect "eventually" grows faster than the competition effect, as $|S|$ becomes large.

The following example illustrates part 2 of the above Proposition:

Example 4 *Assume $supp(F_\theta^+) = \theta$ for all θ, and that $\theta_{hd} \neq \theta_{hd'}$ for all $d \neq d'$. Moreover, assume $\beta_{\theta_{hd}} \equiv \beta$ for all θ_{hd}. In this case, for any set $\underline{S}$, such that $|\underline{S}| = k$, if $\theta_{hd'} < \theta_{hd}$ for all $d \in \underline{S}$, then $\mathbb{E}\big[v_{h(j,\underline{S}\cup\{d\})} - v_{h(j,\underline{S})}\big] = 0$ for all $j \leq k$. As a result, the inequality in (A6) holds with equality for $\varepsilon = 1$. Using this, the marginal interview benefit reduces to*

$$
\begin{aligned}
& V\big(S \cup \{d'\}, \sigma_{-h}\big) - V\big(S, \sigma_{-h}\big) \\
&= \beta\Bigg\{(1-\beta)^n\theta_{hd'} + \sum_{i=1}^{\min(|S|+1,|\mathcal{H}|-1)} \prod_{j\leq i} q_j(\sigma_{-h})\Bigg(\\
&\qquad \binom{n}{i}(1-\beta)^{n-i}\beta^i\theta_{hd} - \binom{n}{i-1}(1-\beta)^{n+1-i}\beta^{i-1}\theta_{hd}\Bigg)\Bigg\} \\
&= \beta\Bigg\{(1-\beta)^n\theta_{hd'} + \sum_{i=1}^{\min(|S|+1,|\mathcal{H}|-1)} \prod_{j\leq i} q_j(\sigma_{-h})\Bigg(\\
&\qquad \frac{\theta_{hd}}{i}\binom{n}{i-1}(1-\beta)^{n-i}\beta^{i-1}\Big[\beta(n+1)-i\Big]\Bigg)\Bigg\}
\end{aligned}
$$

Figure 3: **Strategic substitutes (left) and complementarities (right)**

The direction of the strategic externalities depends on hospital interview costs. When costs are high, a hospital finds it optimal to reduce its number of interviews when its competitors increase their interview activity, while the hospital prefers to increase its own number of interviews when interview costs are low (and the number of doctors is sufficiently large)

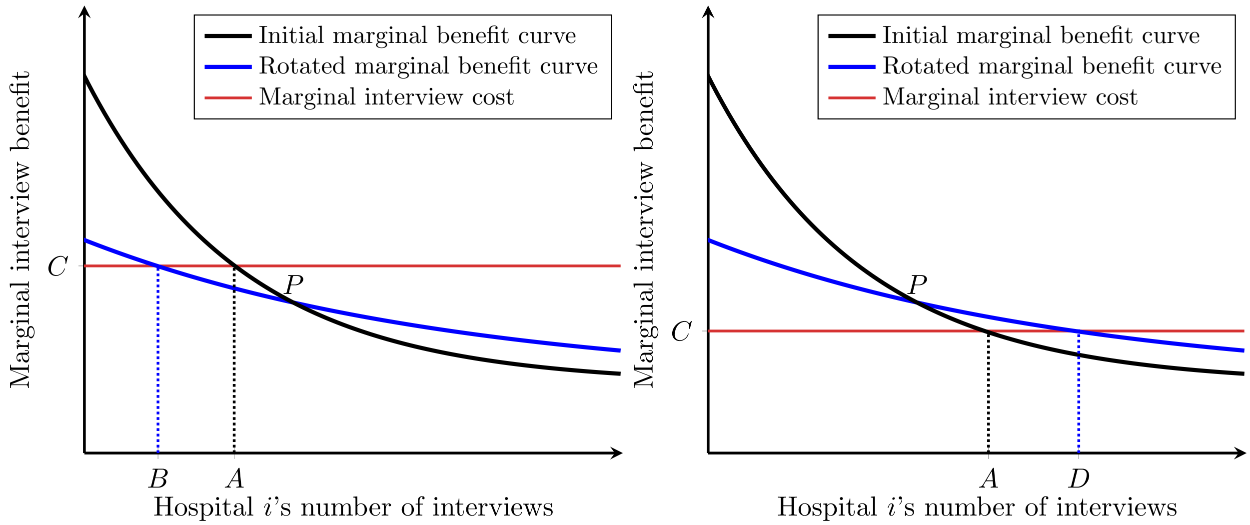

As a result, if $c_{\mathcal{H}}$ is such that h ever finds it optimal to interview a set $S \cup \{d\}$ satisfying $\beta|S \cup \{d\}| \geq |\mathcal{H}| - 1$, then an increase in interview activity by h's competitors will never cause h to best-respond by reducing its own number of interviews.

4.3 Strategic complementarities and negative welfare externalities

In the previous two sections, we analyzed both the welfare externalities that agents on both sides of the market incur as hospitals increase their interview activity, as well as the strategic externalities to which hospitals are exposed when their competitors interview more. Combining these results, we can characterize the set of (pure) strategy equilibria that emerge when hospitals' interviews are sufficiently small. Moreover, these equilibria can be ranked according to the Pareto criterion:

Theorem 4.12 *Restrict attention to equilibria in anonymous strategies, and suppose Assumptions (A5*) and (A7)-(A8) all hold. Suppose $\mathcal{H}$ and $\mathcal{D}$ satisfy $\beta\varepsilon|\mathcal{D}| \geq (1 - \beta + \varepsilon\beta)|\mathcal{H}|$. There exists a threshold $\overline{c_{\mathcal{H}}}$ and a constant $\gamma > 0$ such that if $c_{\mathcal{H}} < \overline{c_{\mathcal{H}}}$*

1. *The set of pure equilibria forms a complete lattice ordered by the number of interviews each hospital conducts. Hospital welfare is decreasing in the equilibrium number of interviews.*

2. *Suppose, in addition, that* $\overline{v}_{\mathcal{D}} \geq \gamma(\overline{v}_{\mathcal{D}} - \underline{v}_{\mathcal{D}})$. *Then there exists* $\underline{c_{\mathcal{D}}} < \overline{c_{\mathcal{D}}}$ *such that*

   (a) *if* $c_{\mathcal{D}} > \underline{c_{\mathcal{D}}}$, *then doctors' welfare is decreasing in the equilibrium number of interviews.*

   (b) *if* $c_{\mathcal{D}} \leq \underline{c_{\mathcal{D}}}$, *it is individually rational for doctors' to accept all interviews.*

Proof:

1. From Proposition 4.11 we know the game exhibits strategic complementarities, meaning the best-response correspondences are increasing. By the Knaster-Tarski fixed-point theorem, the set of fixed points of $\big(B_h(\sigma_{-h})\big)_{h \in \mathcal{H}}$ forms a complete lattice. These fixed-points constitute (pure) equilibria in anonymous strategies of the interview game. For two fixed points/equilibria σ and σ' such that $\sigma' \geq \sigma$, all hospitals weakly conduct more interviews under σ' than under σ. From Corollary 4.2 it follows that all hospitals are worse off under σ' than under σ.

2. By Proposition 4.6, it's sufficient to show that $E[|\mu(\sigma_h(\theta_h))|] > \frac{|\mathcal{D}|}{|\mathcal{D}|+1}$ for all hospital-types (h, θ_h) in all equilibria σ. In the appendix, I show that, conditional on finding $|\mathcal{H}|-1$ doctors acceptable, any hospital is matched with a probability that exceeds $\frac{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}-1}{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}}$. Letting $A(\sigma_h)$ denote the number of doctors h finds acceptable when playing σ_h, then for each θ, $A(\sigma_h(\theta))$ follows a binomial distribution with parameters β and $|\sigma_h|$. Let $|\underline{\sigma_h}|$ denote the minimum number of interviews conducted by h under σ. By Hoeffding's inequality, we have

$$\mathbb{P}\big(A(\sigma_h) \geq |\mathcal{H}| - 1\big) = 1 - \exp\big(-2\eta(\sigma)^2 |\underline{\sigma}_h|\big)$$

   with $\eta(\sigma) := \frac{\beta|\underline{\sigma}_h| - |\mathcal{H}|+1}{|\underline{\sigma}_h|}$. Therefore, as long as $\big[1 - \exp\big(-2\eta^2|\underline{\sigma}_h|\big)\big]\frac{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}-1}{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}} > \frac{|\mathcal{D}|}{|\mathcal{D}|+1}$, the conditions of Proposition 4.6 are satisfied. By Part 1, $|\underline{\sigma}_h|$ increases as we move up the lattice of equilibria, hence it's sufficient to verify that the condition is met for the equilibrium involving the smallest number of interviews. By setting $c_{\mathcal{H}}$ sufficiently low, then either (i) all hospitals will conduct enough interviews in all equilibria such that the above condition is satisfied, or (ii) all hospitals interview all doctors in the market, in which case there's a unique equilibrium and only doctors' individual rationality constraint needs to be satisfied.

   ■

One significant aspect of the above result is that, contrary to classical results in the matching literature, preferences on opposing market sides are not conflicting. Indeed, the final matching stage of our model includes the use of a stable matching mechanism. However, the set of stable matchings have the property that if you make one side of the market worse off, then the other side of the market is necessarily better off (Knuth (1976)). However, as we have just demonstrated, this central result in the theory of stable matchings no longer holds when we also consider interview costs: It is indeed possible to make both sides of the market worse off. Due to the presence of the interview stage, despite the use of the stable matching mechanism in the final matching stage, the resulting match need not be stable in the classical sense: Hospitals may optimally choose not to interview all doctors. Moreover, even if hospitals do interview all doctors, and the resulting match indeed is stable, this does not imply that hospitals could not be made even better off when we account for their interview costs. In short, the introduction of the costly interview stage implies that some of the central results from the matching literature no longer hold when we take an extended view of the entire matching process, which includes the costly process through which agents learn about their preferences.

It is also worth emphasizing that the interview allocations described in Theorem 4.12 are not simply imposed on the agents: These are equilibrium interview allocations that result from hospitals optimally choosing their interview strategies in response to their competitors' strategies. Furthermore, while doctors play a passive role in the interview game, under the conditions of the Theorem, it is still individually rational for doctors to accept all the interviews they are offered.

An immediate consequence of Theorem 4.12 is that, in the case in which hospitals coordinate on the equilibrium that involves the most interviews, then agents on both sides of the market can be made better off with coordinated reduction in interview activity. We will explore this and related points when we discuss the role of market design interventions in Section 6.

Before we proceed with our analysis of comparative statics in Section 5, we will briefly discuss the results in some more detail: Specifically, why is the characterization of strategic complementarities so involved? Does there exist a more straightforward model under which strategic complementarities in hospitals' interview decisions would hold more generally, and would not require additional assumptions on market tightness ($|\mathcal{D}|/|\mathcal{H}|$) and hospitals' interview costs? For instance, in supermodular games, the characterization of strategic complementarities is more straightforward, since it's essentially build into agents' payoff functions. The issue turns out to combine this type of strategic complementarity with negative welfare externalities. Specifically, supermodular games cannot exhibit negative welfare externalities. Since we are considering a game with negative welfare externalities, this explains why hospitals' payoff functions are not supermodular in their own and competitors' interview decisions:

Observation 4.13 *Consider a strategic form game* $\big(I, (S_i)_{i \in I}, (u_i)_{i \in I}\big)$ *where the strategy space has a partial order for each* $i \in I$*, and each player has an "lowest" action,* $a_i^{(0)}$*, that guarantees a constant payoff regardless of others' actions:* $a_i^{(0)} \leq a_i$ *for all* $a_i \in S_i$ *and* $u_i(a_i^{(0)}, a_{-i}) \equiv \overline{u}_i$ *for all* a_{-i}*. Suppose the game exhibits negative welfare externalities; i.e. an increase in the action by other players make all other players worse off. Then the players' payoff functions do not exhibit increasing differences.*

Proof: Consider an action $a_i \neq a_i^{(0)}$ and a sequence of actions $(a_i^{(k)})_{k=0}^n$ such that $a_i^{k-1} \leq a_i^{(k)}$ for $k = 1, ..., n$. Write $u(a_i, a_{-i}) - \overline{u}_i = \sum_{k=1}^n u_i(a_i^{(k)}, a_{-i}) - u_i(a_i^{(k-1)}, a_{-i})$. Consider $a'_{-i} \geq a_{-i}$, then

$$u(a_i, a'_{-i}) - u(a_i, a_{-i}) = \sum_{k=1}^{n} u_i(a_i^{(k)}, a'_{-i}) - u_i(a_i^{(k-1)}, a'_{-i}) - \big(u_i(a_i^{(k)}, a_{-i}) - u_i(a_i^{(k-1)}, a_{-i})\big)$$

Suppose u_i satisfies increasing differences. Then each term on the right-hand side is positive. However, negative welfare externalities imply that the left-hand side is negative, leading to a contradiction. ■

5 Comparative statics

The previous sections have analyzed the structure and properties of hospitals' optimal strategies and best-response correspondences, in addition to determining the qualitative nature of the welfare externalities to which agents are exposed. This analysis considered the case in which the economic environment was held constant. We next turn to an analysis of comparative statics, in which we will explore the effect of changing certain features of the economic environment, in particular the effect of changing agents' interview costs or the number of participants in the market.

5.1 Varying interview costs

We start our analysis by exploring the effect of changing agents' interview costs. Such an analysis seems particularly relevant in light of recent changes to the interview technology: Traditionally, most interviews were conducted in person, with candidates (doctors) usually being responsible for bearing the costs of travel and accommodation. The in-person interviews would also require coordination of the hospital program directors (interviewers) to be physically present for both the interview and any tour of the premises which the candidates would attend. During the 2020-2021 job market cycle, however, due to the COVID-19 pandemic, most interviews moved to a virtual format. It stands to reason that such a change may have altered agents' interview costs. For instance, candidates/doctors would no longer have to pay for expensive air travel. For hospital program directors, the change in the costs may have been less pronounced, but still significant: For hospitals the primary cost of the interview is likely the opportunity cost of the interviewers, who may be high-paying doctors who could otherwise be performing procedures with high profit-margins for the hospital. Still, the ability to conduct virtual interviews likely eases any coordination issues the hospital may face, in particular when more than one hospital employee is involved in the interview process: Allowing program directors to conduct the interviews from the comfort of their personal offices may lower any "down time" otherwise caused by the interview process.

Having argued that the move to virtual interviews likely implied lower interview costs, both for hospitals and for doctors, the obvious question that arises is what effect such a change has on equilibrium welfare. Clearly, holding the interview allocations fixed, a decrease in interview costs make agents better off (as illustrated in the left panel of Figure 4). However, a reduction in (hospital) interview costs may also drive hospitals to increase their interview activity. By Theorem 4.1, this imposes a negative welfare externality on other hospitals. In a symmetric world, in which all hospitals initially conduct, say, k interviews, a decrease in hospital interview costs may drive all hospitals to increase their interview activity. In equilibrium, hospitals gain from the lower costs they incur, but incur a loss from the increased welfare externalities to which they are exposed. The net effect is ambiguous, and depends on the relative sizes of the gain (green area in the right panel of Figure 4) and the negative welfare externalities (red areas in the right panel of Figure 4) to which hospitals are exposed. In Example 5, we show how hospitals may indeed be strictly worse off as they move from an equilibrium with high interview costs and low interview activity, to an equilibrium with lower interview costs and higher interview activity. The insights from this example and the discussion above is summarized in the following observation:

Observation 5.1 *Equilibrium welfare need not be monotonic in agents' interview costs.*

Figure 4: **A reduction in hospitals' interview costs may decrease equilibrium welfare**
Consider a symmetric equilibrium. A reduction in hospital interview cost from c to c' has an unambiguous positive effect (illustrated by the green area in the left figure), as long as the decrease does not change the equilibrium strategy of any hospital. As costs further decrease from c' to c'', a hospital finds it optimal to increase its interview activity. Since this applies to all hospitals, every hospital is also exposed to the negative welfare externalities of the increase in its competitors interview activity (illustrated by the red areas in the figure to the right). The net effect is ambiguous.

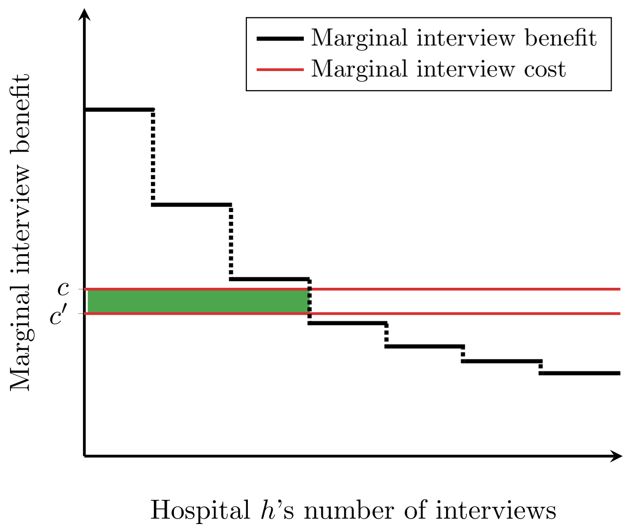

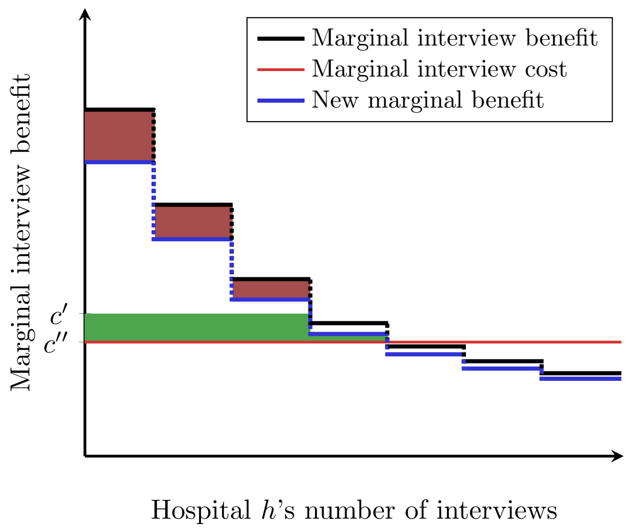

Example 5 *Consider the case with* $|\mathcal{H}| = |\mathcal{D}| = 3$*, with* $v_{hd}, v_{dh} \overset{iid}{\sim} U(0,1)$ *and* $\beta = 1$*. Let* $[j,1]_{-h}$ *denote the anonymous strategy profile in which all of hospital h's competitors conduct j interviews each with probability 1. One can show that* $q_h^{(1)}([1,1]_{-h}) = \frac{76}{108}$*,* $q_h^{(2)}([1,1]_{-h}) = \frac{26}{108}$*, and that* $q_h^{(1)}([2,1)]_{-h}) = \frac{69}{108}$*,* $q_h^{(2)}([2,1]_{-h}) = \frac{26}{108}$*. Based on this, it's an equilibrium for all hospitals to conduct one interview each when* $c_{\mathcal{H}} \in (128/648, 228/648)$*, two interviews each when* $c_{\mathcal{H}} \in (80/648, 121/648)$*, and three interviews each when* $c_{\mathcal{H}} < 80/648$*. Equilibrium expected utilities are*

$$\begin{aligned} V(1,[1,1]) - c_{\mathcal{H}} &= 228/648 - c_{\mathcal{H}} && \text{for } c_{\mathcal{H}} \in (128/648, 228/648) \\ V(2,[2,1]) - 2c_{\mathcal{H}} &= 328/648 - 2c_{\mathcal{H}} && \text{for } c_{\mathcal{H}} \in (80/648, 121/648) \\ V(3,[3,1]) - 3c_{\mathcal{H}} &= 408/648 - 3c_{\mathcal{H}} && \text{for } c_{\mathcal{H}} < 80/648 \end{aligned}$$

Consider a reduction in interview costs from $130/648$ *to* $120/648$*. In equilibrium, hospitals all increase their number of interviews from 1 to 2. Despite the decrease in interview costs, hospitals' equilibrium expected utilities decrease from* $98/648$ *to* $88/648$*, making all hospitals worse off. Since*

$\beta = 1$*, the third interview any hospital conducts imposes no externality on the other hospitals. Thus, when* $c_{\mathcal{H}} < 121/648$*, a reduction in interview costs always make hospitals better off.*

The discussion above focused on hospitals' perspective. However, the logic extends to doctors: Doctors are indeed better off if their interview costs decrease, as long as the (expected) number of interviews they conduct, as well as the expected match utilities they achieve are unchanged. Therefore, if doctors' interview costs only decrease, and hospitals' costs are changed, then the set of equilibria is unchanged, and in each eqquilibrium doctors are unambiguously better off. If both doctors' *and* hospitals' interview costs decrease, the effect is more complicated: If the decrease in hospital interview costs is such that hospitals, in equilibrium, conduct more interviews, the effect on doctors' welfare is ambiguous. From Theorem 4.1, doctors' expected match utilities increase. On the other hand, the total effect on doctors' expected interview costs is ambiguous. On the one hand, the per-interview costs $c_{\mathcal{D}}$ may decrease, but the overall costs may increase as long the increase in the expected number of interviews is sufficiently large. As a result, the increase in doctors' expected match utilities may be too low to offset the increase in their interview costs, similar to the logic of Corollary 4.5 and Proposition 4.6.

5.2 Varying the Number of Market Participants and Market Thickness

We next turn to the analysis of a change in the number of market participants. This is empirically relevant for the entry-level market for doctors in the US. Indeed, Figure 5, taken from *nrmp.org*, documents the number of total applicants and total post-graduate year 1 (PGY-1) positions at US hospitals participating in the Main Residency Natch organized by the NRMP, over the years 1952-2020. As is immediately clear, the number of positions has more than tripled over the almost 70 year period, while the number of applicants has increased by a factor of more than 8. It is worth noting, however, that in recent decades, a non-negligible fraction of applicants represent candidates who received their medical education outside of the US. Few among these candidates are successful at obtaining a match, and foreign applicants typically receive fewer interview offers than their US counterparts.

To get a sense of what the effect of increasing the number of market participation, consider the following: On the one hand, increasing the number of doctors allows hospitals to interview more candidates, which may yield higher expected match utilities. On the other hand, increasing the number of hospitals exposes hospitals to increased negative welfare externalities, potentially making

Figure 5: **Evolution of the number of applicants and positions participating in the NRMP residency match, 1952-2020.** Source: nrmp.org

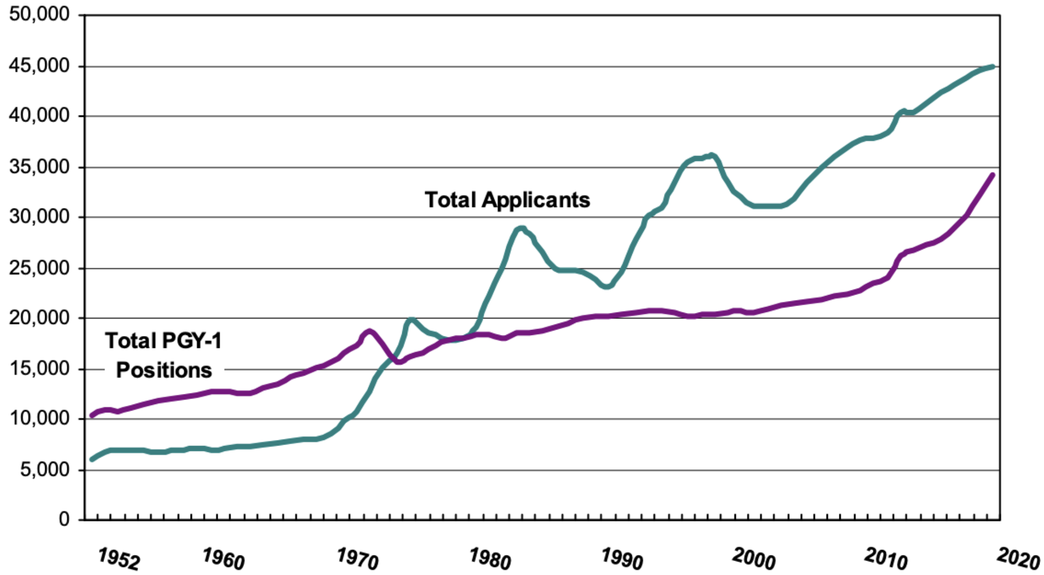

it harder for hospitals to match with their top choices. Which effect dominates in equilibrium is ambiguous, and will depend on hospitals' interview costs. The following simple example illustrates the trade-off that are involved:

Example 6 *Suppose for all $h \in \mathcal{H}$, $d \in \mathcal{D}$, we have $v_{hd} \sim U(0,1)$, with $\beta = 1$. Consider first the case in which $|\mathcal{H}| = |\mathcal{D}| = 1$. If the hospital interviews the doctor, the hospital's expected utility will be $1/2 - c_{\mathcal{H}}$, which is strictly positive as long as $c_{\mathcal{H}} < 1/2$.*

Consider now the case in which $|\mathcal{H}| = |\mathcal{D}| = 2$. Suppose both hospitals conduct one interview each, at random. In this case, the two hospitals will overlap with probability $1/2$, in which case any hospital will "lose" the doctor in question with probability $1/2$, leading to an expected utility of $\frac{3}{4}\ 1/2 - c_{\mathcal{H}}$. If a hospital chooses to interview both doctors, it will not be matched with its most preferred doctor with probability $\frac{3}{4}$, hence the expected utility is $\frac{3}{4}2/3 + \frac{1}{4}1/3 - 2c_{\mathcal{H}}$. Conducting both interviews will be an equilibrium as long as $c_{\mathcal{H}} < 5/24$. In either case, equilibrium welfare under $|\mathcal{H}| = |\mathcal{D}| = 2$ is lower than under $|\mathcal{H}| = |\mathcal{D}| = 1$ as long as $c_{\mathcal{H}} \geq 1/12$.

The purpose of our analysis will be to understand how the number of market participants affects both equilibrium welfare and hospitals strategies, and to quantify the extent of the welfare externalities and strategic externalities to which hospitals are exposed in equilibrium. Our focus will be on proportional changes in the number of hospitals and doctors, usually referred to changing *market thickness*. To facilitate our analysis, we will make additional simplifying assumptions, which are detailed below. First, throughout this section, we will consider the case in which agents have no pre-interview information, in particular $v_{hd} \sim F_{\mathcal{H}}$ for all (h, d). Second, our focus will be on symmetric equilibria in anonymous strategies, i.e. equilibria in which all hosptials play the same strategy. There are multiple reasons for this shift in focus: Restricting attention to environments with no pre-interview information, while also only considering symmetric equilibria, allows us to abstract from the complicating cases in which some hospital types may benefit from increased market thickness, while other hospital types may be hurt by the same change. While these cases are certainly both interesting and important, they lie outside of the scope of this paper. Finally, even when we restrict attention to symmetric equilibria, the set of equilibria is potentially large, and it is not immediately clear how to compare the set of equilibria as we change the number of market participants. To make such comparisons more straightforward and coherent, our focus will be on the symmetric equilibrium in anonymous strategies that maximizes hospital welfare, which we show below, indeed exists.

Recall that any equilibrium in anonymous strategies can be characterized by the number of interviews that each hospital is conducting. By Proposition 3.4, all equilibria are either pure, in which all hospitals conduct exactly k interviews, for some $k \leq |\mathcal{D}|$, or the number of interviews each hospital conducts is mix between two consecutive numbers $k-1$ and k, with $k \leq |\mathcal{D}|$. By Corollary 4.2 hospital welfare is decreasing in the equilibrium number of interviews hospitals conduct. Hence, the hospital welfare-maximizing symmetric equilibrium in anonymous strategies is the one in which hospitals conduct the fewest interviews possible.

To make further progress, we will introduce additional notation: First, due to the absence of pre-interview information, assuming every $h' \neq h$ plays an anonymous strategy, then h is indifferent between interviewing any subsets $S, S' \subseteq \mathcal{D}$ such that $|S| = |S'|$. With some abuse of notation, therefore, for any anonymous strategy profile σ_{-h}, we let $V(k, \sigma_{-h}) := V(S, \sigma_{-h})$ for any S such that $|S| = k$ (where we also omit the dependence on θ). Second, for any $1 \leq k \leq |\mathcal{D}|$, denote by $[k, p]$ the mixed strategy which with probability p randomly and uniformly selects k doctors to interview, and with probability $1-p$ randomly and uniformly selects $k-1$ doctors to interview. In particular, we have $[k+1, 0] = [k, 1]$ and let $[1, 0]$ be the strategy that involves no interviews. Note that $[k, p]$ is an anonymous strategy for every (k, p). Denote by $([k, p])_{h \in \mathcal{H}}$ the strategy profile in which every hospital plays $[k, p]$. Third, using the above, define $g_j([k, p]) := V\big(j, ([k, p])_{-h}\big) - V\big(j-1, ([k, p])_{-h}\big)$. Using this notation, $([k, p])_{h \in \mathcal{H}}$ is an equilibrium if every hospital h is indifferent between conducting

k and $k-1$ interviews, given that all other hospitals are mixing between k and $k-1$ with probability p and $1-p$, respectively. This implies that the marginal benefit of conducting the k-th interview must equal the marginal cost: $g_k([k,p]) = c_{\mathcal{H}}$. Note that in the case of $p \in \{0,1\}$, either $([k-1,1])_{h\in\mathcal{H}}$ or $([k,1])_{h\in\mathcal{H}}$ is a pure strategy equilibrium. It turns out that, as a function of p, $g_k([k,p])$ has certain convenient properties:

Lemma 5.2 *For every $k \leq |\mathcal{D}|$, $g_k([k,p])$ is a polynomial function in p, and hence (uniformly) continuous on* $[0,1]$.

Proof: For any $p \in [0,1]$, the probability that exactly n of h's competitors will conduct k interviews is $\binom{|\mathcal{H}|-1}{n} p^n (1-p)^{|\mathcal{H}|-1-n}$. Since every firm randomly selects the doctors to interview, the identity of the firms that conduct k interviews does not influence h's expected match utility. For $n \leq |\mathcal{H}|-1$, denote by $(k,n)_{-h}$ the strategy in which n of h's competitors conduct k interviews, while the others conduct $k-1$ interviews. With a slight abuse of notation, we can now write

$$V\big(j, ([k,p])_{-h}\big) = \sum_{n \leq |\mathcal{H}|-1} \binom{|\mathcal{H}|-1}{n} p^n (1-p)^{|\mathcal{H}|-1-n} V\big(j, (k,n)_{-h}\big)$$

which is a polynomial in p, and the result follows. ■

This result has immediate consequences for the existence of equilibrium which we will rely on in the following analysis:

Proposition 5.3 *A hospital welfare-maximizing symmetric equilibrium in anonymous strategies always exists.*

Proof: To see that an equilibrium always exists, it's sufficient for some $k \leq |\mathcal{D}|$, that

$$\min\big\{g_k([k,0]), g_k([k,1])\big\} \leq c_{\mathcal{H}} \leq \max\big\{g_k([k,0]), g_k([k,1])\big\}$$

Indeed, since the function $g_k([k,p])$ is continuous in p, thus if the above two inequalities hold, by the Intermediate Value Theorem there exists a $p \in [0,1]$ such that $g_k([k,p]) = c_{\mathcal{H}}$. If the above inequalities do not hold for any k, then it must be that either

$$\max\big\{g_k([k,0]), g_k([k,1])\big\} < c_{\mathcal{H}} \quad \text{for all } k\text{, or}$$
$$\min\big\{g_k([k,0]), g_k([k,1])\big\} > c_{\mathcal{H}} \quad \text{for all } k$$

In the former case, it is an equilibrium for all hospitals to conduct 0 interviews each. In the latter case, it is an equilibrium for all hospitals to conduct $|\mathcal{D}|$ interviews.

Since for each k, the function $g_k([k,p])$ can be written as a finite polynomial, it also follows that the equality $g_k([k,p]) = c_{\mathcal{H}}$ holds for finitely many p. By Corollary 4.2, the welfare-maximizing equilibrium is the one in which hospitals conduct the lowest number of interviews. Since there are finitely many equilibria, there must be an equilibrium that maximizes hospital welfare. ■

Having established the existence of a welfare-maximizing symmetric equilibrium, we next seek to construct a (theoretical) algorithm that can find the equilibrium in question. The idea behind the algorithm, which we describe in detail below, is to increase the number of interviews hospitals conduct, at most by an increment of 1, until we arrive at a situation in which $g_k \leq c_{\mathcal{H}}$. Key to this procedure will be our ability to determine whether we can say with certainty that $g_k([k,p]) > c_{\mathcal{H}}$ for all p in some interval of mixing probabilities $[p_0, p_1]$. The following lemma will be useful:

Lemma 5.4 *For any $k \leq |\mathcal{D}|$ and any $0 \leq p_0 < p_1 \leq 1$ we have $g_k([k,p]) > c_{\mathcal{H}}$ for all $p \in [p_0, p_1]$ as long as*

$$V\big(k, ([k,p_1])_{-h}\big) - V\big(k-1, ([k,p_0])_{-h}\big) > c_{\mathcal{H}}$$

Proof: Using that $V(j, ([k,\hat{p}])_{-h})$ is decreasing in $\hat{p}$, for any $p \in (p_0, p_1)$ we get

$$g_k([k,p]) := V\big(k, ([k,p])_{-h}\big) - V\big(k-1, ([k,p])_{-h}\big) \geq V\big(k, ([k,p_1])_{-h}\big) - V\big(k-1, ([k,p_0])_{-h}\big)$$

■

We next construct an algorithm that allows us to find the hospital welfare-maximizing symmetric equilibrium in anonymous strategies. The idea behind the algorithm is, conditional on reaching k, to check whether $([k,p_0])_{h\in\mathcal{H}}$ is an equilibrium for $p_0 = 0$, by evaluating if $\big(k, ([k,0])_{-h}\big) - V\big(k-1, ([k,0])_{-h}\big) \leq c_{\mathcal{H}}$. If it's not, then algorithm evaluates the condition of Lemma 5.4 for $p_0 = 0$ and $p_1 = 1$. If the condition is satisfied (with $p_1 = 1$) then the algorithm increases k by 1, and starts over with the new value of k. If the condition of the lemma is not satisfied, then the algorithm decreases p_1 to $(p_0 + p_1)/2$, and checks the condition of Lemma 5.4 again with the new values of p_0 and p_1. If the condition of the lemma is satisfied for some $p_1 < 0$, then p_0 is increased to the current value of p_1, and p_1 is increased to its next lowest value. If the condition of the lemma is not satisfied, then if $p_1 - p_0 < \delta$, the algorithm ends. Otherwise, p_0 is held at its current level, while p_1 is decreased to $(p_1 + p_0)/2$, before the algorithm again evaluates the condition of Lemma 5.4.

An algorithm for approximating the welfare-maximizing symmetric equilibrium

Fix $\delta > 0$ and initialize $k = STOP = 0$. While $STOP = 0$

Step 0: $k = k + 1$.

(a) If $k > |\mathcal{D}|$. Stop the algorithm. $[|\mathcal{D}|, 1]$ is an equilibrium.

(b) Else: Initialize $n = 0$, and $p_0^n = 0$, $p_1^n = 1$. Proceed to Step 1.

Step 1: If $n = 0$, and $g_k([k, p_0^0]) \leq 0$: End the algorithm, $[k, 0] = [k-1, 1]$ is an equilibrium. Else:

(a) If $V\big(k, ([k, p_1^n])_{-h}\big) - V\big(k-1, ([k, p_0^n])_{-h}\big) > c_{\mathcal{H}}$:

i. If $p_1^n = 1$: Proceed to Step 0.

ii. Else: $n = n + 1$. Set $p_0^n = p_1^{n-1}$, $p_1^n = \min\limits_{0 \leq i \leq n-2}(p_1^i : p_1^i > p_0^n)$. Proceed to Step 1.

(b) Else if $p_1^n - p_0^n < \delta$: End the algorithm; $STOP = 1$.

(c) Else: $n = n + 1$. Set $p_0^n = p_0^{n-1}$, $p_1^n = (p_1^{n-1} - p_0^n)/2$. Proceed to Step 1.

It turns out that the algorithm indeed converges to the welfare-maximizing symmetric equilibrium as $\delta \to 0$, as the following result show, the proof of which is in the appendix. Results from numerical simulations of the hospital welfare-maximizing symmetric equilibrium are provided in Section 5.2.1.

Lemma 5.5 *For every $\delta > 0$ the algorithm finishes in a finite number of steps. Moreover, for every $\epsilon > 0$ there exists $\delta > 0$ such that the algorithm ends with an evaluation of hospital welfare within ϵ of the welfare achieved in the hospital welfare-maximizing symmetric equilibrium, and within δ of the corresponding mixing probability p.*

5.2.1 Simulation Results of the Hospital Welfare-Maximizing Symmetric Equilibrium

In this section, we present results from numerical simulations of the hospital welfare-maximizing symmetric equilibrium in anonymous strategies, based on the algorithm described in the previous section. Specifically, the algorithm described above is based on knowledge of the values of $V(k, ([k, p])_{-h})$ and $V(k-1, (k, p])_{-h})$ for different values of p. Since these values are unknown, we estimate them using simulations. Since everything is symmetric, it suffices to estimate the values for one hospital, which we will refer to as hospital h_1. Specifically, for every (k, p), we generate N independent draws of the vectors v_h and v_d for each hospital and doctor from the distributions

$F_{\mathcal{H}}$ and $F_{\mathcal{D}}$. For each hospital $h \neq h_1$, we next generate N (independent) vectors i_h of interview indicators which, with probability p randomly assign a 1 for k different doctors, and with probability $k-1$, randomly assign a 1 for $k-1$ doctors, and in both cases assigns a 0 for all other doctors. The hospital's preferences are then calculated as the element-wise product of the vectors v_h and i_h. To estimate $V(k, ([k, p])_{-h})$ we generate N independent draws of interview indicators i_{h_1} for hospital h_1, all indicating exactly k (random) interviews, and generate corresponding interview indicators of length $k-1$ to estimate $V(k-1, ([k, p])_{-h})$. For each strategy $[k, 1]$ and $[k, 0]$ for hospital h_1, we can now calculate the match outcome and corresponding match utility for h_1 in all N market instances. Taking the average over all N instances provides an estimate of h_1's expected match utility, holding the strategies of all other hospitals constant, which allows us to estimate the benefit to h_1 of the k-th interview. Each of the N market instances also allows us to estimate the average expected match utility that doctors achieve in the hospital welfare-maximizing symmetric equilibrium. Throughout this sections, all simulations were based on match utilities drawn from the uniform $U(0, 1)$ distribution, with all doctors acceptable ($\beta = 1$), and with expected utilities estimated from $N = 250,000$ market instances. The precision of the equilibrium-finding algorithm was set to $\delta = 10^{-5}$. Finally, we restricted attention to the case of balanced markets ($|\mathcal{H}| = |\mathcal{D}|$), and simulated equilibrium outcomes for the cases of $|\mathcal{H}| \in \{2, 3, 5, 10, 20, 30, 50, 75, 100\}$. We also considered different levels of interview costs; $c_{\mathcal{H}} \in \{0, 0.01, 0.05\}$.

Figure 6 shows the results from the numerical simulations of the hospital welfare-maximizing symmetric equilibrium for different hospital interview costs, and for different levels of market thickness. We included results for the case with $c_{\mathcal{H}} = 0$, in which all hospitals interview all doctors. The evolution of hospital and doctor welfare in this case is known theoretically, and the results were included as a reference only. Note that with $c_{\mathcal{H}} = 0$ hospital welfare is monotonically increasing in market thickness. This stands in stark contrast to the two other curves we included (the numbers below the curves indicate the equilibrium number of interviews hospitals conduct):

Panel (a) shows hospitals expected match utilities, net of interview costs. For $c_{\mathcal{H}} = 0.01$, hospitals all optimally choose to interview all the doctors in the market as long as $|\mathcal{D}| \leq 10$, and hospitals' equilibrium welfare is increasing to this point. From $|\mathcal{H}| = |\mathcal{D}| = 20$ onward, however, hospitals no longer find it optimal to interview all the doctors in the market. Moreover, hospital equilibrium welfare is lower at this level of market thickness than under $|\mathcal{H}| = |\mathcal{D}| = 10$, and continues to fall as the market further thickens. In particular, equilibrium welfare under $|\mathcal{H}| = |\mathcal{D}| = 100$ is barely higher than under $|\mathcal{H}| = |\mathcal{D}| = 2$. As a result, the gap between the welfare curve under $c_{\mathcal{H}} = 0.01$ and $c_{\mathcal{H}} = 0$ is widening as the market thickens. Within the range of market sizes we're considering, the results indicate that the equilibrium number of interviews (with $c_{\mathcal{H}} = 0.01$) increases as the market thickens, albeit very slowly: The equilibrium number of interviews are 17 when $|\mathcal{H}| = |\mathcal{D}| = 20$ and agents mix between 19 and 20 interviews as $|\mathcal{H}| = |\mathcal{D}| = 100$, meaning that the fraction of doctors that each hospital interviews decreased from 85% to less than 20%. One possible explanation for the non-monotonic nature of hospitals' expected utilities as a function of market thickness could be the increase in interview costs that they incur. Panel (b), which displays hospitals' expected match utilities shows that this is only part of the story: Expected match utility under $|\mathcal{H}| = |\mathcal{D}| = 20$ is indeed higher than under $|\mathcal{H}| = |\mathcal{D}| = 10$, but is monotonically decreasing beyond $|\mathcal{H}| = |\mathcal{D}| = 20$, with expected match utilities being similar under $|\mathcal{H}| = |\mathcal{D}| = 100$ and $|\mathcal{H}| = |\mathcal{D}| = 10$. Therefore, not only does market thickness provide incentives for hospitals to increase their interview activity and thus incur additional interview costs, it also exposes them to increased welfare externalities, lowering their expected match utilities.

The simulation results for the higher interview costs ($c_{\mathcal{H}} = 0.05$) indicate a similar tendency, although the contrast to the case with zero interview costs is even starker. Equilibrium welfare (Panel (a)) indicates that welfare is barely higher under $|\mathcal{H}| = |\mathcal{D}| = 3$ than under $|\mathcal{H}| = |\mathcal{D}| = 2$, and is monotonically decreasing thereafter. Since interview costs are relatively large, hospitals do not find it in their best interest to increase their interview activity by a lot; within the range of market sizes under consideration, hospitals at most mix between 5 and 6 interviews. As Panel (b) indicates, hospital expected match utility is also monotonically decreasing beyond $|\mathcal{H}| = |\mathcal{D}| = 3$, with an ever widening gap between the case of zero interview costs.

Figure 7 shows the expected match utilities achieve in the hospital welfare-maximizing symmetric equilibrium for different interview costs and market sizes. As for hospitals, doctors' expected match utilities are monotonically increasing for $c_{\mathcal{H}} = 0$, although at a different rate than for hospitals. When hospitals' interview costs are $c_{\mathcal{H}} = 0.01$, the simulation results suggest doctors' expected match utilities are still monotonically increasing, at least within the range of market sizes we are considering. Moreover, the gap between the case with low interview costs ($c_{\mathcal{H}} = 0.01$) and zero interview costs is much less pronounced than for hospitals. For the higher hospital interview costs ($c_{\mathcal{H}} = 0.05$), however, doctors' equilibrium expected match utilities evolve very differently. Beyond $|\mathcal{H}| = |\mathcal{D}| = 10$, doctors' expected match utilities seem to have reached a plateau, and very little variation is discernible. That said, the numerical estimate of the expected match utilities under $|\mathcal{H}| = |\mathcal{D}| = 100$ is indeed strictly lower than under $|\mathcal{H}| = |\mathcal{D}| = 30$.

Figure 6: **Hospital equilibrium net utility, match utility and equilibrium strategies, for different market sizes and interview costs**

Notes: Simulation of the hospital welfare-maximizing symmetric equilibrium, based on balanced markets ($|\mathcal{H}| = |\mathcal{D}|$), with $\beta = 1$ and $v_{hd}, v_{dh} \overset{iid}{\sim} U(0,1)$, and $\delta = 10^{-5}$. Expected match utilities estimated from 250,000 realization of market instances, given the (mixed) strategies. The numbers in square brackets represent the equilibrium number of interviews each hospitals conducts.

(a) Equilibrium expected match utility net of interview costs

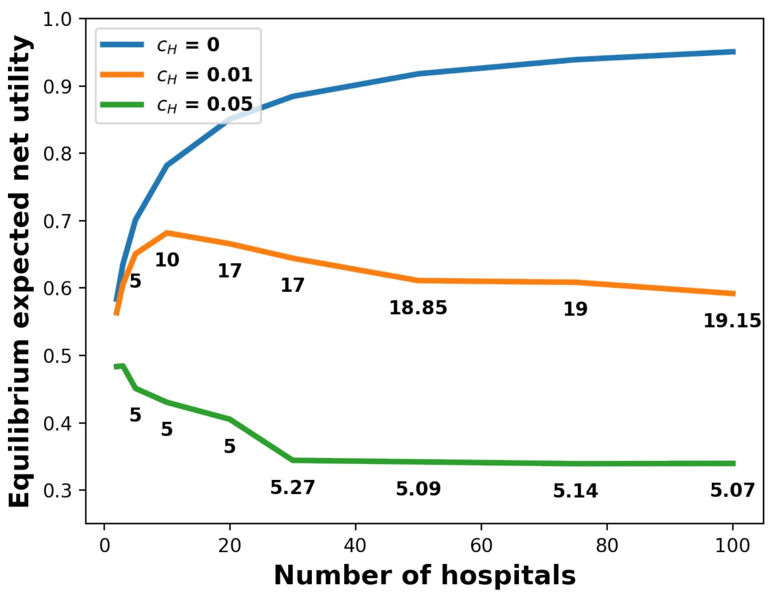

(b) Equilibrium expected match utility

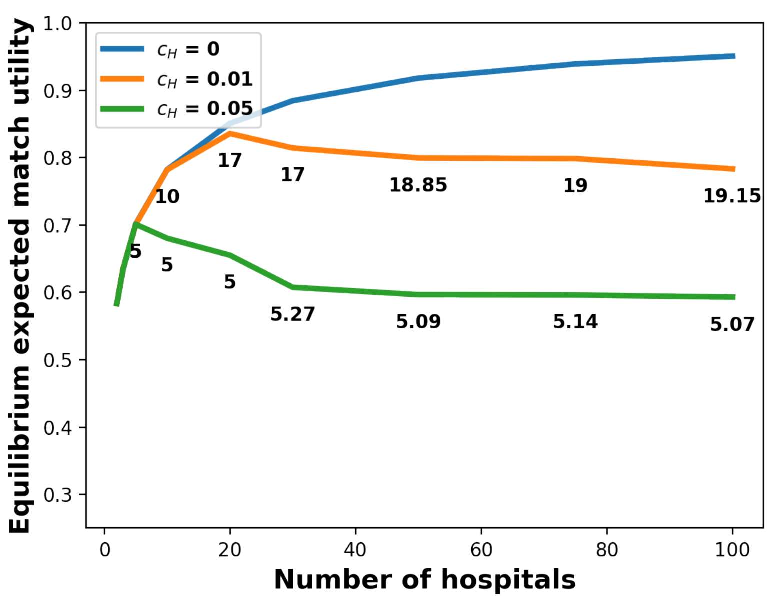

Figure 7: **Doctor equilibrium match utility, for different market sizes and hospital interview costs**

Notes: Simulation of the hospital welfare-maximizing symmetric equilibrium, based on balanced markets ($|\mathcal{H}| = |\mathcal{D}|$), with $\beta = 1$ and $v_{hd}, v_{dh} \overset{iid}{\sim} U(0,1)$, and $\delta = 10^{-5}$. Expected match utilities estimated from 250,000 realization of market instances, given hospitals' (mixed) strategies.

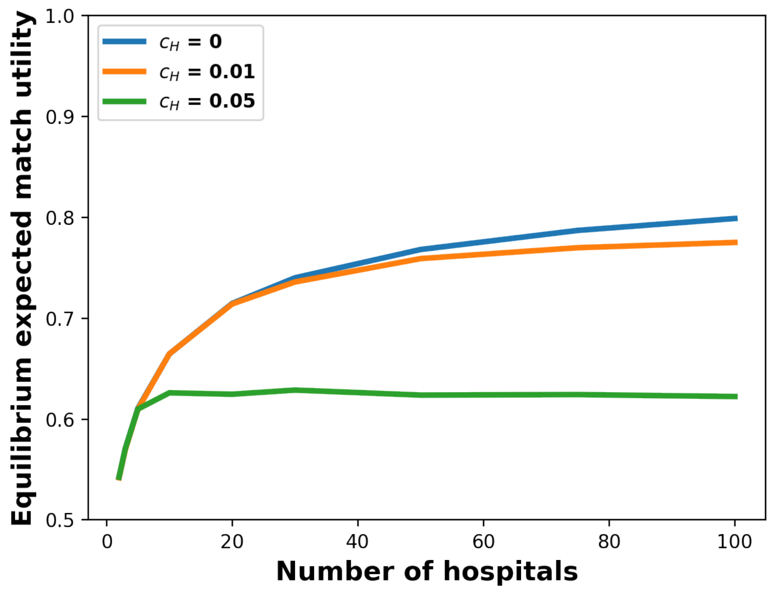

The two previous figures provides a convenient comparison to the "idealized" scenario in which there are no interview costs, and all hospitals interview all doctors. However, the simulation results themselves say little about the exact extent of the welfare externalities and strategic externalities to which hospitals are exposed. To investigate this in more detail, for every level of market thickness and interview costs, we consider the case in which only one hospital conducts interviews. That is, we consider the optimal number of interviews a hospital would choose, as well as the resulting welfare the hospital would achieve, in the case in which none of its competitors conducted any interviews. In this case, a hospital is always matched to its highest ranked doctor according to its post-interview preferences. Using the fact that match utilities are drawn from the uniform distribution, if the hospital conducts k interviews, it will achieve an expected match utility of $\frac{k}{k+1}$. The increase from $k-1$ to k interviews is therefore $\frac{1}{k(k+1)}$. Formally, the hospital's optimal strategy is therefore given by the set of k's satisfying

$$g_k([1,0]) = \frac{1}{k(k+1)} \geq c_{\mathcal{H}} \geq \frac{1}{(k+1)(k+2)}$$

Figure 8: **Hospital optimal strategy and net utility under when no other hospital conducts interviews, relative to hospital equilibrium match utility and equilibrium strategies, for different market sizes and interview costs**

Notes: Simulation of the welfare-maximizing symmetric equilibrium, based on balanced markets ($|\mathcal{H}| = |\mathcal{D}|$), with $\beta = 1$ and $v_{hd} \sim U(0,1)$, and $\delta = 10^{-5}$. Match utilities estimated from 250,000 independent market simulations. The numbers in square brackets represent the equilibrium strategies under different costs.

(a) Expected net utility and strategies with $c_{\mathcal{H}} = 0.01$

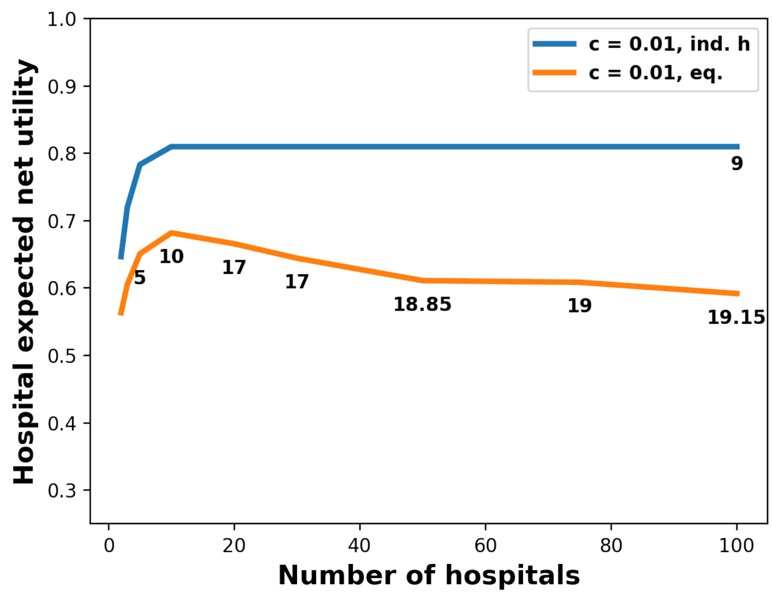

(b) Expected net utility and strategies with $c_{\mathcal{H}} = 0.01$

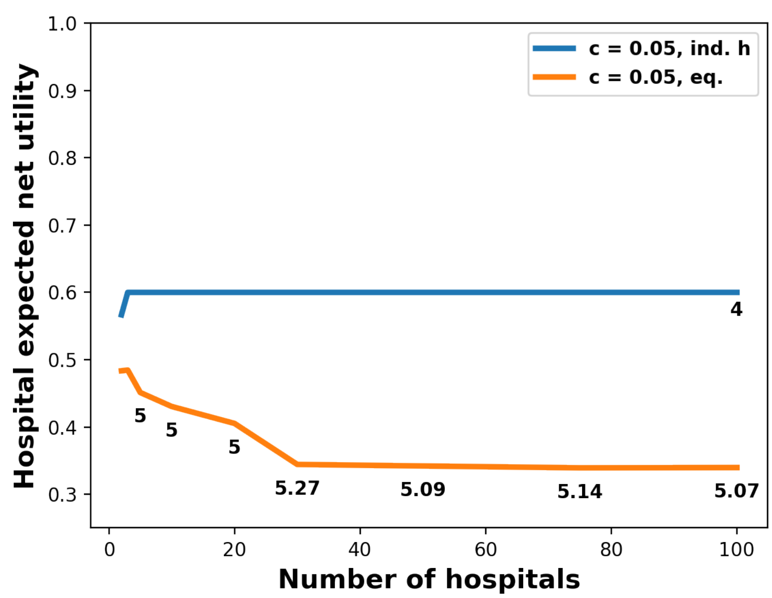

In the case of $c_{\mathcal{H}} = 0.01$, the unique optimal number of interviews for the single hospital is $k = \min\{9, |\mathcal{D}|\}$, i.e. the hospital interviews all doctors as long as there are no more than 9, but stops interviewing at 9 whenever the number exceeds this level. The achieved expected utility net of interview costs is 0.81 (for $|\mathcal{D}| \geq 9$). The comparison to the hospital welfare-maximizing symmetric equilibrium is illustrated in Panel (a) of Figure 8. Unsurprisingly, there's a substantial gap between the welfare the hospital obtains as the sole interviewing hospital and the welfare obtained in the symmetric equilibrium. For $|\mathcal{H}| = |\mathcal{D}| \geq 50$, in the symmetric equilibrium the hospital is exposed to net welfare externalities that amount to around 25% of the utility achieved when the hospital is alone in the market. The strategic externalities are also quantifiably large: Compared to the case in which the hospital is alone in the market, the hospital more than doubles its number of interviews in the hospital welfare-maximizing symmetric equilibrium.

When $c_{\mathcal{H}} = 0.05$ and the hospital is alone in the market, the hospital is indifferent between conducting 4 and 5 interviews, and the maximum expected utility is given by 0.6. Panel (b) of Figure 8 plots the equilibrium utility relative to the utility achieved under the hospital welfare-maximizing equilibrium. For $|\mathcal{H}| = |\mathcal{D}| \geq 30$, the net welfare externalities are around 40% of the utility the hospital achieves when it's alone, which is higher than in the case of $c_{\mathcal{H}} = 0.01$. However, the strategic externalities are less substantial with $c_{\mathcal{H}} = 0.05$, as the optimal number of interviews increase from an indifference between 4 and 5 interviews to an indifference between 5 and 6 interviews.

Before we proceed, we point out that Lee and Schwarz (2017) also did include simulation results for equilibrium strategies for different market sizes. However, their focus were on fixing hospitals' strategies, and then derive conditions on interview costs under which the strategy profile in question constituted an equilibrium. This makes direct comparisons of a change in market size, holding interview costs constant, particularly difficult. Our approach, which is based on a particular equilibrium selection criterion, makes the interpretability of the comparative statics much more transparent, and is one of the major advantages of our analysis: We can independently change hospitals' interview costs and market size, and explore the resulting equilibrium strategy and equilibrium welfare for each parameter choice.

5.2.2 Incentives to Match Outside the Centralized Clearinghouse

The simulation results in the previous sections illustrate how hospitals and doctors, when the number of interviews they conduct is limited by agents' interview costs, may not be able to reap the full benefits of market thickness. Indeed, the simulation results suggest there is a widening gap between the welfare that would result in the case of zero interview costs, and the equilibrium welfare that results under strictly positive interview costs. More precisely, due to the presence of the costly interviews, agents on both sides of the market may be worse off as the market thickens. This raises questions about the usefulness of the centralized match in thick markets. At least since Roth (1991), a widely held view has been that the use of a stable matching mechanism is crucial to prevent the centralized market from disintegrating, which among other things, could lead the market to unravel. Roth (2008) further argue that market thickness is a crucial component for the success of the centralized clearinghouse: For agents to find it useful to participate in a centralized match, the clearinghouse must attract a sufficiently large proportion of the market participants. Numerous authors have pointed to the benefits of thick markets.[7] However, Roth (2008) also point out that the clearinghouse must "overcome the *congestion* that thickness can bring". Early work on congestion considered the cases in which congestion resulted from the bottlenecks caused by the available communication technology (Roth and Xing (1997)). In the case of interviews, the primary issue may not be a technological bottleneck, but rather the costs incurred by agents as they acquire information about their preferences.[8]

While formally extending the model to account for agents' participation decision in the centralized match is beyond the scope of this paper, a brief discussion is still warranted. In many cases, hospitals and doctors may have access to better information about their match qualities with certain potential market participants than with others. For instance, a doctor who worked as an intern at a hospital during their medical education may have a decent idea about what life as a resident at this particular hospital may be like. Similarly, the doctor may have interacted with the residency program director at the hospital, and the program director may have formed an opinion about the medical student's suitability for their program. If both the doctor and the residency program director expect that they will have to go through a particularly costly interview process in order to successfully obtain a good match outcome through the centralized clearinghouse, both the medical student and the residency program director may find it in their best interest to agree to match and simply circumvent the centralized match. In such a scenario, the presence of the

[7] See, for instance, Allen and Gale (1994), Ngai and Tenreyro (2014), Gan and Li (2016), Akbarpour et al. (2020), and Loertscher and Muir (2021)

[8] There may be other technological constraints that are relevant for interviews: For instance, the practical considerations around the decentralized scheduling of interviews may be an additional issue which we ignore in this paper.

costly interviews may lower participation in the centralized match, which potentially could cause the market to unravel.

The above discussion opens the door to explore the role of market design interventions, which we will discuss in more detail in the next section.

6 Lessons for Market Design Interventions

In recent years, alongside various descriptions of the issues related to the interview process in the market for residency and fellowship positions, many authors have also proposed market design interventions to address some of these issues. Many of these interventions have been suggested in medical journals, and relatively little work on the topic has been published in the economics literature. Some recent proposals in the medical literature include the geographic fragmentation of the residency market, through a "geographically linked consortium match" (Wong (2016)), such that hospitals in a given geographical location all end up interviewing the same candidates; the introduction of a preference signalling mechanism, similar to the one used by the American Economic Association for the job market for economists (Talcott and Evans (2021)); the introduction of an early decision option, similar to the early decision used by many colleges for in their admissions process (Monir (2020)); the introduction of a limit on the number of applications doctors can send (Burbano et al. (2019)); and the coordination of participants' interview activity through a centralized mechanism that imposes restrictions on the number of interviews participants are allowed to conduct (Melcher et al. (2018)). Some of these proposals draw on insights developed in the economics literature, either explicitly or inadvertently. For instance, Burbano et al. (2019) cite the work on application restrictions by Arnosti et al. (2021) as a justification for the usefulness of their proposal. Similarly, fragmenting the market such that all hospitals in a given geographic area interview the same set of applicants is reminiscent of the idea of *interview overlap* introduced by Lee and Schwarz (2017).

It turns out that our model of matching with interview provide new insights into some of the market design interventions mentioned above. In this section we will consider two such interventions: The introduction of either one- or two-sided restrictions on market participants' interview activity; and the introduction of a limit on the number of applications that doctors can send.

6.1 One- vs Two-sided Restrictions on Interview Activity

As shown in Section 4, one source of equilibrium inefficiency is the inability of hospitals to coordinate on the most efficient equilibrium: Indeed, it would be in the interest of all hospitals to coordinate on an equilibrium where each hospital conducts as few interviews as possible. Moreover, even when hospitals are able to coordinate on the most efficient equilibrium, the equilibrium may be inefficient, and both hospitals and doctors could potentially benefit from reduced interview activity by the hospitals. One way to achieve more efficient outcomes would be to impose restrictions on the number of interviews hospitals conduct. As the following example illustrates, however, there may be only a limited benefit to imposing such one-sided restricitons on interview activity:

Example 7 *Consider again the market in Example 5 with* $|\mathcal{H}| = |\mathcal{D}| = 3$*, with* $F = U(0,1)$ *and* $\beta = 1$*, in which we found symmetric pure-strategy equilibria involving either 1, 2, or 3 interviews per hospital, depending on interview costs. The equilibrium expected utilities were given by*

$$\begin{aligned} V(1,[1,1]) - c_{\mathcal{H}} &= 228/648 - c_{\mathcal{H}} && \textit{for } c_{\mathcal{H}} \in (128/648, 228/648) \\ V(2,[2,1]) - 2c_{\mathcal{H}} &= 328/648 - 2c_{\mathcal{H}} && \textit{for } c_{\mathcal{H}} \in (80/648, 121/648) \\ V(3,[3,1]) - 3c_{\mathcal{H}} &= 408/648 - 3c_{\mathcal{H}} && \textit{for } c_{\mathcal{H}} < 80/648 \end{aligned}$$

Since $\beta = 1$*, the third interview any hospital conducts imposes no externality on the other hospitals. Thus, from the hospitals' perspective, it is never in their interest to limit their number of interviews to 2. Consider a restriction on interview activity that limit hospitals to conduct no more than 1 interview each. This one-sided restriction only improves on the pure-strategy equilibrium whenever* $c_{\mathcal{H}} \in (100/648, 121/648)$*. When interview costs are such that in equilibrium hospitals conduct 3 interviews each, no one-sided restriction on interview activity improves hospital welfare.*

One of the issues highlighted in the previous example is that the one-sided restriction on interview activity does not allow hospitals to better coordinate on the limited number of interviews they conduct: Indeed, when hospitals are restricted to one interview each, it is still possible that all hospitals end up interviewing the same doctor, which would be inefficient. Such an outcome would imply that one doctor conducts three interviews, while the others conduct zero. Intuitively, to reap the benefits of reduced interview activity, one needs to limit the number of interviews on both sides of the market, as the following example shows:

Example 8 *Consider again the market in Example 7. Suppose a two-sided interview restriction is imposed, with an interview capacity of 1 interview per agent, such that each hospital ends up interviewing a unique doctor. From each hospital's perspective, this is equivalent to the case in which* $|\mathcal{H}| = |\mathcal{D}| = 1$, *with hospital welfare at* $1/2 - c_{\mathcal{H}}$. *Such a restriction always improves hospital welfare relative to the pure-strategy equilibria involving 1 or 2 interviews, and is preferred by hospitals to the equilibrium with 3 interviews each as long as* $c_{\mathcal{H}} > 42/648$. *Furthermore, it always improves upon the one-sided interview restriction in which hospitals are limited to one interview each.*

Having reached the conclusions that the benefits from a reduction in interview activity likely are larger if the reduction is coordinated on both sides of the market, two questions immediately arise: (1) what is the "best" interview allocation on which market participants should strive to coordinate? and (2) what are mechanisms through which such interview allocations can be implemented? Lee and Schwarz (2017) provide partial answers to both questions: First, an interview allocation that provides full *overlap* minimizes the number of unmatched agents.[9] Second, under no pre-interview information, full overlap interview allocations can be sustained in equilibrium for certain values of interview costs.[10] Hence, in principle, no complicated mechanism design is needed to implement the unemployment-minimizing interview allocation.

While Lee and Schwarz (2017) provide invaluable insights into how to address the mis-coordination of interview activity, there are several issues with their proposed solution: First, their theory is only explored under restrictive assumptions involving balanced markets in which agents have no pre-interview information about their match utilities. Second, even under the assumption of balanced markets and no pre-interview information, their results say nothing about what avenues to pursue when interview costs are such that the unemployment-minimizing interview allocation does not form an equilibrium. To illustrate this issue, consider an extension of Example 8 above:

Example 9 *Consider again the market in Examples 7 and 8, and consider the interview allocation in which each hospital and doctor conducts only one interview each. Under what values of* $c_{\mathcal{H}}$ *does the allocation form an equilibrium (in non-anonymous strategies)? If, say,* h_1 *decides to add the doctor interviewed by* h_2 *to its interview list, then* h_1 *would match with this doctor* 25% *of the*

[9] In simple terms, an interview allocation with full overlap is one in which, whenever two hospitals h_1 and h_2 both interview the same doctor d, then any doctor d' that is interviewed by h_2 is also interviewed by h_1. In balanced markets ($|\mathcal{H}| = |\mathcal{D}|$) with $\beta = 1$, if all hospitals conduct the same number of interviews, such interview allocations guarantee that every hospital will be matched.

[10] This would constitute an equilibrium in non-anonymous strategies.

time (h_1 needs to prefer this doctor to "its own" doctor, and the doctor needs to prefer h_1 over h_2), but would prefer "its own" doctor 50% *of the time. Thus, h_1 would achieve expected utility of* $\frac{3}{4}2/3 + \frac{1}{4}1/3 - 2c_{\mathcal{H}}$ *with the two interviews. Hence, the two-sided interview restriction in which all agents conduct 1 interview each does not form a non-anonymous equilibrium as long as* $c_{\mathcal{H}} < 5/24$*, which covers the range* $(42/648, 121/648)$ *over which the two-sided restriction would be beneficial.*

When interview costs are such that an interview allocation that improves on the equilibrium in anonymous strategies cannot itself be sustained in equilibrium, other means are needed. For instance, if a social planner was able to increase hospitals' interview costs, the increased interview costs could be used to prevent deviations from the welfare-improving interview allocation. However, it is unclear how this could be implemented in practice. Another option would be for the social planner to simply impose the welfare-improving interview allocation. However, this would also not be straightforward to implement in cases in which agents have private pre-interview information about their match values, in which case it may be difficult for a social planner to find a welfare-improving interview allocation.

Several authors have considered the case in which interviews are allocated by a deferred-acceptance style interview assignment mechanism (e.g. Melcher et al. (2018) and Manjunath and Morrill (2021)). Such interview assignment mechanisms have the potential to allow market participants to coordinate on a welfare-improving interview allocation when such interview allocations do not form an equilibrium. The idea behind a centralized ordinal interview assignment mechanism is to impose an interview limit/cap for each agent in the market, elicit participants' pre-interview preferences over agents on the other side, and allocate interviews according to an algorithm that takes as input the interview limits and agents' preferences.

Thus far, work that has considered the use of a centralized match for allocating interviews, which includes Manjunath and Morrill (2021) and Lee and Schwarz (2007), has focused on the use of a deferred-acceptance algorithm with exogenously given interview capacities. Despite the reliance on the deferred-acceptance procedure, the incentive properties of such mechanisms are not well understood. One potential problem is the fact that such mechnaisms are based on *ordinal* information about agents' preferences, while, as we have seen, the interview game involves the trade-off of agents' (cardinal) match utilities and the interview costs they incur. The following example illustrates that, even when the mechanism designer has the ability to prevent participants from conducting interviews outside of those recommended by the mechanism, it is still not straightforward to ensure that participants necessarily conduct all the interviews recommended by the mechanism:

Example 10 *Consider the case with* $|\mathcal{H}| = 2$ *and* $|\mathcal{D}| = 4$*, with* $\beta = 1$*,* $v_{dh} \stackrel{iid}{\sim} U(0,1)$ *and* $v_{hd} \stackrel{iid}{\sim} \theta_h U(0,1)$ *with* $\mathbb{P}(\theta_h = 1) = \mathbb{P}(\theta_h = \frac{1}{3}) = 1/2$ *and* $\theta_h \perp \theta_{h'}$*. Assume* $c_{\mathcal{H}} \in (1/18, 1/12)$*. Due to the abundance of doctors relative to hospitals, one possible interview recommendation would be to divide the doctors equally between hospitals, to avoid any negative welfare externalities. Note that when* $\theta_{hd} = 1$*, the hospital is happy to follow the recommendation from the mechanism, and conduct both interviews. However, when* $\theta_{hd} = 1/3$*, the benefit of the second interview to the hospital is* $\frac{1}{3}(2/3 - 1/2) = 1/18 < c_{\mathcal{H}}$*, meaning the hospital would not want to follow the recommendation. On the other hand, conditional on the other hospital only conducting one interview, a hospital with* $\theta_{hd} = 1$ *would want to conduct three interviews, since the benefit of the third interview is* $3/4 - 2/3 = 1/12 > c_{\mathcal{H}}$*. Hence, in the case where the two hospitals have different draws of* θ*, the interview recommendation that assigns two (different) doctors to each hospital will not be followed, and results in an inefficiency. However, the cases with* $\theta_h = 1$ *and* $\theta_h = \frac{1}{3}$ *contain exactly the same ordinal information about the hospital's match utilities, and therefore would be treated equally by a mechanism that only takes as input ordinal information about agents' preferences.*

Hence, such interview assignment mechanisms may have advantages over relying on market participants coordinating on an equilibrium of their own accord. However, a centralized interview allocation mechanism does not in itself guarantee that market participants will adhere to the mechanisms recommendations.

6.2 Limiting Interview Activity vs. the Number of Applications

Another market design intervention that has been proposed is the introduction of a limit on the number of applications doctors are allowed to send to hospitals. Indeed, some see the introduction of the Electronic Residency Application Service (ERAS) in 1995 as an important contributor to the issues market participants are currently experiencing in the residency matching in the US. The argument is based on the idea that ERAS lowered doctors' costs of sending residency applications, increasing the number of applications doctors would send. As a result, hospitals would, on average, receive more applications. Moreover, hospitals knew that the doctors from whom they received applications likely had also sent applications to many other hospitals. As Burbano et al. (2019) write: "The high application volume imparts significant time and financial burden for applicants and programs alike. Furthermore, it makes distinguishing between applicants with a genuine interest in a specific program and those who are merely hoping to improve their chances vastly more difficult." Based on this argument, Burbano et al. (2019) argue that the increased volume of applications increases hospitals' incentives to conduct more interviews. They cite the work by Arnosti et al. (2021) to conclude that "an application limit is the most reasonable approach to address this issue". Indeed, Arnosti et al. (2021) consider a two-sided matching problem in which firms seek to hire workers through a platform. Firms need time to process applications from applicants, and are willing to hire the first applicant they find acceptable. However, since each worker sends multiple applications, by the time a firm finds an applicant acceptable, with some probability the worker has already been hired by another firm, and the firm must return to its pile of applications. Arnosti et al. (2021) indeed show that the market may benefit from imposing a limit on the number of applications workers are allowed to send. Intuitively, whenever a firm finds a worker acceptable, the application cap reduces the probability that the worker is hired by another firm.

Introducing a limit on the number of applications doctors are allowed to send may have several benefits. If hospitals' costs of processing the applications are high, then an application limit should, on average, reduce hospitals processing costs. With regards to the effect on hospitals' interview decisions, the effect of an application limit is less clear-cut: First, as suggested by Burbano et al. (2019), an application cap may also drive hospitals to lower the number of interviews they conduct. However, this conclusion does not necessarily follow from the analysis in Arnosti et al. (2021): Their model is intended to capture firms' application processing decisions, and not interview decisions. Moreover, firms have no pre-interview preferences, and are equally likely to hire any of the workers from whom they receive an application. Given the complicated nature of the strategic externalities we explored in Section 4.2, it is not immediate that an application limit will necessarily lower the number of interviews hospitals conduct in equilibrium. Second, even if the application limit leads to a reduction in hospital interview activity, it does not in itself help market participants better coordinate their interviews. Third, an application limit opens up an additional source of mis-coordination: Doctors do not necessarily know where other doctors will send their applications. As a result, the distribution of applications across hospitals may have a high variance, which may lead to inefficient outcomes at the interview stage. In this section, we will formally extend our model of interviews to include an application stage for doctors. We will argue that the application limit indeed may improve upon the decentralized interview equilibrium in which all doctors apply to all hospitals. However, we will also illustrate some limitations of an application limit, and show that a two-sided interview restriction may, in principle, lead to even better outcomes. As such, we will debunk the claim that "an application limit is the most reasonable approach" to address the issues currently experienced in the entry-level market for doctors in the US.

In the model analyzed in the previous sections doctors played an entirely passive role, accepting all interview offers they receive. By allowing doctors to choose the hospitals to which they will send their applications, doctors immediately take on a more active role. Several issues emerge: Intuitively, given the underlying uncorrelated and uniform preference structure, if doctors are constrained to only send $K_{\mathcal{D}} < |\mathcal{H}|$ applications, then doctors should want to send their applications to their $K_{\mathcal{D}}$ highest ranked hospitals. However, due to a logic similar to the one showing that the deferred-acceptance algorithm is not strategy-proof for agents on the receiving side, this need not be the case: it will depend on the distributions from which doctors' match utilities are drawn (and how informative doctors' pre-interview preferences are). Moreover, as with hospitals' interview decisions, a doctor's application sending strategy will, in general, depend on the strategy pursued by the other doctors. For the main body of the paper, we will simply assume that doctors send their applications to hospitals at random, reminiscent of an anonymous interview strategy with no pre-interview information. In this case, we can again treat doctors as passive players in the game. In the appendix, we show that under certain conditions on the distribution of doctors' preferences, we construct an extended game in which doctors play an active role, and show that whenever hospitals play anonymous interview strategies, and all doctors play anonymous application-sending strategies, anonymous application-sending strategies are indeed optimal from doctors' perspective. The equilibria in anonymous strategies of this extended game are therefore outcome-equivalent to the equilibria in anonymous strategies in the following simplified game:

Timing of the interview game with application limits $K_{\mathcal{D}}$

1. Doctors send $K_{\mathcal{D}}$ applications at random to hospitals.
2. Hospitals learn their pre-interview information θ_h for the doctors from whom they received an application.
3. Hospitals simultaneously choose which doctors to interview among the doctors from whom they received an application. Doctors accept all interviews.
4. Hospitals and doctors learn their match utilities, v_h and v_d, for those whom they interviewed.
5. Hospitals and doctors all submit rank-order lists, only listing agents (i) with whom they interviewed and (ii) found acceptable.

The definition of equilibrium in this alternative game is identical to before, with hospitals now restricted to only interviewing those doctors from whom they received an application.

Definition 6.1 *Let $\big(\mathcal{H}, \mathcal{D}, c_{\mathcal{H}}, G_{\mathcal{H}}, K_{\mathcal{D}}\big)$ be an interview game with application caps. A* ***(mixed) Bayesian interview equilibrium with application caps*** *is a tuple $\big(\sigma_h(\theta_h, \sigma_d(v_d)^{-1})\big)_{h\in\mathcal{H},\theta_h\in\Theta}$ such that for each $h \in \mathcal{H}$, each $\theta_h \in \Theta$, and for each $(v_d)_{d\in\mathcal{D}}$,*

$$\sigma_h(\theta_h, \sigma_d(v_d)^{-1}) \in \Delta \underset{S\subset\big(\sigma_d(v_d)\big)^{-1}_{d\in D}(h)}{\arg\max} \left\{ \sum_{d\in S} \mathbb{P}\Big(\mu(h) = d \,\big|\, \sigma_{-h}\Big)\mathbb{E}\big[v_{hd} \,\big|\, \mu(h) = d,\, \sigma_{-h}\big] - c_{\mathcal{H}}|S| \right\}$$

where $\big(\sigma_d(v_d)\big)^{-1}_{d\in\mathcal{D}}(h)$ denotes the applications h receives from all the doctors at the preference profile $(v_d)_{d\in\mathcal{D}}$, with $\sigma_d(v_d)$ denoting the $K_{\mathcal{D}}$ highest ranked hospitals according to v_d.

In words, an equilibrium with applications is one in which hospitals, for any pre-interview preferences, and for any set of applications they receive, optimally choose which doctors to interview among those from whom they received applications, given the strategies of all other hospitals. The previous definition of anonymous strategies given in Section 3 still applies to this alternative game.

Note: Receiving an application from a doctor provides a signal about the doctors' interest in the hospital. In particular, the hospital knows it's only competing with $K_{\mathcal{D}} - 1$ other hospitals for the doctor in question. On the other hand, under the same assumption as we introduced in Section 3, receiving an application from, say, doctor d does not provide any additional information about (i) which other hospitals received the remaining $K_{\mathcal{D}} - 1$ applications from d, and (ii) how other hospitals rank d relative to all the other doctors from whom they received applications. Therefore, the model with application limits is equivalent to the original interview game with the addition that hospitals will treat a doctor as unacceptable as long as they did not receive an application from said doctor. Formally:

Lemma 6.2 *Suppose pre-application preferences satisfy Assumptions (A1)-(A4), (A5*), and (A6). Then the extended game with applications is equivalent to the standard interview game (with Assumptions (A1)-(A6) in which hospitals find unacceptable all doctors from whom they didn't receive an application.*

It suffices to verify that the model satisfies Assumption (A5). Note that the applications do introduce a dependency in hospitals' pre-interview preferences. Specifically, conditional on receiving an application from a doctor d, hospital h knows that only $K_{\mathcal{D}} - 1$ other hospitals received an application from this doctor. On the other hand, if h did not receive an application from d, then more other hospitals will have received an application from this doctor. As a result, the stronger independence assumption in (A5*) cannot be satisfied. The proof is in the appendix. Note that the definition of anonymous strategies introduced in Section 3 can immediately be applied to the model with interview caps. As a consequence, we have the following:

Corollary 6.3 *All the results from Sections 3.2 and 4.1 hold for the model with application caps.*

Since the model with interview caps satisfy all of Assumptions (A1)-(A6) of the original interview game, all the results pertaining to negative welfare externalities still hold. The extent of the externalities may be lower, as hospitals face "less competition" for each doctor they interview, but the negative welfare externalities are still present. In particular, this means that a game with application caps has the potential to lead to equilibria with inefficiently many interviews.

We wrap up our discussion of the introduction of application limits by considering an example that both illustrates the benefits of the application restriction, points out the potential for miscoordination of the applications from doctors' perspective, and compares the application restriction to the two-sided interview restriction:

Example 11 *Consider again the market in Example 1 with* $|\mathcal{H}| = |\mathcal{D}| = 2$*, with* $v_{hd}, v_{dh} \overset{iid}{\sim} U(0,1)$ *and* $\beta = 1$*. As long as* $c_{\mathcal{H}} < 5/24$*, the unique equilibrium in anonymous strategies involves both hospitals conducting two interviews each, yielding equilibrium welfare of* $7/12 - 2c_{\mathcal{H}}$*.*

Assume now that doctors are constrained to sending only one application each, while no restriction is placed on interview activity. Doctors send their applications at random, or, alternatively, to their most preferred hospital according to their pre-interview preferences. With probability $1/2$*, the doctors send their applications to different hospitals, and the hospitals conduct one interview each. With probability* $1/2$*, the doctors will send their applications to the same hospital, in which case the hospital in question will interview both doctors if* $c_{\mathcal{H}} < 1/6$*, or otherwise interview one doctor at random. In either case, the hospital that didn't receive any applications will not conduct any interviews. As a result, hospital expected utility under the restricted application scenario is*

$$\frac{1}{2}\Big(1/2 - c_{\mathcal{H}}\Big) + \frac{1}{4}\max\left\{1/2 - c_{\mathcal{H}}, 2/3 - 2c_{\mathcal{H}}\right\} = \begin{cases} \frac{3}{4}\big(1/2 - c_{\mathcal{H}}\big) & \text{if } c_{\mathcal{H}} \in (1/6, 1/2) \\ 5/12 - c_{\mathcal{H}} & \text{if } c_{\mathcal{H}} < 1/6 \end{cases}$$

The application restriction improves upon the equilibrium as long as $c_{\mathcal{H}} \in (1/6, 1/2)$*, while hospitals prefer the equilibrium for* $c_{\mathcal{H}} < 1/6$*.*

Consider now a two-sided restriction on interview activity that limits agents to one interview each, which yield hospital welfare of $1/2 - c_{\mathcal{H}}$*. The two-sided interview restriction is preferred to the application restriction for any* $c_{\mathcal{H}} < 1/2$*, and is only dominated by the equilibrium when* $c_{\mathcal{H}} < 1/12$*.*

As the example above makes clear, an application restriction alone does not prevent market participants from coordinating on an equilibrium that may lower the expected number of matched agents compared to the decentralized equilibrium. Ignoring any costs associated with hospitals' processing of the applications they receive, the example shows the existence of cases in which an application restriction is *not* "the most reasonable approach" to address the issues caused by the costly process through which agents' learn about their preferences.

7 Concluding Remarks

This paper has highlighted new features of matching markets with interviews which to date have not been explored. In particular, hospitals' interview decisions impose negative welfare externalities on other hospitals. Moreover, we have identified two channels that may amplify the extent of these welfare externalities. First, hospitals' decisions to increase their interview activity can prompt other hospitals to also increase their interview activity. Second, hospitals may be exposed to larger welfare externalities as markets become thicker. In both cases, in conjunction with the strategic externalities, the interview strategies that can be supported equilibrium may potentially be Pareto ranked, meaning there some equilibria may be better than others from the perspective of all hospitals. Moreover, even the most efficient equilibrium (in anonymous strategies) may be inefficient and, in particular, involve an inefficiently high number of interviews. Furthermore, when both sides of the market incur interview costs, the equilibrium may be inefficient from the perspective of all market participants, meaning both sides of the market could potentially be made better off with an appropriate market design intervention. Moreover, absent a market design intervention, agents' incentives to participate in the centralized match may decrease as the negative welfare externalities to which they are exposed magnify.

Many market design interventions to address the negative aspects of the pre-match interview process have been proposed, especially in the medical literature, but careful analysis of the proposed mechanisms' properties is still needed to properly evaluate their effects on the market outcomes. In this paper, we have pointed to the benefits and limitations of two interentions that have recently been proposed: We showed that imposing a limit on the number of applications doctors are allowed to send may improve the market outcome, mirroring some earlier results. However, we also pointed out that an application limit does not specifically target the mis-coordination of interviews, and other market interventions may be better suited to achieve this goal. Moreover, the introduction of an application may create a new source of mis-coordination, leading to an inefficient distribution of applications across hospitals. We also discussed the use of interview assignment mechanisms that limit agents' number of interviews, and argued that they are likely to be more successful if they limit/coordinate interview activity on both sides of the market. Early work on the use of such mechanisms has failed to consider their incentive properties. Since mechanisms that are based on variations of the deferred-acceptance algorithm usually take as input *ordinal* information about participants' preferences, we illustrated some difficulties that arise as a result of the *cardinal* nature of agents' true preferences. The analysis and design of interventions that could help market participants better coordinate on more efficient interview allocations remains a promising direction for future research.

A Appendix

A.1 Proofs for Section 3.2

Proof of Proposition 3.3: The proof is in three steps:

1) First, conditional on all other hospitals playing anonymous strategies, each hospital has an optimal strategy that is anonymous: Indeed, conditional on θ_h, the expected utility of any interview set S is equal to the expected utility of the set πS at $\pi\theta_h$, for any permutation π. Therefore, if S maximizes utility at θ_h given the anonymous strategies σ_{-h}, then πS maximizes utility at $\pi\theta_h$ given σ_{-h}.

2) We next define an "expanded game" and apply the standard Nash equilibrium existence theorem to the expanded game: The set of players (hospitals) remains the same. Define $\mathcal{S}$ as the set of all anonymous strategies, and hospitals' payoff functions as

$$\hat{g}(s_h, s_{-h}) = \sum_{\theta_h \in \Theta} \sum_{\theta_{-h} \in \Theta^{|\mathcal{H}|-1}} \prod_{h' \in \mathcal{H}} \mathbb{P}_G(\theta_{h'}) g(s_h(\theta_h), s_{-h}(\theta_{-h}))$$

This expanded game consists of finitely many players with finite strategies, and hence an equilibrium exists.

3) Finally, because each type has positive probability, this ex ante formulation requires that for each type θ_h, h's best response specify an interview set S such that

$$\sum_{\pi} \mathbb{P}_G(\pi\theta_h) \sum_{\theta_{-h}} \mathbb{P}_G(\theta_{-h}) g(\pi S, s_{-h}(\theta_h)) \geq \sum_{\pi} \mathbb{P}_G(\pi\theta_h) \sum_{\theta_{-h}} \mathbb{P}_G(\theta_{-h}) g(\pi S', \sigma_{-h}(\theta_{-h}))$$

for all S', whenever σ_{-h} is anonymous. But any such S must also maximize h's utility at θ_h: If not, by the argument in 1), there exists an alternative interview set S'' which maximizes utility at θ_h, such that $\pi S''$ maximizes utility at $\pi\theta_h$, in which case S would not satisfy the above inequality. Therefore, any Nash equilibrium in the expanded game in 2) specifies, for each hospital h, an anonymous strategy that for each hospital type θ_h stipulates a (mixed) interview set that is optimal given the hospital's type.

■

Proof of Lemma 3.5: Denote by $\mu(h|S)$ hospital h's match when interviewing the set S. Suppose when interviewing the set S', h is matched with d', and when interviewing S, h is matched with d. Let $S \subset S'$.

First, consider a preference realization ω in which $\mu(h|S) = d$ and $\mu(h|S') = d'$. Then $d' \in \{d\} \cup S' \setminus S$: Since the HPDA algorithm is strategy proof, h cannot be worse off when interviewing S', since h could then profitably manipulate the algorithm by only listing the doctors in S. If $d' = d$, then there's noting to prove, so assume $d' \neq d$. If $d' \in S$, then h must prefer d' to d also when only interviewing the set S. If h matches to d' when interviewing S' but not when interviewing S, then h could profitably manipulate the algorithm by listing the doctors in $S' \setminus S$ in the appropriate order, violating the strategy-proofness of the algorithm. So we cannot have $d' \in S \setminus d$. Since this must hold for each such preference realization, it follows that we can write, for each $d \in S \cup \{\emptyset\}$

$$\mathbb{1}\Big(\mu(h|S) = d\Big) = \sum_{d' \in \{d\} \cup S' \setminus S} \mathbb{1}\Big(\mu(h|S) = d,\, \mu(h|S') = d'\Big) \quad \text{and, for each } d' \in S' \cup \{\emptyset\}$$

$$\mathbb{1}\Big(\mu(h|S') = d'\Big) = \begin{cases} \mathbb{1}\Big(\mu(h|S') = d',\, \mu(h|S) = d\Big) & \text{if } d' = d \in S \\ \sum_{d \in S} \mathbb{1}\Big(\mu(h|S) = d,\, \mu(h|S') = d'\Big) & \text{if } d \neq d' \in S' \setminus S \end{cases}$$

Second, consider still a realization of preferences ω in which $\mu(h|S) = d$ and $\mu(h|S') = d'$. Suppose $d = d'$. If, at the same preference realization, we have $\mu(h|S' \cup \{d_k\}) = d_k$, then d_k must be preferred to d. Then we must also have $\mu(h|S \cup \{d_k\}) = d_k$: If h could not match with d_k when only listing the doctors in S, this would violate the strategy-proofness of the HPDA algorithm.

Suppose instead $d' \neq d$. Then d' is preferred to d and $d' \notin S'$ (by strategyproofness of $\mu(\cdot|S')$ for hospitals). If, at the same preference realization, we have $\mu(h|S' \cup \{d_k\}) = d_k$, then d_k must be preferred to d', and hence also preferred to d. Then we must also have $\mu(h|S \cup \{d_k\}) = d_k$. Indeed, if h could not match with d_k when only listing the doctors in S, this would violate the strategy-proofness of the HPDA algorithm. In both cases ($d = d'$ and $d \neq d'$), since this must hold at every such preference realization, it follows that for every $d \in S$ and $d' \in \{d\} \cup S' \setminus S$

$$\begin{aligned} &\mathbb{1}\Big(\mu(h|S' \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big) \\ \leq\ &\mathbb{1}\Big(\mu(h|S \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big) \end{aligned}$$

Third, while the match outcome μ does reveal information about others' preferences, it does not reveal additional information about h's preferences, other than h's ordinal ranking of doctors: Suppose h matches with d when interviewing S, with d' when interviewing $S' \supset S$, and with d_k when interviewing $S' \cup \{d_k\}$. Then the match outcome μ does not reveal any other information about the differences $v_{hd_k} - v_{hd'}$ and $v_{hd_k} - v_{hd}$ than the fact that $v_{hd_k} > v_{hd'} \geq v_{hd}$. It therefore follows that

$$\begin{aligned} &E\Big[v_{hd_k} - v_{hd'} \,\Big|\, \mu(h|S' \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big] \\ = &E\Big[v_{hd_k} - v_{hd'} \,\Big|\, v_{hd_k} > v_{hd'} \geq v_{hd}\Big] \qquad \text{and} \\ &E\Big[v_{hd_k} - v_{hd} \,\Big|\, \mu(h|S \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big] \\ = &E\Big[v_{hd_k} - v_{hd} \,\Big|\, v_{hd_k} > v_{hd'} \geq v_{hd}\Big] \end{aligned}$$

By Assumption (A1), it now follows that

$$\begin{aligned} &E\Big[v_{hd_k} - v_{hd'} \,\Big|\, \mu(h|S' \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big] \\ = &E\Big[v_{hd_k} - v_{hd'} \,\Big|\, v_{hd_k} > v_{hd'} \geq v_{hd}\Big] \\ \leq &E\Big[v_{hd_k} - v_{hd} \,\Big|\, v_{hd_k} > v_{hd'} \geq v_{hd}\Big] \\ = &E\Big[v_{hd_k} - v_{hd} \,\Big|\, \mu(h|S \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big] \end{aligned}$$

Finally, combining the above steps, we get

$$
\begin{aligned}
& V(S \cup \{d,' d_k\}, \sigma_{-h}) - V(S \cup \{d'\}, \sigma_{-h}) - \Big(V(S \cup \{d_k\}, \sigma_{-h})) - V(S, \sigma_{-h})\Big) \\
= & \sum_{d \in S \cup \{\emptyset\}} \sum_{d' \in \{d\} \cup S' \setminus S} \mathbb{P}\Big(\mu(h|S') = d',\, \mu(h|S) = d\Big) \mathbb{P}\Big(\mu(h|S' \cup \{d_k\}) = d_k \Big| \mu(h|S') = d',\, \mu(h|S) = d\Big) \\
& \qquad \times E\Big[v_{hd_k} - v_{hd'} \,\Big|\, \mu(h|S' \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big] \\
& - \sum_{d \in S \cup \{\emptyset\}} \sum_{d' \in \{d\} \cup S' \setminus S} \mathbb{P}\Big(\mu(h|S') = d',\, \mu(h|S) = d\Big) \mathbb{P}\Big(\mu(h|S \cup \{d_k\}) = d_k \Big| \mu(h|S') = d',\, \mu(h|S) = d\Big) \\
& \qquad \times E\Big[v_{hd_k} - v_{hd} \,\Big|\, \mu(h|S \cup \{d_k\}) = d_k,\, \mu(h|S') = d',\, \mu(h|S) = d\Big]\Bigg\} \\
= & \sum_{d \in S \cup \{\emptyset\}} \sum_{d' \in \{d\} \cup S' \setminus S} \mathbb{P}\Big(\mu(h|S') = d',\, \mu(h|S) = d\Big)\Bigg\{ \\
& \quad \mathbb{P}\Big(\mu(h|S' \cup \{d_k\}) = d_k \Big| \mu(h|S') = d',\, \mu(h|S) = d\Big) E\Big[v_{hd_k} - v_{hd'} \,\Big|\, v_{hd_k} > v_{hd'} \geq v_{hd}\Big] \\
& \quad - \mathbb{P}\Big(\mu(h|S \cup \{d_k\}) = d_k \Big| \mu(h|S') = d',\, \mu(h|S) = d\Big) E\Big[v_{hd_k} - v_{hd} \,\Big|\, v_{hd_k} > v_{hd}\Big]\Bigg\} \\
\leq & \, 0
\end{aligned}
$$

$\blacksquare$

Proof of Lemma 3.6: Consider any set S and rank-order list P that ranks a subset of the doctors in S. Let $d', d'' \notin S$ and let P' and P'' be two rank-order lists that rank the doctors in S the same way as P but also rank d' and d'', respectively, in the same spot. First, by assumption (A5) and the fact that h's competitors are playing anonymous strategies, all competitors are equally likely to interview any of the doctors in the market, and hence equally likely to rank any doctor in the deferred acceptance stage. Second, by assumption (A6), no doctor is more likely to rank h higher than any other doctor. As a result, the probability h is matched to d' under P' is the same as the probability h is matched with d'' under P''. Moreover, for any $\tilde{P}''$ that ranks all doctors in S the same as P'', but ranks d'' higher than under P'', the probability that h matches with d'' is weakly higher than under P'' (by the strategy-proofness of the hospital-proposing deferred-acceptance algorithm). This holds for any rank-order list P that ranks a subset of the doctors in S. As a result, it follows

$$
\begin{aligned}
& \mathbb{P}\big(\mu(h|S \cup \{d''\} = d'' \,\big|\, v_{hd''} > v_{h\mu(h|S)},\, v_{h\mu(h|S)}\big) \\
= & \, \mathbb{P}\big(\mu(h|S \cup \{d'\} = d' \,\big|\, v_{hd'} > v_{h\mu(h|S)},\, v_{h\mu(h|S)}\big)
\end{aligned}
$$

Moreover, the increment in the match utility from matching with d when interviewing $S \cup \{d\}$ relative to matching with $\mu(h|S)$ when only interviewing S, does not depend on other agents' preferences, only on h's own match utilities. Hence

$$\mathbb{E}\big[\mathbb{1}\big\{v_{hd} > v_{h\mu(h|S)}\big\}\big(v_{hd} - v_{h\mu(h|S)}\big) \,|\, \mu(h|S \cup \{d\} = d, v_{h\mu(h|S)}\big]$$
$$\mathbb{E}\big[\mathbb{1}\big\{v_{hd} > v_{h\mu(h|S)}\big\}\big(v_{hd} - v_{h\mu(h|S)}\big) \,|\, v_{h\mu(h|S)}\big]$$

Note that for every y, the function $\mathbb{1}\{x > y\}(x - y)$ is increasing in x. By first-order stochastic dominance, we therefore get

$$\begin{aligned}
&V(S \cup \{d''\}, \sigma_{-h}) - V(S, \sigma_{-h}) = \mathbb{E}\big[\mathbb{1}\big\{\mu(h|S \cup \{d''\}) = d''\big\}\big(v_{hd''} - v_{h\mu(h|S)}\big)\big] \\
&= \mathbb{E}\Big[\mathbb{E}\big[\mathbb{1}\big\{\mu(h|S \cup \{d''\}) = d'',\, v_{hd''} > v_{h\mu(h|S)}\big\}\big(v_{hd''} - v_{h\mu(h|S)}\big) \,\big|\, v_{h\mu(h|S)}\big]\Big] \\
&= \mathbb{E}\Big[\mathbb{P}\big(\mu(h|S \cup \{d''\}) = d'' \,|\, v_{hd''} > v_{h\mu(h|S)}\big)\, \mathbb{E}\big[\mathbb{1}\big\{v_{hd''} > v_{h\mu(h|S)}\big\}\big(v_{hd''} - v_{h\mu(h|S)}\big)\big]\Big] \\
&\geq \mathbb{E}\Big[\mathbb{P}\big(\mu(h|S \cup \{d'\}) = d' \,|\, v_{hd'} > v_{h\mu(h|S)}\big)\, \mathbb{E}\big[\mathbb{1}\big\{v_{hd'} > v_{h\mu(h|S)}\big\}\big(v_{hd'} - v_{h\mu(h|S)}\big)\big]\Big] \\
&= V(S \cup \{d'\}, \sigma_{-h}) - V(S, \sigma_{-h})
\end{aligned}$$

■

Proof of Proposition 3.4: For every $\theta \in \Theta$, and for every $k \in \{1, ..., |\mathcal{D}|\}$, let $d_\theta^{(k)}$ denote (any of) the k-th highest ranked doctors according to θ, such that $S_\theta^k := \{d_\theta^{(1)}, ..., d_\theta^{(k)}\}$ is a selection of the k highest ranked doctors according to θ.

1. For any $k \leq |\mathcal{D}|$, conditional on conducting k interviews, it is always optimal to interview a selection of the k highest ranked doctors. Suppose not, then there exists a k and a set $S \subset \mathcal{D}$ with $|S| = k$ such that $V(S, \sigma_{-h}) > V(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h})$, but S is not a selection of the k highest ranked doctors. Let $\{d_1, ..., d_m\}$ be an ordering of the doctors in S according to θ. Then for every $j = 1, ..., k$, $\theta_h^j \geq \theta_{hd_j}$. Replacing, for each j, doctor d_j in S with $d_\theta^{(j)}$ in $S_\theta^{(k)}$ we find, using Lemma 3.6

$$\begin{aligned}
&V\Big(\big\{d_{\theta_h(1)}, ..., d_{\theta_h(k)}\big\}, \sigma_{-h}\Big) - V\Big(S, \sigma_{-h}\Big) \\
&= \sum_{n=1}^{k} V\Big(S \cup_{j=1}^{n} \big\{d_\theta^{(j)}\big\} \setminus \cup_{j=1}^{n}\big\{d_j\big\}, \sigma_{-h}\Big) - V\Big(S \cup_{j=1}^{n-1} \big\{d_\theta^{(j)}\big\} \setminus \cup_{j=1}^{n-1}\big\{d_j\big\}_{j=1}^{n-1}, \sigma_{-h}\Big) \geq 0
\end{aligned}$$

contradicting $V(S, \sigma_{-h}) > V(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h})$

2. Combining Lemma 3.5 with Part 1 above, we find that the marginal benefit curve, while interviewing a selection of the k highest ranked doctors, is a decreasing function of k:

$$\begin{aligned} & V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k+1)}\}, \sigma_{-h}\Big) - V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h}\Big) \\ \geq\ & V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k-1)}\} \cup \{d_\theta^{(k+1)}\}, \sigma_{-h}\Big) - V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k-1)}\}, \sigma_{-h}\Big) \\ \geq\ & V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h}\Big) - V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k-1)}\}, \sigma_{-h}\Big) \end{aligned}$$

Whenever the marginal benefit exceeds the interview costs

$$V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h}\Big) - V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k-1)}\}, \sigma_{-h}\Big) > c_{\mathcal{H}}$$

it is optimal to interview at least k doctors. If

$$V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h}\Big) - V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k-1)}\}, \sigma_{-h}\Big) < c_{\mathcal{H}}$$

then it's optimal to interview strictly less than k doctors. If

$$V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k)}\}, \sigma_{-h}\Big) - V\Big(\{d_\theta^{(1)}, ..., d_\theta^{(k-1)}\}, \sigma_{-h}\Big) = c_{\mathcal{H}}$$

then hospital h is indifferent between conducting k and $k-1$ interviews. ■

A.2 Proofs for Section 4.1

For the proof of Lemma 4.1 we will use a known result pertaining to agents' welfare when a hospital extends its rank-order list in the final matching stage (see Theorem 2 in Gale and Sotomayor (1985) or Theorem 2.24 in Roth and Sotomayor (1992). For convenience, we state their result here, rephrased to fit the context of hospitals and doctors. For notation, let P_i denote the rank-order list of agent i, and write $P_i' \geq P_i$ if i adds previously unacceptable agents to the bottom of P_i (i.e. appending to the rank-order list, leaving the ranking of acceptable agents under P_i unchanged):

Lemma A.1 (Gale and Sotomayor (1985)) *Suppose $P_h' \geq P_h$ for all $h \in \mathcal{H}$ and let $\mu_{\mathcal{H}}'$, $\mu_{\mathcal{H}}$, $\mu_{\mathcal{D}}'$, and $\mu_{\mathcal{D}}$ be the corresponding hospital- and doctor-optimal matchings. Then under the preferences P the hospitals are not worse off and the doctors are not better off in $(\mathcal{H}, \mathcal{D}, P)$ than in $(\mathcal{H}, \mathcal{D}, P')$, no matter which of the two optimal matchings are considered. That is*

$$\mu_{\mathcal{H}}(h) P_h\, \mu_{\mathcal{H}}'(h)\ \forall\, h \in \mathcal{H} \quad \textit{(so } \mu_{\mathcal{H}}(d) P_d\, \mu_{\mathcal{H}}'(d)\ \forall\, d \in \mathcal{D} \textit{ by the stability of } \mu_{\mathcal{H}}'\textit{), and}$$

$$\mu_{\mathcal{D}}(d)' P_d\, \mu_{\mathcal{D}}(d)\ \forall\, d \in \mathcal{D} \quad \textit{(so } \mu_{\mathcal{D}}(h) P_h\, \mu_{\mathcal{D}}'(h)\ \forall\, h \in \mathcal{H} \textit{ by the stability of } \mu_{\mathcal{D}}\textit{)}$$

Proof of Theorem 4.1: First, if $\sigma'_{-h} \geq \sigma_{-h}$, then we can create a finite sequence $\left(\sigma_{-h}^{(i)}\right)_{i=0}^{n}$ with $\sigma_{-h}^{(0)} = \sigma_{-h}$ and $\sigma_{-h}^{(n)} = \sigma'_{-h}$ such that for any i, $\sigma_{-h}^{(i)}$ and $\sigma_{-h}^{(i-1)}$ differ for only one competitor hospital, such that $\sigma_{-h}^{(i)} \geq \sigma_{-h}^{(i-1)}$. The conjectured result must hold for every i, and it therefore suffices to prove the result for $\sigma'_{-h} \geq \sigma_{-h}$, with $\sigma'_{-h} = \left(\sigma'_{h'}, \sigma_{-h,h'}\right)$ for any h', i.e. when only one of the hospitals increases its interview activity.

Every $\sigma_{h'}$ induces a distribution $\mathbb{P}_{\sigma_{h'}}$ over rank-order lists $\mathcal{P}$, which can further be refined by conditioning on $\theta_{h'}$, which we denote by $\mathbb{P}_{\sigma_{h'}(\theta_{h'})}$. For any θ, denote by $\mathbb{P}_{\sigma_{h'}(\mathbb{S}(\theta))}$ the distribution over rank-order lists conditional on $\theta_h \in \mathbb{S}(\theta)$, the set of all permutations of θ. If $\sigma_{h'}$ is anonymous, as a consequence of Assumption (A5), for every rank-order list P and every permutation π we have $\mathbb{P}_{\sigma_{h'}}(P) = \mathbb{P}_{\sigma_{h'}}(\pi P)$. More specifically, for every $\theta \in \Theta$, and every $P \in supp\left(\mathbb{P}_{\sigma_{h'}(\theta)}\right)$, we have $\mathbb{P}_{\sigma_{h'}(\mathbb{S}(\theta))}(\pi P) = \mathbb{P}_{\sigma_{h'}(\mathbb{S}(\theta))}(P)$.

If $\sigma'_{h'} \geq \sigma_{h'}$, then $\sigma'_{h'}$ induces a distribution over rank-order lists whose lengths first-order stochastically dominate those of the rank-order lists induced by $\sigma_{h'}$. Moreover, if $\sigma_{h'}, \sigma'_{h'}$ are both anonymous and $\sigma'_{h'} \geq \sigma_{h'}$, then for every $\theta \in \Theta$, for every $P \in supp\left(\mathbb{P}_{\sigma_{h'}(\mathbb{S}(\theta))}\right)$ and $P' \in supp\left(\mathbb{P}_{\sigma'_{h'}(\mathbb{S}(\theta))}\right)$, if $|P'| \geq |P|$ then $\exists \pi$ such that $P \leq \pi P'$. By Assumption (A5), if $\sigma'_{h'} \geq \sigma_{h'}$ it therefore follows that for every $\theta \in \Theta$, and every $P \in supp\left(\mathbb{P}_{\sigma_{h'}(\theta)}\right)$ we have $\mathbb{P}_{\sigma_{h'}(\mathbb{S}(\theta))}(P) \leq \sum_{P':P' \geq P} \mathbb{P}_{\sigma'_{h'}(\mathbb{S}(\theta))}(P')$. Since this holds for all P, by Lemma A.1, it follows all hospitals other than h' are worse off under $\sigma'_{h'}$ than under $\sigma_{h'}$.

Note: Lemma A.1 also implies that doctors' expected match utilities are higher under $\sigma'_{h'}$ than under $\sigma_{h'}$. ■

Proof of Lemma 4.3:

$$\mathbb{E}\left[|\mu(\sigma')| - |\mu(\sigma)|\right] = \mathbb{E}\Big[\sum_{h \in \mathcal{H}} \left|\mu_h(\sigma')\right| - \left|\mu_h(\sigma)\right|\Big] \leq \mathbb{E}\Big[\sum_{h \in \mathcal{H}} \left|\mu_h\left(\sigma'_h, \sigma_{-h}\right)\right| - \left|\mu_h(\sigma)\right|\Big]$$
$$\leq \sum_{h \in \mathcal{H}} \mathbb{E}\Big[\left(1 - \left|\mu_h(\sigma)\right|\right)\left(\left|\sigma'_h\right| - \left|\sigma_h\right|\right)\Big] = \sum_{h \in \mathcal{H}} \sum_{\theta \in \Theta} \mathbb{E}\left[\left(1 - \mu_h(\sigma_h(\theta), \sigma_{-h})\right)\left(\left|\sigma'_h(\theta)\right| - \left|\sigma_h(\theta)\right|\right)\right]$$
$$= \sum_{h \in \mathcal{H}} \sum_{\theta \in \Theta} \mathbb{E}\left[\left(1 - \mu_h(\sigma_h(\theta), \sigma_{-h})\right)\right]\mathbb{E}\left[\left|\sigma'_h(\theta)\right| - \left|\sigma_h(\theta)\right|\right] \leq \sum_{h \in \mathcal{H}} \mathbb{E}\left[\left|\sigma'_h\right| - \left|\sigma_h\right|\right]\left(1 - \underline{\mu_h}(\sigma)\right)$$

The first inequality follows from Theorem 4.1: the negative welfare externalities of other hospitals' interview activity can never increase the match rate of any individual hospital. The second inequality uses that an increase in h's interviews only affects h's match rate when both (i) the preference realizations are such that h would otherwise be unmatched, i.e. $|\mu_h(\sigma)| = 0$, in which case the match rate can only increase by 1, and (ii) the strategy is such that at the particular instance of pre-interview information, h is increasing its interview activity, i.e. $|\sigma_h(\theta_h)| - |\sigma_h(\theta_h)| > 0$. The last equality follows from the fact that $\big(1 - \mu_h(\sigma_h(\theta), \sigma_{-h})\big)$ and $\big|\sigma'_h(\theta)\big| - \big|\sigma_h(\theta)\big|$ are uncorrelated, since σ' and σ are anonymous. The third inequality is immediate from the definition of $\underline{\mu_h}(\sigma)$. ■

Proof of Lemma : First, from Theorem 4.1 we know that an increase in hospital interview activity can be seen as hospitals appending previously unacceptable doctors to the bottom of their rank-order lists, implying that $\mathbb{E}\big[v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)} \,|\, \mu_d(\sigma) \neq \emptyset\big] \geq 0$. Second, since doctors' match utilities are bounded above $\overline{v}_{\mathcal{D}}$, any increase in utility must be bounded by $\overline{v}_{\mathcal{D}}$. Third, as hospitals extend their rank-order lists, any rejection chain can only be set off by a hospital which is unmatched under σ. (Conditional on some hospital setting off a rejection chain, other hospitals may benefit from extending their rank-order lists.) Hence, a doctor's match utility increases only in the cases in which (i) some hospital h is unmatched under σ, i.e. $|\mu_h(\sigma)| = \emptyset$, and (ii) h increases its own interview activity. Fourth, an increase in interview activity by any hospital h increases the expected number of acceptable doctors by $\beta\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big]$. The last inequality follows from the fact that σ' and σ are anonymous, hence for all $\theta \in \Theta$, the random variables $|\mu_h(\sigma_h(\theta), \sigma_{-h})|$ and $|\sigma'_h(\theta)| - |\sigma_h(\theta)|$ are uncorrelated. ■

Proof of Corollary 4.5: First, by Assumption (A5), any doctor is equally likely to be matched, meaning $\mathbb{P}(\mu_d(\sigma)) = \mathbb{E}[|\mu(\sigma)|]/|\mathcal{D}|$ for every $d \in \mathcal{D}$. Second, as the number of matches increases, every previously unmatched doctor is equally likely to "obtain" any of the new matches: the probability of matching under σ', conditional on not matching under σ, equals the expected increments in matches relative to the initial match rate. Third, combining the first two points, the probability of being matched under σ' but not under σ is equal to the incremental number of matches, divided by the number of doctors. Conditional on matching under σ', the match utility can be no higher than the maximum $\overline{v}_{\mathcal{D}}$. Using this, together with Lemmas 4.3 and A.2 we get

$$
\begin{aligned}
&\mathbb{E}\big[v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)}\big] \\
&= \mathbb{E}\big[\mathbb{1}\{\mu_d(\sigma') \neq \emptyset, \mu_d(\sigma) = \emptyset\} v_{d\mu_d(\sigma')}\big] + \mathbb{E}\big[\mathbb{1}\{\mu_d(\sigma'), \mu_d(\sigma) \neq \emptyset\}(v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)})\big] \\
&= \mathbb{P}\big(\mu_d(\sigma) = \emptyset\big)\mathbb{P}\big(\mu_d(\sigma') \neq \emptyset \,|\, \mu_d(\sigma) = \emptyset\big)\mathbb{E}\big[v_{d\mu_d(\sigma')} \,\big|\, \mu_d(\sigma) = \emptyset \neq \mu_d(\sigma')\big] \\
&\qquad + \mathbb{P}\big(\mu_d(\sigma) \neq \emptyset\big)\mathbb{E}\big[v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)} \,\big|\, \mu_d(\sigma) \neq \emptyset\big] \\
&\leq \frac{\beta \overline{v}_{\mathcal{D}}}{|\mathcal{D}|} \sum_{h \in \mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] + \frac{\mathbb{E}[|\mu(\sigma)|]}{|\mathcal{D}|}\beta(\overline{v}_{\mathcal{D}} - \underline{v}_{\mathcal{D}}) \sum_{h \in \mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] \\
&\leq \frac{\beta}{|\mathcal{D}|}\Big[\overline{v}_{\mathcal{D}} + |\mathcal{H}|(\overline{v}_{\mathcal{D}} - \underline{v}_{\mathcal{D}})\Big] \sum_{h \in \mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] \qquad \blacksquare
\end{aligned}
$$

Proof of Proposition 4.6: First, by Assumptions (A5) and (A4), every doctor is equally likely to receive any interview and equally likely to be hired when hospitals play an anonymous strategy σ. Since $v_{dh} \in [\underline{v_{\mathcal{D}}}, \overline{v_{\mathcal{D}}}]$, the worst that can happen to a matched doctor is to achieve a utility of $\underline{v_{\mathcal{D}}}$. At the same time, any doctor d can receive an offer from at most all of the hospitals. Using that $\overline{v_{\mathcal{D}}} \geq \gamma(\underline{v_{\mathcal{D}}} - \overline{v_{\mathcal{D}}}$, the expected utility of doctors under any (anonymous) strategy σ can be bounded below by

$$
\begin{aligned}
\mathbb{E}\Big[v_{d\mu(\sigma)} - \frac{c_{\mathcal{D}}}{|\mathcal{D}|}\sum_{h \in \mathcal{H}} |\sigma_h|\Big] &\geq \underline{v_{\mathcal{D}}} \frac{\sum_{h \in \mathcal{H}} \mathbb{E}\big[|\mu_h(\sigma)|\big]}{|\mathcal{D}|} - c_{\mathcal{D}} \frac{\sum_{h \in \mathcal{H}} \mathbb{E}\big[|\sigma_h|\big]}{|\mathcal{D}|} \\
\geq \frac{\underline{v_{\mathcal{D}}}}{|\mathcal{D}|}\mathbb{E}\big[|\mu(\sigma)|\big] - c_{\mathcal{D}}|\mathcal{H}| \geq \frac{\overline{v_{\mathcal{D}}}}{|\mathcal{D}|}\frac{\gamma - 1}{\gamma}\mathbb{E}\big[|\mu(\sigma)|\big] - c_{\mathcal{D}}|\mathcal{H}| &\geq \frac{\overline{v_{\mathcal{D}}}}{|\mathcal{D}|}\frac{\gamma - 1}{\gamma}|\mathcal{H}|\underline{\mu} - c_{\mathcal{D}}|\mathcal{H}|
\end{aligned}
$$

where $|\underline{\mu_{\mathcal{H}}}(\sigma)|$ is the lowest match probability for any hospital type under σ. Assuming $|\underline{\mu_{\mathcal{H}}}(\sigma)|, |\underline{\mu_{\mathcal{H}}}(\sigma')| > \underline{\mu}$, we find that it's individually rational for doctors to accept all interviews under both σ and σ' as long as

$$
c_{\mathcal{D}} \leq \frac{\overline{v_{\mathcal{D}}}}{|\mathcal{D}|}\frac{\gamma - 1}{\gamma}\underline{\mu}
$$

Second, every doctor is equally likely to receive any of the additional interview under σ' relative to σ, hence the expected increase in interview costs are given by $\frac{c_{\mathcal{D}}}{|\mathcal{D}|}\sum_{h \in \mathcal{H}} \mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big]$. By Corollary 4.5, we can bound the increase in doctors' expected match utilities from above as a function of

$\overline{v_{\mathcal{D}}} - \underline{v_{\mathcal{D}}}$ and hospitals' match probabilities. Using that $\overline{v_{\mathcal{D}}} > \gamma(\overline{v_{\mathcal{D}}} - \underline{v_{\mathcal{D}}})$ and $|\mu_h(\sigma)| > \underline{\mu}$, we get

$$\begin{aligned}
&\mathbb{E}\big[v_{d\mu_d(\sigma')} - v_{d\mu_d(\sigma)}\big] \leq \frac{\beta}{|\mathcal{D}|}\Big[\overline{v}_{\mathcal{D}} + (\overline{v}_{\mathcal{D}} - \underline{v}_{\mathcal{D}})\mathbb{E}\big[|\mu(\sigma)|\big]\Big] \sum_{h\in\mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] \\
&\leq \frac{\beta\overline{v}_{\mathcal{D}}}{|\mathcal{D}|}\Big[1 + \frac{|\mathcal{H}|}{\gamma}\Big] \sum_{h\in\mathcal{H}} \big(1 - |\underline{\mu_h}(\sigma)|\big)\mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] \\
&\leq \frac{\beta\overline{v}_{\mathcal{D}}}{|\mathcal{D}|}\Big[1 + \frac{|\mathcal{H}|}{\gamma}\Big](1 - \underline{\mu}) \sum_{h\in\mathcal{H}} \mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] \leq \frac{c_{\mathcal{D}}}{|\mathcal{D}|} \sum_{h\in\mathcal{H}} \mathbb{E}\big[|\sigma'_h| - |\sigma_h|\big] \\
\Leftrightarrow \quad & c_{\mathcal{D}} \geq \beta\overline{v_{\mathcal{D}}}\Big[1 + \frac{|\mathcal{H}|}{\gamma}\Big](1 - \underline{\mu})
\end{aligned}$$

Combining the two inequalities, we can find an open interval between both constraints as long as

$$\begin{aligned}
&\beta\overline{v_{\mathcal{D}}}\Big[1 + \frac{|\mathcal{H}|}{\gamma}\Big](1 - \underline{\mu}) < \frac{\overline{v_{\mathcal{D}}}}{|\mathcal{D}|}\frac{\gamma - 1}{\gamma}\underline{\mu} \\
\Leftrightarrow \quad & \underline{\mu} > \frac{\beta|\mathcal{D}|\big(1 + |\mathcal{H}|/\gamma\big)}{1 - 1/\gamma + \beta|\mathcal{D}|\big(1 + |\mathcal{H}|/\gamma\big)} \quad \xrightarrow{\gamma\to\infty} \quad \frac{\beta|\mathcal{D}|}{\beta|\mathcal{D}| + 1}
\end{aligned}$$

That is, by choosing γ sufficiently large, there exists $\varepsilon > 0$ such that if $\underline{\mu} \in (\frac{\beta|\mathcal{D}|}{\beta|\mathcal{D}|+1} + \varepsilon, 1)$, then both constraints on $c_{\mathcal{D}}$ are satisfied, and the result follows.

A.3 Proofs for Section 4.2

Proof of Lemma 4.8: Fix θ_h. Let $S, S' \subseteq \mathcal{D}$ and $\overline{S} \subseteq S$, $\overline{S}' \subseteq S'$ and orderings P, P' over S and S', respectively.

σ_{-h} induces a distribution $\mathbb{P}_{\sigma_{-h}}$ over other hospitals' rank-order lists. By Assumption (A4), the realization of h's match utilities, for any set of doctors interviewed by h, does not provide any additional information about the (distribution over the) other hospitals' rank-order lists. By Assumption (A5) and the fact that σ_{-h} is anonymous, for any P_{-h} in the support of the distribution of rank-order lists induced by σ_{-h}, we have $\mathbb{P}_{\sigma_{-h}}(\tilde{\pi}P_{-h}) = \mathbb{P}_{\sigma_{-h}}(P_{-h})$ for any permutation $\tilde{\pi}$. As a result, regardless of h's actual realization of preferences, any rank-order list P submitted by h of length k must lead to a match with the j-th ranked doctor with the same probability, for $j \leq k$. Moreover, suppose $P' \geq P$, i.e. P' appends to (the bottom of) P doctors listed as unacceptable under P. If for some $j \leq |P|$, the probability h matches with its j-th highest ranked doctor differs under P and P', then there must exist a realization of preferences at which h could profitably deviate by either reporting P' when h's real preferences are P, or vice versa. However, this contradicts the hospital-proposing deferred-acceptance algorithm being strategy-proof for the hospitals. The result follows. ■

Proof of Proposition 4.11: For every $\theta \in \Theta$, and for every $k \in \{1, ..., |\mathcal{D}|\}$, let $d_\theta^{(k)}$ denote (any of) the k-th highest ranked doctors according to θ, such that $S_\theta^k := \{d_\theta^{(1)}, ..., d_\theta^{(k)}\}$ is a selection of the k highest ranked doctors according to θ.

1. Since $|\Theta|$ is finite, there exists a $\theta \in \Theta$ such that $V_{h,\theta}(S_\theta^2, \sigma_{-h}) - V_{h,\theta}(S_\theta^1, \sigma_{-h})$ for every strategy profile σ_{-h} played by the competitors: The benefit of the additional interview is a continuous function in the mixing of the competitors' strategies. Since every competitor has a finite number of pure strategies and a finite number of pre-interview information θ_{-h} with which they could be endowed, for every hospital-type (h, θ) there exists a competitor strategy profile σ_{-h} which maximizes $V_{h,\theta}(S_\theta^2, \sigma_{-h}) - V_{h,\theta}(S_\theta^1, \sigma_{-h})$. Since there's a finite number of types θ, there also exists a type that achieves the largest maximum of $V_{h,\theta}(S_\theta^2, \sigma_{-h}) - V_{h,\theta}(S_\theta^1, \sigma_{-h})$ (as a function of σ_{-h}. Letting $\underline{c}$ exceed this maximum, no hospital with any pre-interview information would ever interview more than one doctor for any $c_\mathcal{H} > \underline{c}$. As competitors' increase their number of interviews, because of Lemma 4.9, no hospital will optimally increase its own interview activity in response, and the result follows.

2. There are two cases to consider:

Case (i): $\beta \in (0, 1)$: The benefit of an additional interview is always strictly positive. Since $|\Theta|$ is finite, we can choose $\bar{c}$ low enough such that for any $c_\mathcal{H} < \bar{c}$, any hospital-type pair (h, θ) would, when no other hospital is conducting any interview, want to interview at least a set of doctors S satisfying $\beta\varepsilon|S| \geq (1 - \beta + \varepsilon)|\mathcal{H}|$. Using Lemma 4.9, hospital h's benefit of an additional interview beyond S never deacreases as its competitors increase their interview activity, and the result follows.

Case (ii): $\beta = 1$: The condition reduces to $|\mathcal{D}| \geq |\mathcal{H}|$. Since $|\Theta|$ is finite, we can find $\bar{c}_1$ small enough such that if $c_\mathcal{H} < \bar{c}_1$, then for every $\theta \in \Theta$ and any strategy σ_{-h} we have $V_{(h,\theta)}\big(\{d_\theta^{(1)}\}, \sigma_{-h}\big) > c_\mathcal{H}$, that is, if interview costs are sufficiently low, then every hospital will always (strictly) want to conduct at least one interview, regardless of other hospitals' strategies. For any hospital conducting no more than $|\mathcal{H}| - 1$ interviews, with positive probability, the hospital will be unmatched, while interviewing $|\mathcal{H}|$ doctors will guarantee a match. By assumption $(A2)$, $V_\theta\big(S_\theta^{(|\mathcal{H}|)}, \sigma_{-h}\big) - V_\theta\big(S_\theta^{(|\mathcal{H}|-1)}, \sigma_{-h}\big) > 0$ for any competitor strategy profile σ_{-h}. By an argument similar to the previous, we can find $\bar{c}_2 > 0$ such that every hospital will want to conduct at least $|\mathcal{H}|$ interviews. But then, for any S with $|S| \geq |\mathcal{H}| - 1$, we have

$$V_{\theta_h}\big(S \cup \{d\}, \sigma_{-h}\big) - V_{\theta_h}\big(S, \sigma_{-h}\big)$$
$$= \mathbb{E}\big[v_{h(1,S\cup\{d\})} - v_{h(1,S)}\big] + \sum_{i=1}^{|\mathcal{H}|-1} \prod_{j \leq i} q_j(\sigma_{-h}) \Big($$
$$\mathbb{E}\big[v_{h(i+1,\underline{S}\cup\{d\})} - v_{h(i+1,\underline{S})} - \big(v_{h(i,\underline{S}\cup\{d\})} - v_{h(i,\underline{S})}\big)\big]\Big)$$

Since $\prod_{j \leq i} q_j(\sigma_{-h})$ is increasing in σ_{-h}, and $E\big[v_{h(i+1,\underline{S}\cup\{d\})} - v_{h(i+1,\underline{S})} - \big(v_{h(i,\underline{S}\cup\{d\})} - v_{h(i,\underline{S})}\big)\big] \geq 0$, a hospital who finds it optimal to interview any $k \geq |\mathcal{H}|$ when all hospitals interview at least 1 doctor each, will never want to reduce its number of interviews as its competitors increase their interview activity. ■

Lemma A.2 *Suppose $|\mathcal{D}| \geq |\mathcal{H}|$. For every $h \in \mathcal{H}$, $\theta_h \in \Theta$ we have* $\sum_{i=1}^{|\mathcal{H}|-1} q_{\theta_h,i} \geq \frac{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}-1}{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}}$.

Proof: By Lemma 4.8, for every permutation π we have $q_{\pi\theta,i} = q_{\theta,i}$ for every i. Consider the case in which h lists exactly $|\mathcal{H}| - 1$ doctors at the final matching stage. There are $\binom{|\mathcal{D}|}{|\mathcal{H}|-1}$ different combinations of $|\mathcal{H}| - 1$ doctors that could be listed by h (i.e. different hospital h types). For a hospital h type to be unmatched, all of the $|\mathcal{H}| - 1$ doctors listed by h must be matched to another hospital. But then every other h type that lists at least one other doctor not in the set of $|\mathcal{H}| - 1$ doctors matched to the other hospitals will necessarily be matched. This holds for every realization of preferences: At most one of the $\binom{|\mathcal{D}|}{|\mathcal{H}|-1}$ combinations of h types will be unmatched, hence at least a fraction $\frac{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}-1}{\binom{|\mathcal{D}|}{|\mathcal{H}|-1}}$ of the hospital h types will be matched at every preference realization. The result follows.

A.4 Proofs for Section 5.2

Proof of Lemma 5.5: For each k, p_0 is monotonically increasing, and strictly increases after no more than n steps defined as the smallest n such that $1 - p_0 < \delta$. Every time p_0 increases, the step size is at least δ. Hence, for every k, the algorithm finishes in a finite number of steps. Since k is increases monotonically from 0 to no more than $|\mathcal{D}|$, the algorithms converges in finite time.

Hospital welfare in the symmetric equilibrium equals $V(k, ([k,p])_{-h}) - kc_{\mathcal{H}}$: Either the equilibrium is pure, and all hospitals conduct exactly k interviews, or the equilibrium is mixed, and every hospital is indifferent between conducting the $(k + 1)$-th interview or not.

In the former case, we have $g_k([k,0]) \leq c_{\mathcal{H}}$, in which case the algorithm ends at step 1 (upon reaching k), and the algorithm finds the exact optimum hospital welfare in any symmetric equilibrium in anonymous strategies. In the latter case, the welfare-maximizing equilibrium satisfies $g_k([k,p^*]) = c_{\mathcal{H}}$ for some $p \in (0,1)$. For a given k, $g_k([k,p])$ is a (finite) polynomial function in p, and hence uniformly continuous on $[0,1]$. If p^* is the minimum p satisfying the equilibrium condition, there exists $\epsilon_g > 0$ such that $g_k([k,p]) > c_{\mathcal{H}} + \varepsilon_g$ for all $p < p^*$. Moreover, there exists δ such that $|p - p^*| < \delta_1$ then $|g_k([k,p]) - g_k([k,p^*])| = |g_k([k,p]) - c_{\mathcal{H}}| < \epsilon_g$. Similarly, $V(k, ([k,p])_{h' \neq h})$ is continuous, meaning for every $\epsilon > 0$ there exists $\delta_2 > 0$ such that $|V(k, ([k,p])_{h' \neq h}) - V(k, ([k,p^*])_{h' \neq h})| < \epsilon$ whenever $|p - p^*| < \delta_2$. By choosing $\delta = \min(\delta_1, \delta_2)$, the algorithm ends at the first time $p^* - p < \delta$, hence the evaluated welfare at p will satisfy $|V(k, ([k,p])_{h' \neq h}) - V(k, ([k,p^*])_{h' \neq h})| < \epsilon$.

A similar argument shows that if there is no equilibrium for $j < k$, then the algorithm will not converge for j and continue to $j+1$ when δ is chosen sufficiently small, meaning the algorithm eventually arrives at k and stops within ϵ of the equilibrium value.

Proof of Lemma 6.2: It suffices to show that Assumptions (A5) and (A6) are satisfied. Since doctors' send applications at random to hospitals, receiving an application does not provide additional information to hospitals about doctors' preferences over hospitals. As a result, Assumption (A6) is satisfied. Similarly, since hospitals underlying pre-application preferences satisfy Assumption (A5*), conditional on receiving applications from a set $D \subset \mathcal{D}$, hospital h does not learn anything about (i) the identity of the hospitals who received the remaining $K_{\mathcal{D}} - 1$ applications from the doctors in D, (ii) the identity of the hospitals who received the $K_{\mathcal{D}}$ applications from the doctors in $\mathcal{D} \setminus D$, and (iii) how any other hospital rank the doctors from whom they received applications. As a result, both parts (a) and (b) of Assumption (A5) are satisfied.

A.5 An extended interview game with applications

Definition A.3 *Let $(\mathcal{H}, \mathcal{D}, c_{\mathcal{H}}, G_{\mathcal{H}}, K_{\mathcal{D}})$ be an interview game with application caps. A **(mixed) Bayesian interview equilibrium with application caps** is a tuple $\left(\left(\sigma_h(\theta_h)\right)_{h \in \mathcal{H}, \theta_h \in \Theta}, \left(\sigma_d(v_d)\right)_{d \in \mathcal{D}, v_d \in F_{\mathcal{D}}}\right)$ such that for each $h \in \mathcal{H}$, each $\theta_h \in \Theta$, and for each*

$(v_d)_{d\in\mathcal{D}}$,

$$\sigma_d(v_d) \in \Delta \underset{S\subset\mathcal{H}}{\arg\max} \Big\{ \sum_{h\in S} \mathbb{P}\Big(\mu(d) = h \,\big|\, (\sigma_{h'})_{h'\in\mathcal{H}}, \sigma_{-d}\Big) v_{dh}$$

$$\sigma_h(\theta_h) \in \Delta \underset{S\subset \left(\sigma_d(v_d)\right)^{-1}_{d\in D}(h)}{\arg\max} \Big\{ \sum_{d\in S} \mathbb{P}\Big(\mu(h) = d \,\big|\, \sigma_{-h}\Big) \mathbb{E}\big[v_{hd} \,\big|\, \mu(h) = d,\, \sigma_{-h}\big] - c_{\mathcal{H}}|S| \Big\}$$

where $\left(\sigma_d(v_d)\right)^{-1}_{d\in\mathcal{D}}(h)$ denotes the applications h receives from all the doctors at the preference profile $(v_d)_{d\in\mathcal{D}}$. That is, doctors optimally choose which hospitals to apply to, and, given the received applications, hospitals optimally choose which doctors to interview, given the strategies of all others.

Proposition A.4 *Assume doctors' match values are bounded within some interval $[\underline{v}_{\mathcal{D}}, \overline{v}_{\mathcal{D}}]$. Suppose hospitals are playing anonymous strategies as described in Section 3. There exists $\alpha > 0$ and $\overline{c}$ such that if $\underline{v}_{\mathcal{D}} > \alpha(\overline{v}_{\mathcal{D}} - \underline{v}_{\mathcal{D}})$ and $c_{\mathcal{D}} < \overline{c}$, then doctors find it optimal to send their applications to their $K_{\mathcal{D}}$ highest ranked hospitals.*

Proof: A doctor may deviate from the prescribed strategy in two ways: Either (i) sending strictly fewer applications or (ii) sending an application to a hospital h' and not to h even though $v_{dh} > v_{dh'}$. Since every application is equally likely to result in an interview, for a given number of applications the doctor sends, deviation (ii) is equivalent to the doctor misreporting their preferences at the final matching stage. However, as shown in Lee and Schwarz (2017) (see the proof of Lemma 1 in their appendix), one can find $\alpha > 0$ such that if the distribution of doctors' match values satisfies the above condition, such a deviation is never optimal. Deviation (i) is similar to listing a hospital as unacceptable at the final matching stage, with the added benefit that the doctor also avoids the potential interview cost incurred by sending the application. However, since every application has the chance of being the sole proposal the doctor receives during the deferred-acceptance algorithm at the final matching stage, there exists a sufficiently low interview cost, in combination with the restriction on doctors' match utilities, such that the doctor finds it optimal to send all applications.

References

Agarwal, Nitin, Phillip A Choi, David O Okonkwo, Daniel L Barrow, and Robert M Friedlander, "Financial burden associated with the residency match in neurological surgery," *Journal of neurosurgery*, 2017, *126* (1), 184–190.

Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan, "Thickness and information in dynamic matching markets," *Journal of Political Economy*, 2020, *128* (3), 783–815.

Allen, Franklin and Douglas Gale, "Limited market participation and volatility of asset prices," *The American Economic Review*, 1994, pp. 933–955.

Arnosti, Nick, Ramesh Johari, and Yash Kanoria, "Managing congestion in matching markets," *Manufacturing & Service Operations Management*, 2021.

Bernstein, Joseph, "Not the Last Word: Predicting Chaos in the Residency Match," *Clinical Orthopaedics and Related Research®*, 2021, *479* (3), 452–457.

Beyhaghi, Hedyeh and Éva Tardos, "Randomness and Fairness in Two-Sided Matching with Limited Interviews," in James R. Lee, ed., *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, Vol. 185 of *Leibniz International Proceedings in Informatics (LIPIcs)* Schloss Dagstuhl–Leibniz-Zentrum für Informatik Dagstuhl, Germany 2021, pp. 74:1–74:18.

Burbano, Felipe Molina, Amy Yao, Nikki Burish, Michael Ingargiola, Matthew Freeman, Jeffrey Stock, and Peter J Taub, "Solving congestion in the plastic surgery match: A game theory analysis," *Plastic and reconstructive surgery*, 2019, *143* (2), 634–639.

Chade, Hector, Gregory Lewis, and Lones Smith, "Student portfolios and the college admissions problem," *Review of Economic Studies*, 2014, *81* (3), 971–1002.

Das, Sanmay and Zhuoshu Li, "The role of common and private signals in two-sided matching with interviews," in "International Conference on Web and Internet Economics" Springer 2014, pp. 492–497.

Gale, David and Lloyd S Shapley, "College admissions and the stability of marriage," *The American Mathematical Monthly*, 1962, *69* (1), 9–15.

_ **and Marilda Sotomayor**, "Some remarks on the stable matching problem," *Discrete Applied Mathematics*, 1985, *11* (3), 223–232.

Gan, Li and Qi Li, "Efficiency of thin and thick markets," *Journal of Econometrics*, 2016, *192* (1), 40–54.

Kadam, Sangram V, "Interviewing in Matching Markets with Virtual Interviews," 2015.

Knuth, Donald Ervin, "Marriages stables," *Technical report*, 1976.

Lee, Robin S. and Michael A. Schwarz, "Signalling Preferences in Interviewing Markets," in Peter Cramton, Rudolf Müller, Eva Tardos, and Moshe Tennenholtz, eds., *Computational Social Systems and the Internet*, number 07271. In 'Dagstuhl Seminar Proceedings.' Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany Dagstuhl, Germany 2007.

Lee, Robin S and Michael Schwarz, "Interviewing in two-sided matching markets," *The RAND Journal of Economics*, 2017, *48* (3), 835–855.

Lien, Yuanchuan, "Costly Interviews and Non-Assortative Matchings," 2009.

Loertscher, Simon and Ellen V Muir, "The benefits of market thickness for niche products," 2021.

Manjunath, Vikram and Thayer Morrill, "Interview Hoarding," *arXiv preprint arXiv:2102.06440*, 2021.

Melcher, Marc L, Itai Ashlagi, and Irene Wapnir, "Matching for fellowship interviews," *Jama*, 2018, *320* (16), 1639–1640.

Monir, Joseph G, "Reforming the Match: a proposal for a new 3-phase system," *Journal of graduate medical education*, 2020, *12* (1), 7–9.

Ngai, L Rachel and Silvana Tenreyro, "Hot and cold seasons in the housing market," *American Economic Review*, 2014, *104* (12), 3991–4026.

Program, Data Release National Residency Matching and Research Committee, "Results of the 2017 NRMP Applicant Survey," 2017.

Rastegari, Baharak, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown, "Two-sided matching with partial information," in "Proceedings of the fourteenth ACM conference on Electronic Commerce" 2013, pp. 733–750.

_ , _ , _ , Robert Irving, and Kevin Leyton-Brown, "Reasoning about optimal stable matchings under partial information," in "Proceedings of the fifteenth ACM conference on Economics and computation" 2014, pp. 431–448.

Roth, Alvin E, "A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom," *The American economic review*, 1991, pp. 415–440.

_ , "What have we learned from market design?," *The Economic Journal*, 2008, *118* (527), 285–310.

_ **and Elliott Peranson**, "The redesign of the matching market for American physicians: Some engineering aspects of economic design," *American economic review*, 1999, *89* (4), 748–780.

_ **and Marilda Sotomayor**, "Two-sided matching," *Handbook of game theory with economic applications*, 1992, *1*, 485–541.

_ **and Xiaolin Xing**, "Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists," *Journal of political Economy*, 1997, *105* (2), 284–329.

Strand, Eric A and Tammy S Sonn, "The residency interview season: time for commonsense reform," *Obstetrics & Gynecology*, 2018, *132* (6), 1437–1442.

Talcott, Wesley J and Suzanne B Evans, "I Need a Sign: The Growing Need for a Signaling Mechanism to Improve the Residency Match," *International Journal of Radiation Oncology, Biology, Physics*, 2021, *109* (2), 329–331.

USMLE, United States Medical Licensing Examination, "Change to pass/fail score reporting for Step 1," *https://www.usmle.org/incus/#decision*, 2020.

Wapnir, Irene, Itai Ashlagi, Alvin E Roth, Erling Skancke, Akhil Vohra, and Marc L Melcher, "Explaining a Potential Interview Match for Graduate Medical Education," *Journal of Graduate Medical Education*, 2021, *13* (5).

Watson, Shawna L, Robert H Hollis, Lasun Oladeji, Shin Xu, John R Porterfield, and Brent A Ponce, "The burden of the fellowship interview process on general surgery residents and programs," *Journal of surgical education*, 2017, *74* (1), 167–172.

Wong, Brian JF, "Reforming the match processâearly decision plans and the case for a consortia match," *JAMA Otolaryngology–Head & Neck Surgery*, 2016, *142* (8), 727–728.