Theoretical Economics 11 (2016), 1–39 1555-7561/20160001

Strategic complementarities and unraveling in matching markets

Federico Echenique
Division of Humanities and Social Sciences, California Institute of Technology

Juan Sebastián Pereyra
ECARES—Solvay Brussels School of Economics and Management, Université libre de Bruxelles and F.R.S.-FNRS

We present a theoretical explanation of inefficient early matching in matching markets. Our explanation is based on strategic complementarities and strategic unraveling. We identify a negative externality imposed on the rest of the market by agents who make early offers. As a consequence, an agent may make an early offer because she is concerned that others are making early offers. Yet other agents make early offers because they are concerned that others worry about early offers, and so on and so forth. The end result is that any given agent is more likely to make an early offer than a late offer.

Keywords. Stable matching, Gale–Shapley, two-sided matching, unraveling, strategic complementarities, assortative matching.

JEL classification. C72, D78, D82.

1. Introduction

We study unraveling in labor markets and in matching markets in general. Unraveling is a phenomenon by which matches are made too early. They are made at a point in time when there is too little information about the quality of a match. The literature has documented many episodes of unraveling: the market for medical interns is a famous example in which labor contracts for interns were signed two years before the future interns would graduate (see Roth 1984 or Roth and Oliveira Sotomayor 1990). Other examples of unraveling include the market for federal court clerks (Avery et al. 2001, Roth

Federico Echenique: fede@hss.caltech.edu
Juan Sebastián Pereyra: jpereyra@ulb.ac.be
We are grateful to David Cantala, Serafín Frache, Drew Fudenberg, Andrés Sambarino, Juan Gabriel Brida, Doron Zeilberger, Szilvia Pápai, Renzo Sprugnoli, Christoph Koutschan, Kaniska Dam, Francis Bloch, Estelle Cantillon, Patrick Legros, Alvaro Forteza, Juan Dubra, Néstor Gandelman, Andrei Gomberg, Julio Davila, César Martinelli, Gastón Cayssials, and Alberto Pinto for their comments and suggestions, as well as seminar participants at Paris School of Economics, 19th Coalition Theory Network Workshop, ITAM, ORT–Uruguay, DECON–Uruguay, Universidad de Montevideo, CORE–UCL, ECARES, and Computer Science Department at ULB, the Free University of Bolzano, and El Colegio de México. Finally, we are very grateful to the co-editor and two anonymous referees for their suggestions, which have led to a much improved version of the paper. Pereyra gratefully acknowledges financial support from ERC Grant 208535. This paper was previously circulated under the title "Strategic uncertainty and unraveling in matching markets."

DOI: 10.3982/TE1831

2013), for gastroenterology fellows (Niederle and Roth 2003, 2004), for college football games (Fréchette et al. 2007, Roth 2012), and for placement in sororities (Mongell and Roth 1991).

We explain unraveling of the timing of offers as the result of *strategic unraveling*. If some agents go early, it becomes more attractive for other agents to go early, which makes it more attractive for even more agents to go early. Our explanation is reminiscent of models of bank runs, where strategic complementarity makes agents undertake an inefficient action because they are concerned that others may take this inefficient action (Diamond and Dybvig 1983). Our theory of unraveling also relies crucially on strategic complementarities, but the matching environment is quite different from models of bank runs. The basic logic of strategic unraveling is, however, similar to bank runs.

Strategic unraveling in our model proceeds as follows. There is a loss in efficiency when some agents go early: Information about the quality of the matches arrives late, so it is better for efficiency to wait until the information has arrived to make a match. If some agents go early anyway, the strategic complementaries built into our model force later matches to be less efficient (and a quantification of this effect is the main technical result of this paper). The consequence is a negative externality that makes it more tempting for all agents to go early. So the negative externality may push some additional agents over the threshold by which they decide to go early. In turn, these additional agents going early makes it even more tempting to go early—and so on and so forth.

It should be intuitive that the negative externality causes unraveling, but how far does it go? In our model, we can precisely *calculate* the extent to which strategic unraveling pushes agents to go early. It turns out that unraveling goes all the way to making each individual more likely to go early than to go late.

Our model assumes that there are two periods, and that there is incomplete information over the agents' discount factor. If an agent goes early, she has no information about the quality of a match. If an agent goes late, then all information has been released and matching is assortative on the quality of an agent as a partner (highest quality agents match with each other, the second highest match with each other, and so on).

Incomplete information about the discount factor is how we "seed" unraveling. It makes some agents go early, which makes other agents go early, and so on. We need to motivate this role of the discount factor, i.e., to explain what the assumed uncertainty over the discount factor corresponds to in actual instances of unraveling in matching markets. There are two possible explanations.

The first explanation is that the discount factor captures a pure preference for going early. It could be that agents start to produce as soon as they are matched, but in many matching markets, early matching is not associated with producing earlier. Furthermore, we do not assume that all information about match qualities are resolved earlier. A pure preference for early matching is, however, reasonable even if there is no early production and even if match qualities are not resolved earlier. A pure preference for early matching could reflect a preference for early resolution of uncertainty as in Kreps and

Porteus (1978).[1] Such preferences could result from the ability to make relationship-specific investments, such as arranging to move cities or learning new skills. Matching early allows agents to undertake match-specific investments that can only be undertaken once the identity of the match partner is resolved. Doctors may learn about the particular techniques or personnel at a hospital once they learn where they will be matched. Investments take time to realize, and matching early is therefore valuable. Finally, a desire to smooth consumption would mean that the early elimination of uncertainty about income, for example, could be beneficial to credit constrained agents.

The second explanation is that we use incomplete information purely as a modeling device. Equilibrium theory requires that agents be sure of the behavior of others. We use incomplete information to introduce uncertainty over others' actions (even though, in equilibrium, they know everyone's strategies). In this interpretation, our model says that when agents are not totally sure of the actions of other agents, then unraveling is likely to result.

A more precise statement of our results follows. We first assume that only firms are strategic. Workers always accept the offers they receive. In this environment, we show that there is always a full unraveling Bayesian Nash equilibrium in which all firms make early offers. Further, in *any* symmetric Bayesian Nash equilibrium, a firm makes an early offer with probability at least $\frac{3}{4}$.

If we assume that the prior over discount factors is uniform, we can say more. There are exactly two symmetric equilibria when the size of the market is at least 11. One is the full unraveling equilibrium, but it is unstable. In the second equilibrium, which is stable, agents go early with probability larger than $\frac{3}{4}$. As the size of the market grows, the probability of going early in the second equilibrium converges to $\frac{3}{4}$. If the number of agents is lower than or equal to 10, the unique symmetric Bayesian Nash equilibrium is the full unraveling equilibrium.

We then consider a version of the model in which both sides are strategic. Our results continue to apply (there is actually not a substantial conceptual difference between the two models). Among other things, we prove that in any symmetric Bayesian Nash equilibrium, the expected proportion of agents who match early is at least $\frac{1}{2}$.

Our results reveal that there may exist an equilibrium pattern of adherence and non-adherence to the hiring dates. The market may become divided in equilibrium, with one segment hiring early and the other waiting to match in the final period with full information about agents' qualities. We demonstrate that a mixed level of adherence can be sustainable in an equilibrium, which is consistent with the empirical evidence (Avery et al. 2001).

[1] As anyone who has gone through the job market in economics know, there is a value to early resolution of uncertainty, even if all jobs start in the Fall of the next academic year. The same is true of academic departments as well, not only of job candidates: If we know early on that we have hired in one field, we can use our slots to hire in other fields. One can think of uncertainty in our model as being resolved when agents match, even if production occurs later in time.

Finally, we should emphasize that there are ways in which our model is rigged against unraveling. It makes late matching particularly attractive and rules out unraveling purely as the result of coordination failure (see Section 4). Yet the model produces early contracting as the modal outcome.[2]

We should also emphasize that the notion of efficiency used in the paper reflects the standard use of the term in discussions of unraveling in labor market. It is efficient for all agents to wait and match late. Strictly speaking, the presence of time discounting may entail some agents going early even in an efficient outcome. The negative externality identified in our paper means, however, that the equilibrium outcomes will in any case involve inefficiently many agents going early.

1.1 *Related literature*

Ours is the first theoretical study that identifies strategic complementarities as the main force behind the unraveling of matching markets. One empirical investigation of the market for medical interns also attributes unraveling to strategic complementarities: Wetz et al. (2010) write that early contracting is motivated by concerns over losing interns to other programs that operate outside of the centralized algorithm. Their explanation, based on agents' observed behavior in the market, is essentially what we have tried to capture formally in the present paper.

The best known episode of unraveling is the case of the market for hospital interns before 1945 (Roth 1984, 2002, Roth and Oliveira Sotomayor 1990). There is evidence that unraveling still exists in this market: Wetz et al. (2010) study out-of-match residency offers during the year 2007. In the market for interns, some interns are allowed to take outside-the-match offers (for instance, osteopathic medical students and international medical graduates). Wetz et al. (2010) find that 15.7% of the total number of postgraduate year-1 positions available in the three primary care and four procedural and/or lifestyle-oriented specialities studied, were offered outside the match. The authors conclude that about one in five positions in nonprocedural, primary care specialties were offered outside the match and, thus, the situation is similar to that which existed before 1952.

One classic explanation of unraveling is the "stability hypothesis" as formulated by Roth (1991) and Kagel and Roth (2000). This hypothesis affirms that unraveling will be prevented if once the relevant information is revealed, a stable matching is implemented through a clearinghouse. The idea is that, in some sense, the market is trying to establish a stable matching. It simply may be doing so in an inefficient manner. Our paper provides some justification for central clearinghouses. There is a clear efficiency gain from late contracting in our model, and late contracting equals a stable matching. The agents'

[2]Continuing with the similarity with bank runs, the result is reminiscent of the literature on global games, where basic assumptions on the structure of signals give a precise calculation of how far iterated elimination of dominated strategies will go (Frankel et al. 2003). There is, however, a clear difference with the literature on bank runs. A run can be explained purely by coordination failure. Agents' payoffs in our model are biased against unraveling, and coordination failure alone would not suffice to make agents unravel (see Section 4 for a discussion of this issue).

strategic behavior prevents the market from reaching this stable matching and makes the market unravel.

A handful of other papers provide theoretical explanations for unraveling. They focus on different mechanisms than the one we study here.

Li and Rosen (1998) and Li and Suen (2000) study a model with transfers (a model based on Shapley and Shubik's (1972) assignment game) in which early contracting provides insurance. They show that unraveling may occur among workers who appear to be most promising a priori before full information is revealed. In a similar framework, Li and Suen (2004) allow for unproductive firms and find multiple equilibria with unraveling. They show that more firms and workers will contract early if the uncertainty about the number of productive workers is higher and the more risk-adverse agents are. As we explain in Section 4, our model does not have an insurance motive for early contracting and focuses on a different explanation for unraveling.

Damiano et al. (2005) present an explanation of unraveling that is based on search and matching. Agents know their qualities, so there is no informational gain from matching late, but an agent may not meet a partner of sufficiently high quality in a given period. If there are costs to searching, then there is unraveling in how willing agents are to accept a partner. In Damiano et al. (2005), unraveling is triggered by search costs. In our model, it is triggered by incomplete information.

Du and Livne (2014) consider the role of transfers in unraveling. They show that in the absence of transfers and in the limit as the market size grows, a substantial number of agents will contract early. Unraveling in their paper happens because new agents arrive over time, and agents who are in relatively high positions may want to contract early because the new arrivals may be of higher match qualities. In contrast, in a flexible-transfer regime, agents will not unravel.

Niederle et al. (2013) explain unraveling as the result of an imbalance between demand and supply. Unraveling arises when there is a surplus of applicants, but a shortage of high quality applicants. When a worker does not know if she will be in the long or short side of the market, she may find early offers made by low quality firms attractive. For such firms, early offers is the only way to employ high quality workers.

Hałaburda (2010) proposes that the key to explaining unraveling is the similarity of firms' preferences. Workers' preferences for firms are identical, and known from the start, but firms learn their preferences for workers in the second period. If firms' preferences are similar, then firms tend to prefer the same workers. Thus, worse firms may have better chances to hire their most preferred candidates if they make early offers. So, if firms' preferences are sufficiently similar, it is likely that some firms will go early. In our model, although preferences are identical, this feature does not explain unraveling. An agent may be concerned about being one of the worst agents in the market, but she would still prefer to wait and contract in the second period. Early contracting in our model is inefficient for every agent. As we show below, the uncertainty over how many other agents go early is the main mechanism behind incentives for some agents to match early.

2. The model and results

We present a model of one-to-one matching between workers and firms. In describing the model, we adopt the language of the medical interns market. The workers are doctors and the firms are hospitals.

Let H and D be two finite and disjoint sets: H is the set of hospitals and D is the set of doctors. Suppose that $|H| = |D| = n$, so we can identify H and D with (copies of) $\{1, \dots, n\}$.

A *matching* is a function $\mu : H \cup D \to H \cup D$ such that, for all $h \in H$ and $d \in D$,

(a) $\mu(h) \in D \cup \{h\}$ and $\mu(d) \in H \cup \{d\}$

(b) $d = \mu(h)$ if and only if $h = \mu(d)$.

The meaning of $\mu(h) = h$ is that the position of hospital h remains unfilled, and $\mu(d) = d$ means that doctor d does not find a job.

Each doctor d and hospital h is assigned a *quality*

$$\pi^D(d) \in \{1, \dots, n\} \quad \text{and} \quad \pi^H(h) \in \{1, \dots, n\}.$$

Suppose that π^H and π^D are permutations of $\{1, \dots, n\}$, so we can think of quality as the rank of a hospital or doctor in the market. The highest ranked hospital is h such that $\pi^H(h) = n$, for example. If doctor d is hired by hospital h, then they obtain utilities that depend on their qualities: $u_d(\pi^D(d), \pi^H(h))$ is the utility to d and $u_h(\pi^D(d), \pi^H(h))$ is the utility to h. If an agent remains unmatched, then she obtains a utility of zero.

A matching μ is *stable* if there is no pair (h, d) such that

$$u_d(\pi^D(d), \pi^H(h)) > u_d(\pi^D(d), \pi^H(\mu(d)))$$
$$u_h(\pi^D(d), \pi^H(h)) > u_h(\pi^D(\mu(h)), \pi^H(h)).$$

We assume that u_d and u_h are *multiplicative*, that is, $u_d(i, j) = u_h(i, j) = ij$.

Remark 1. There is a unique stable matching—the matching defined as $\mu(d) = h$ if and only if $\pi^D(d) = \pi^H(h)$ (the identity matching).

2.1 *Matching over time: Early or late offers*

The model is a stylized environment with two periods. In the first period, match qualities π^H and π^D are not known. In the second period, a pair (π^H, π^D) is drawn at random, uniformly and independently. A match is formed among the agents who wish to match in period $t = 0$: all agents are identical at that point, so the matching is purely random. In the second period, when match qualities are known, a stable matching is formed among the agents who did not match in the first period.

Our purpose is to focus on the strategic motivations for going early. We begin with a simultaneous-move game in which only hospitals decide whether to go early and match at time $t = 0$ or to wait and match at time $t = 1$. In particular, we assume that only hospitals are strategic and that matchings are automatic. In period $t = 1$, the matching is assortative among the agents who have not matched in period $t = 0$; the assortative matching is the unique stable matching under our assumptions. In period $t = 0$, matching is random because no agent has any information on match qualities.

In Section 3, we present results where both doctors and hospitals are strategic. Our results essentially continue to hold when both sides are strategic, but we choose to present first the model in which only hospitals are strategic. The reason is twofold. First, there is no deep conceptual difference between the two cases. Indeed, we use the results in this section to prove the results of Section 3. Second, the discussion of unraveling in Roth (1984) suggests that in the hospital–interns market, only hospitals are strategic.

Each agent $i \in H \cup D$ has a discount factor δ_i. The utility at $t = 0$ when h and d match in period t is given by

$$\delta_h^t u_h(\pi^D(d), \pi^H(h)) = \delta_h^t \pi^D(d)\pi^H(h)$$
$$\delta_d^t u_d(\pi^D(d), \pi^H(h)) = \delta_d^t \pi^D(d)\pi^H(h),$$

to h and d, respectively.[3]

The following time line describes how events unfold:

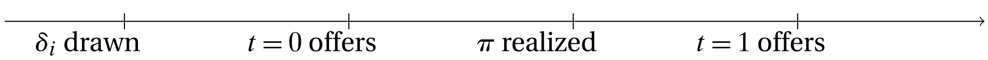

We proceed to describe the payoffs from making an early versus a late offer to match. At time $t = 0$, qualities are purely random. So if a hospital h matches in period 0, its expected utility is $\mathcal{U}_e = (1/n^2)\sum_{i=1}^n \sum_{j=1}^n ij$, the expected value of the product ij when i and j are random.

In period 1, agents have learned the values of π^D and π^H. The matching will be assortative among the agents who have not matched early. Assortative means that the doctor with the highest value of $\pi^D(d)$ will match with the hospital with the highest value of $\pi^H(h)$, the doctor with the next-highest value of $\pi^D(d)$ will match with the hospital with the next-highest value of $\pi^H(h)$, and so on.

Now, it is complicated to calculate the expected utility of going late because the calculation depends on how many agents go early. If m agents have left the market in time $t = 0$, then the assortative matching matches the highest *available* hospital and doctor, but the actual highest quality matches may have left early. The problem is compounded as we consider the second-highest qualities, the third-highest, and so on.

One special case is simple to calculate. Consider a given hospital h. If *all other* hospitals wait to make offers in period $t = 1$, then the expected utility to hospital h in period 0 of waiting for period 1 is $\delta_h(1/n)\sum_{i=1}^n i^2$.

In general, if m hospitals have left the market, we write $\mathcal{U}_m$ for the expected value of $\pi^H(h)\pi^D(\tilde{\mu}(h))$, where π^H and π^D are random, and $\tilde{\mu}$ is the (random) assortative

[3]As discussed in the Introduction, the discount factor is simply a modeling device to seed the strategic complementarities over agents going early. It reflects individual-level heterogeneity in pure preferences for going early.

matching in period 1. The matching $\tilde{\mu}$ is determined by the realization of match qualities π^H and π^D, including the qualities of the m hospitals, with corresponding doctors, who have left the market. That is, when m hospitals exit the market at $t = 0$, $\mathcal{U}_m$ is the expected utility to a hospital of waiting for $t = 1$.

The following is an important technical result in our paper.

Lemma 1. *We have*

$$\mathcal{U}_m = \frac{(n+1)^2(2(n-m)+1)}{6(n-m+1)}.$$

An important consequence of Lemma 1 is that $\mathcal{U}_m > \mathcal{U}_{m+1}$. The difference $\mathcal{U}_{m+1} - \mathcal{U}_m$ is the *negative externality* imposed by a hospital–doctor pair who match early on the agents who decide to match late, when m pairs have already decided to go early. It is important to note that the negative externality, that is, $\mathcal{U}_{m+1} - \mathcal{U}_m$, increases with m, so that additional agents going early increase the incentives of any given agent to go early. This effect vanishes as the market grows large, which helps to stabilize the number of agents who go early in a large market (see the discussion after Corollary 3).

Section 5 gives a precise definition of the quantity $\mathcal{U}_m$ and presents a proof of Lemma 1.

2.2 *Complete information*

As a simple benchmark, consider the model when $\delta_i = 1$ for all $i \in H \cup D$. Note that the expected utility of an early offer is $\mathcal{U}_e = \frac{1}{4}(n+1)^2$, and the expected utility of waiting when m hospitals go early is $\mathcal{U}_m$, with $m \in \{0, \dots, n-1\}$. In particular, if all hospitals decide to go early, the expected utility of waiting is $\mathcal{U}_{n-1} = \mathcal{U}_e$. When $\delta = 1$, given that $\mathcal{U}_m$ is decreasing, we have that $\mathcal{U}_e \leq \sum_{i=0}^{n-1} P(m = i)\mathcal{U}_i$. So it is a weakly dominant strategy to wait. Alternatively, when all hospitals decide to go early, the hospital is indifferent between going early or waiting (in both cases it receives $\mathcal{U}_e = \mathcal{U}_{n-1}$).

Thus, when $\delta = 1$, all hospitals waiting to match is an equilibrium, but there is also an "unraveling" equilibrium in which all agents match early. This unraveling, or early matching, equilibrium is in weakly dominant strategies and is unstable, but it illustrates the effect of strategic complementarities on agents' incentives to wait or to match early. We shall see that when we introduce uncertainty into the game, then unraveling becomes unavoidable in equilibrium.

2.3 *Incomplete information*

We now introduce a Bayesian game in which hospitals may make early offers due to the uncertainty over how many other hospitals go early.

We assume that $\delta_h \in [0, 1]$ is the private information of hospital h. The *type* of an agent h is therefore δ_h. All agents share the prior that the different δ_h are drawn independently from a distribution over $[0, 1]$ with cumulative distribution function (c.d.f.) F. We assume that $x \leq F(x)$ for all $x \in [0, 1]$: the assumption is satisfied for any distribution with a concave c.d.f. For example, the uniform, or truncated normal, distributions on [0, 1] satisfy our assumption.

A strategy for a hospital h is a function

$$s_h : [0,1] \to \{0,1\},$$

where $s_h(\delta_h)$ is the period in which hospital h makes its offer. In our model, there is no decision to be made other than *when* to match.[4]

Given a profile of strategies $s = (s_1, \dots, s_n)$, we write s_{-h} for the profile of strategies of hospitals other than h. Given a profile s_{-h}, for each realization of δ_{-h}, s_{-h} determines m, the number of hospitals, other than h, that go early. Thus, s_{-h} defines a probability distribution for m. Given a profile s_{-h}, m is a random variable, and so is $\mathcal{U}_m$ with a distribution defined by F. Then we can compute the expected value of $\mathcal{U}_m$ given s_{-h} (see Lemma 1), which is denoted by $\mathbf{E}_{s_{-h}}\mathcal{U}_m$. We write $\delta_h \mathbf{E}_{s_{-h}}\mathcal{U}_m$ for the expected utility at time 0, to hospital h, of waiting for $t = 1$ to make an offer, if all hospitals other than h have the profile of strategies s_{-h}: $\mathbf{E}_{s_{-h}}\mathcal{U}_m = \sum_{i=0}^{n-1} \Pr(m = i)\mathcal{U}_i$, where as we just noted, $\Pr(m = i)$ is calculated from s_{-h} and F.

Given a profile s_{-h}, a hospital h will decide to go early if and only if

$$\mathcal{U}_e \geq \delta_h \mathbf{E}_{s_{-h}}\mathcal{U}_m \tag{1}$$

(recall that $\mathcal{U}_e$ is the expected utility of making an early offer). So a strategy s_h is a best response to s_{-h} if for every δ_h, $s_h(\delta_h) = 0$ if and only if (1) is satisfied.

A profile of strategies $s = (s_1, \dots, s_n)$ is a *Bayesian Nash equilibrium* (BNE) if (1) is satisfied for each $h \in H$. A BNE is *symmetric* if $s_h = s_{h'}$ for all $h, h' \in H$. A BNE is *full unraveling* if $s_h = 0$ for all $h \in H$. Thus, in a full unraveling BNE, all agents go early no matter their type.

Theorem 1. *If $n \leq 10$, then the unique symmetric BNE is the full unraveling BNE. If $n > 10$, then there is at least one symmetric BNE, namely the full unraveling BNE; moreover, in any symmetric BNE $s = (s_1, \dots, s_n)$, we have that*

$$\Pr(s_h = 0) \geq F\left(\frac{3}{4}\right) \geq \frac{3}{4}$$

for all $h \in H$, where $\Pr(s_h = 0)$ is the probability that a hospital decides to match early.

Theorem 1 says that any hospital in any symmetric BNE is more likely to go early than late. The equilibrium probability of going early is at least $\frac{3}{4}$. The following corollary is therefore immediate.

Corollary 1. *In any symmetric BNE, the expected number of hospitals that go early is at least $nF(\frac{3}{4}) \geq \frac{3}{4}n$.*

[4] We focus on pure strategies. Mixed strategy equilibria tend to be unstable in games of strategic complementarities (Echenique and Edlin 2004).

2.4 *Stability of BNE: Uniform F*

In this section, we entertain an additional assumption. We suppose that the prior distribution F is the uniform c.d.f. In this case, we can make more precise statements about the set of BNE in our game. We can also talk about the stability of equilibria.

As we shall see, for large n, in the unique stable equilibrium, the market is divided. Most of the market ($\frac{3}{4}$ of all hospitals) go early, while the rest wait and contract late. Thus our results with a uniform F can explain some of the empirical findings where only part of the market unravels.

Theorem 2. *Let F be the uniform c.d.f. If $n \leq 10$, then the unique symmetric BNE is the full unraveling BNE. If $n > 10$, then there are exactly two symmetric BNE. One is the full unraveling BNE. The second is a BNE $s^n = (s_1^n, \dots, s_n^n)$ in which, for every $h \in H$,*

$$\Pr(s_h^n = 0) \geq \tfrac{3}{4} = \lim_{n\to\infty} \Pr(s_h^n = 0),$$

where $\Pr(s_h^n = 0)$ is the probability that a hospital decides to match early.

Remark 2. The proof of Theorem 1 actually follows from Theorem 2. We lay out the details in Section 7.

We discuss a notion of stability of BNE. Stability allows us to select a symmetric BNE in the cases in which there is more than one. It turns out that the full unraveling BNE is stable when $n \leq 10$, and the equilibrium denoted by s^n in Theorem 2 is the unique stable symmetric BNE when $n > 10$.

A strategy s_h satisfying (1) is characterized by a *threshold* $\bar{\delta}_h \in [0, 1]$ such that $s_h(\delta_h) = 0$ if $\delta_h \leq \bar{\delta}_h$ and $s_h(\delta_h) = 1$ if $\delta_h > \bar{\delta}_h$. Given identical thresholds $\bar{\delta}_{-h} = \bar{\delta}$ for all hospitals other than h, we can let $\beta^n(\bar{\delta})$ be the threshold for hospital h defined by (1).

A symmetric BNE is then described by a single $\bar{\delta} \in [0, 1]$ with the property that

$$\bar{\delta} = \beta^n(\bar{\delta}).$$

The function β^n is the best-response function of our game. The symmetric BNE are the fixed points of β^n. Figure 1 presents the graph of β^n for $n = 3, 7, 11, 15, 17$, and, as we will show, in general it holds that $\beta^n \geq \beta^m$ if $n \leq m$.

A symmetric BNE $\bar{\delta}$ is *stable* if there is an open interval I of $\bar{\delta}$ in $[0, 1]$ such that for all $\delta \in I$

(a) $\delta < \beta^n(\delta)$ when $\delta < \bar{\delta}$

(b) $\delta > \beta^n(\delta)$ when $\delta > \bar{\delta}$.

A symmetric BNE that is not stable is *unstable*.

In the examples in Figure 1, it is evident that the full unraveling BNE is stable when it is unique. For larger n, we have two BNEs. The smaller BNE is stable, while the full unraveling BNE is unstable. The picture that emerges from Figure 1 holds more generally.

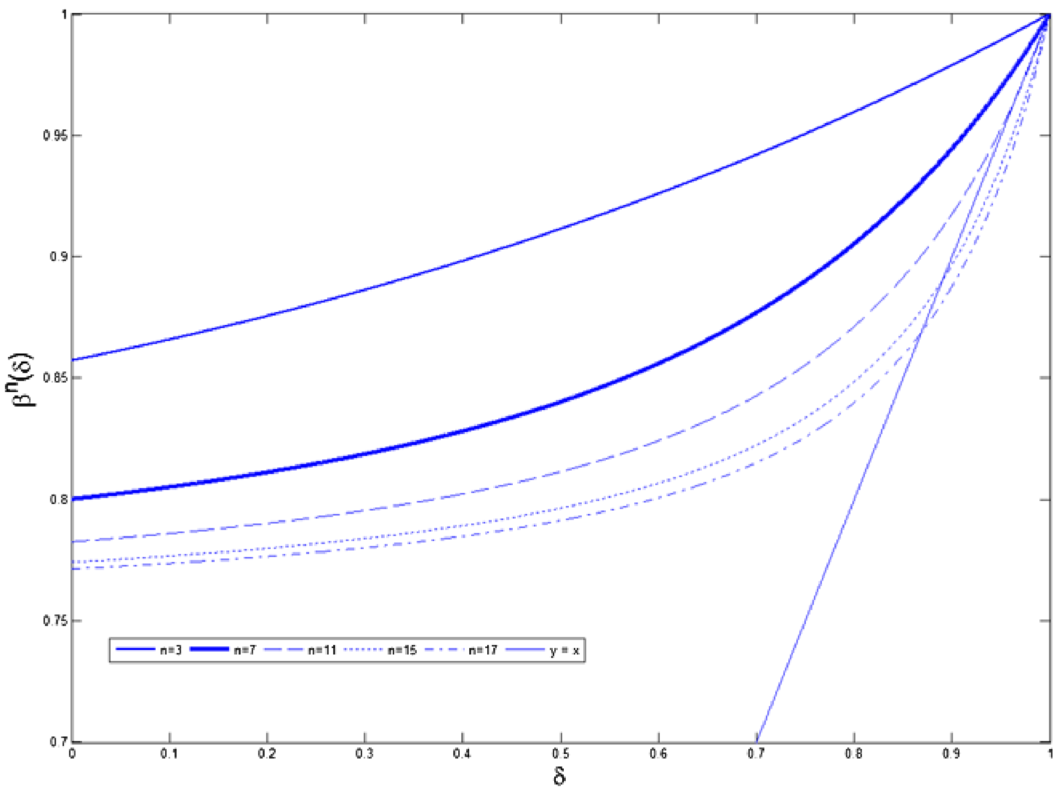

FIGURE 1. The graph of β^n for $n = 3, 7, 11, 15, 17$ and $y = x$.

REMARK 3. Let F be the uniform c.d.f. If $n \leq 10$, then the full unraveling BNE is stable. If $n > 10$, then the symmetric BNE denoted by s^n in Theorem 2 is stable while the full unraveling BNE is unstable.

3. Strategic doctors

We now assume that doctors are strategic as well. We consider the simultaneous-move game in which the players are $H \cup D$. Each agent has to decide whether to match in period $t = 0$ or $t = 1$. When doctors are strategic, the probability that m agents go early is the probability that the *minimum* between the hospitals and the doctors that make offers at period $t = 0$ equals m. Thus early matchings are rationed by the side with the shortest number of agents who go early. For the side with the most number of agents who go early, a random subset of them are indeed matched early, while the rest must stay and match late.

The set of available actions is $\{0, 1\}$ to each player. Agents' strategies are functions $s_i : [0, 1] \to \{0, 1\}$, with $i \in H \cup D$. For any profile of strategies s and any realization of types (δ_i), the number of agents who exit the market is the minimum of two quantities: the number of hospitals h with $s_h(\delta_h) = 0$ and the number of doctors d with $s_d(\delta_d) = 0$.

Thus, given a profile of strategies of *all* agents other than h, the expected value of $\mathcal{U}_m$, $\mathbf{E}_{s_{-h}}\mathcal{U}_m$, involves the probability distribution of the minimum of two independent binomial random variables instead of a single binomial random variable as in the previous case. The number m is drawn according to the minimum of two binomial distributions.

The calculations performed in the proof of Theorem 1 are still sufficient to give us the following result.

Theorem 3. *There is at least one symmetric BNE, namely the full unraveling BNE. In any symmetric BNE $s = (s_i)_{i \in H \cup D}$ for every $i \in H \cup D$, we have that*

$$\Pr(s_i = 0) \geq F\left(\frac{1}{2}\right) \geq \frac{1}{2}.$$

Corollary 2. *In any symmetric BNE, the expected number of agents who go early is at least $nF(\frac{1}{2}) \geq \frac{1}{2}n$.*

The results in Section 2.4 extend to the case when doctors are strategic. We obtain the following result.

Theorem 4. *Let F be the uniform c.d.f. If $n > 10$, then there are* at least *two symmetric BNE. One is the full unraveling BNE, which is unstable. The second is a stable BNE $s = (s_i)_{i \in H \cup D}$ such that $\Pr(s_i = 0) \geq \frac{1}{2}$ for every $i \in H \cup D$.*

Theorem 3 in fact follows from Theorem 4. So we present the proof of Theorem 4 (see Section 8) before that of Theorem 3.

4. A discussion of our model

Our model has two specific assumptions that merit some additional discussion.

Payoffs. We assume that payoffs are multiplicative, a common assumption in applied matching theory (see, e.g., Bulow and Levin 2006, Damiano et al. 2005, and many other papers). In our particular case, there are two reasons for working with multiplicative payoffs. First, a parametric assumption about payoffs is unavoidable when we are trying to precisely calculate the probability that an agent will go early. As such, the multiplicative form is natural.

The multiplicative assumption also makes sense as a way to abstract from other possible explanations of unraveling. We did not want an explanation of unraveling that was based on the insurance value of going early (an avenue explored by Li and Rosen 1998). We assume payoffs for which there is a clear advantage to going late, not early. In our model, agents are risk neutral. The multiplicative model implies that even though an agent may be concerned about a bad draw of their quality, the gains from matching assortatively outweigh the temptation to match to an average partner in $t = 0$.

One might be concerned about the robustness of the results to the specification of multiplicative payoffs. We carried out a simulation for a constant elasticity of substitution (CES) specification $f(\pi(h), \pi(d)) = (\pi(h)^\rho + \pi(d)^\rho)^{1/\rho}$ in which we assumed modest levels of complementarities for agents' decisions. The simulation uses the algorithm described in Proposition 1. The qualitative features of our results are preserved in our simulations: see the discussion at the end of Section 5 for details.

Roth (1984) suggests that unraveling is the result of a prisoners' dilemma game among the hospitals. The implication is that it is a dominant strategy for the hospitals to go early. Our focus is on the strategic channel, whereby agents go early because of their concerns that others go early (and the consequent negative externality). By our assumptions on preferences, we rule out that it is dominant for agents to go early.

It is still possible to generate unraveling by way of a coordination failure, as in the literature on bank runs (Diamond and Dybvig 1983). In our model, however, and in contrast to the model of bank runs, such unraveling is unstable. Only when all agents are *certain* that all other agents want to go early are they willing to go early. This would be an unstable situation: It is easy to rule out such an outcome if agents' beliefs may depart from certainty that everyone goes early. In contrast, we show that there is in our model a stable equilibrium in which agents are more likely to go early than to go late. Coordination failure is still present in that equilibrium, but unraveling arises through the channel of strategic unraveling.

Finally, the multiplicative model also captures the negative externality imposed by agents who go early on the rest of the market. There is an efficiency loss when some agents go early; they hurt the rest of the agents (even in a model without transfers like ours).

Information. The second assumption that deserves mention is our informational assumption. We assume that agents are completely ignorant about match qualities at date $t = 0$. The assumption is extreme, and it is meant to focus the model on the trade-off between the value of the information revealed at $t = 1$ and the incentives to go early. By assuming that there is no information at time $t = 0$ and full information at $t = 1$, we have biased the model against the unraveling outcome.

That said, it may not be an unrealistic assumption. Roth and Xing (1994) claim that "offers are being made so early that there are serious difficulties in distinguishing among the candidates." So our assumption of complete ignorance over match qualities may reflect the actual situation in the markets where we observe unraveling.

Finally, we use an assumption of the c.d.f. F that allows us to exploit the results obtained in the case when F is uniform. The assumption that $x \leq F(x)$ means that F is smaller in the sense of first-order stochastic dominance than the uniform distribution. Again, we need the uniform distribution to make precise calculations, and then the inequality on F allows us to obtain bounds. As we remarked above, the assumption on F is satisfied when F is concave.

Amplification. There is an "amplification" of incentives to going early due to the uncertainty in how many other agents go early. The reason is that the payoffs to going late are concave in the number of agents who go early (even though, as we have emphasized, agents are risk neutral in the model).

Compare the equilibrium probability that an agent goes early with the probability that she goes early if a given number m of agents go early. We take this given number of agents to equal the expected number of agents in equilibrium. The calculations are as follows. The function $x \mapsto \mathcal{U}_x = \frac{1}{6}(n+1)^2([2(n-x)+1]/(n-x+1))$ is strictly concave because its derivative is $\frac{1}{6}(n+1)^2((-1)/(n-x+1)^2)$, which is decreasing. Then Jensen's inequality means that $\mathbf{E}\mathcal{U}_m < \mathcal{U}_{\bar{m}}$, where $\bar{m} = \mathbf{E}m$ and we are taking expectations with respect to the equilibrium distribution over agents who exit the market. Therefore, the cutoff resulting from a known number $\bar{m}$ of agents who exit the market is

$$\delta_{\bar{m}} = \frac{\mathcal{U}_e}{\mathcal{U}_{\bar{m}}} < \frac{\mathcal{U}_e}{\mathbf{E}\mathcal{U}_m} = \delta^*,$$

meaning that the equilibrium cutoff δ^* is always above the cutoff that results from a known number $\bar{m}$ of agents exiting the market. The difference between these two magnitudes illustrates the amplification inherent in our model. It illustrates how uncertainty over the number of agents to exit the market induces any given agent to want to leave the market early.

5. Proof of Lemma 1

In this section, we present, in the first place, a formula for $\mathcal{U}_1$ that clarifies the meaning of this quantity. Then an algorithm to compute $\mathcal{U}_m$ in the general case is introduced (Proposition 1). Lemmas 2 and 3 deduce a simple formula for $\mathcal{U}_m$.

Recall that $\mathcal{U}_0$ is the expected utility from waiting when all other hospitals wait. Then

$$\mathcal{U}_0 = \frac{1}{n}\sum_{i=1}^{n} i^2 = \frac{(n+1)(2n+1)}{6}.$$

5.1 *Computing* $\mathcal{U}_1$

We compute the expected utility from waiting when only one pair of hospital–doctor goes early. In period 1, after permutations π^H and π^D are drawn, sets H and D can be ordered according to agents' quality. Then consider the sets H and D described as $H = \{1, 2, \dots, n\}$ and $D = \{1, 2, \dots, n\}$, where the first agent is the lowest quality agent, and the last agent is the highest quality agent.

First, conditional on being of quality i, the leaving hospital is of a higher quality than i with probability $(n-i)/(n-1)$ and is of a lower quality than i with probability $(i-1)/(n-1)$. This is deduced from the fact that there are $n-1$ possible qualities for the hospital that leaves early, $(n-i)$ of those are higher than i, and $(i-1)$ are lower than i.

If the leaving hospital is of a higher quality than i, this means that hospital i is better off unless the doctor who leaves with the hospital is also a "good" doctor—unless the doctor who leaves is one who would be matched in the second period with a hospital better than i (in which case the leaving hospital does not affect i). This happens with probability $(n-i)/n$. With the complementary probability, i/n, hospital i is better off by the better hospital leaving. Being better off means that hospital i will be matched in the second period with a doctor with a quality 1 unit higher than i (i.e., a doctor of quality $(i+1)$), which is worth i to a hospital of quality i.

If the leaving hospital is of a lower quality than i, then this does not affect hospital i and it gets i^2, unless the doctor who leaves used to be with a better hospital or with i, in which case hospital i goes down one step. To a hospital of quality i, losing one step is worth $-i$. So in the event that a hospital of lower quality than i leaves (which has probability $(i-1)/(n-1)$), it gets i^2 for sure but it loses $-i$ with probability $(n-i+1)/n$, the probability that the partner of the hospital that goes early is of a quality greater than or equal to i.

So

$$\mathcal{U}_1 = \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{n-i}{n-1}\left[i^2+\frac{i}{n}i\right]+\frac{i-1}{n-1}\left[i^2-\frac{n-i+1}{n}i\right]\right\}.$$

Since the terms that multiply i^2 add to 1, this gives

$$\mathcal{U}_1 = \frac{1}{n}\sum_{i=1}^{n}\left[i^2+\frac{i}{n(n-1)}(n-2i+1)\right] = \frac{(2n-1)(n+1)^2}{6n}.$$

Note that $\mathcal{U}_1$ can be also expressed as

$$\mathcal{U}_1 = \mathcal{U}_0 + \frac{1}{n}\sum_{i=1}^{n}\left[\frac{(n-i)}{(n-1)}\frac{i}{n}i - \frac{(i-1)}{(n-1)}\frac{(n-i+1)}{n}i\right] = \mathcal{U}_0 - \frac{n+1}{6n}.$$

The intuition behind this equation is the following. Notice that with probability $(n-i)/(n-1)$, the hospital that leaves early is of a higher quality than i and with probability $(i-1)/(n-1)$ is of a lower quality than i. Then $((n-i)/(n-1))(i/n)$ is the probability that the hospital that leaves early is of a higher quality than i and the doctor it hires is of a quality lower than or equal to i. In this event, hospital i increases its utility by i. If the hospital that goes early is of quality lower than i and it hires a doctor of quality higher than or equal to i, which happens with probability $((i-1)/(n-1))((n-i+1)/n)$, then hospital i decreases its utility by i. Therefore, $\mathcal{U}_1$ can be expressed as $\mathcal{U}_0$ plus the expected utility derived from the leaving of a pair of hospital–doctor. Moreover, $-(n+1)/(6n)$ is the negative externality imposed on the rest of the market by the first pair of hospital–doctor that decides to match early.

Clearly, this argument is very hard to generalize if we consider more than one pair of hospital–doctor that goes early. In the following section, we develop an algorithm to compute the expected utility from waiting when m pairs of hospital–doctor leave the market at $t=0$.

5.2 *An algorithm to compute* $\mathcal{U}_m$

In this section, we introduce an algorithm to compute the value of $\mathcal{U}_m$ in the general case. First, we define the payoff matrix $\mathbb{U}$ as the element (i, j) of $\mathbb{U}$ is the utility that a doctor of quality i has when she is hired by a hospital of quality j (which is also the utility of the hospital). In particular, the elements of the first column of $\mathbb{U}$ are the utilities that the hospital of quality 1 has if it hires a doctor of quality $1, 2, \ldots, n$. Note that the elements of the main diagonal of $\mathbb{U}$ are $1, 4, \dots, i^2, \dots, n^2$, which are the payoffs that each agent has when no pair of hospital–doctor leaves early. Thus, matrix $\mathbb{U}$ is

$$\begin{pmatrix} 1 & 2 & 3 & \cdots & (n-1) & n \\ 2 & 4 & 6 & \cdots & 2(n-1) & 2n \\ 3 & 6 & 9 & \cdots & 3(n-1) & 3n \\ \vdots & & & & & \\ (n-1) & 2(n-1) & 3(n-1) & \cdots & (n-1)^2 & n(n-1) \\ n & 2n & 3n & \cdots & n(n-1) & n^2 \end{pmatrix}.$$

When a hospital makes an offer at $t = 0$ and hires a doctor, both the hospital and the doctor may be of any quality. So, to compute the expected utility of a hospital that waits, we have to consider all possible qualities combinations. Assume that the hospital who leaves is of quality j and the doctor who it hires is of quality i. If only this pair of hospital–doctor leaves the market at $t = 0$, in the second period the utilities of hospitals and doctors that do not leave the market are given by the assortative matching. Indeed, the highest quality hospital (between those that remain in the market) will hire the highest quality doctor of those who do not exit the market. The same argument holds for all agents.

Therefore, when doctor i is hired at $t = 0$ by hospital j, the utilities of hospitals and doctors that remain in the market in the second period are the elements of the main diagonal of the submatrix of $\mathbb{U}$ that it is obtained from deleting the row i and the column j. To consider all possible combinations for the quality of the hospital that leaves early and the doctor who it hires, we have to go over all the elements of $\mathbb{U}$. Thus, to compute the expected utility from waiting when only one pair of hospital–doctor leaves at $t = 0$, we have to compute all the submatrices of $\mathbb{U}$ obtained by deleting one row and one column, for each one of these submatrices, we find its trace, we sum all these traces and, finally, we have to divide the sum by $n^2(n-1)$, since there are n^2 possible pairs of qualities for the hospital and the doctor that go early, and $n-1$ possible qualities that a hospital that waits may be assigned to in the second period.

If m hospitals make an offer at $t = 0$, we generalize the previous argument as follows. Consider all submatrices of $\mathbb{U}$ that result when m rows and m columns are deleted. There are $\binom{n}{m}\binom{n}{m}$ submatrices that can be found. In each case, there are $(n-m)$ possible qualities for a hospital that waits. Thus, for each submatrix, compute its trace; $\mathcal{U}_m$ is the sum of all the computed traces after dividing it by $\binom{n}{m}\binom{n}{m}(n-m) = (n^2(n-1)^2 \cdots (n-m+1)^2/(m!)^2)(n-m)$.

The following proposition states this result.[5]

Proposition 1. *Let $\mathcal{U}_m$ be the expected utility to a hospital of waiting for the second period when m hospitals (with their respective doctors) have left the market at $t = 0$. Denote by $T(n,m)$ the* sum *of the traces of all submatrices of $\mathbb{U}$ when m rows and m columns are*

[5]The algorithm can be also applied with other functions u_h and u_d whenever the functions are strictly supermodular on the lattice $\{1, 2, \dots, n\}^2$.

deleted. Then

$$\mathcal{U}_m = \frac{T(n,m)(m!)^2}{n^2(n-1)^2\cdots(n-m+1)^2}\frac{1}{(n-m)}.$$

To come up with an expression for $\mathcal{U}_m$, the next step involves the computation of $T(n,m)$. The following lemma finds a formula for $T(n,m)$. Then we obtain a reduced expression of the formula by means of some combinatorial identities.

Lemma 2. *Denote by $T(n,m)$ the sum of the traces of all submatrices of $\mathbb{U}$ obtained by deleting m rows and m columns. Then*

$$\begin{aligned} T(n,m) &\\ &= \sum_{i=1}^{n}\left[i^2\sum_{k=0}^{m}\left(\binom{i-1}{k}\binom{n-i}{m-k}\right)^2\right] \\ &\quad + 2\sum_{j=1}^{m}\left[\sum_{i=1}^{n-j} i(i+j)\left(\sum_{k=j}^{m}\binom{i+j-1}{k}\binom{n-(i+j)}{m-k}\binom{n-i}{m-k+j}\binom{i-1}{k-j}\right)\right].\end{aligned}$$

Proof. First we consider the elements of the main diagonal of $\mathbb{U}$ and then consider the remaining elements.

The (ii) elements: Consider an element ii of the matrix, and suppose we delete m rows and m columns. Note that there are $i-1$ rows (columns) above (at the left of) the element ii and $n-i$ rows (columns) below (at the right). When we delete columns and rows, the element ii remains in the main diagonal if the number of rows that are deleted above ii is equal to the number of columns that are deleted from the left of ii. That is, if we delete k rows above ii and $m-k$ rows below, then we have to delete k columns at the left and $m-k$ columns at the right. Thus, the number of submatrices in which the element ii is in the main diagonal is

$$\sum_{k=0}^{m}\left(\binom{i-1}{k}\binom{n-i}{m-k}\right)^2.$$

Since the element ii in the matrix is i^2, the share of $T(n,m)$ that corresponds to the elements of the main diagonal of $\mathbb{U}$ is

$$\sum_{i=1}^{n}\left[i^2\sum_{k=0}^{m}\left(\binom{i-1}{k}\binom{n-i}{m-k}\right)^2\right].$$

The (ij) elements: Since $\mathbb{U}$ is a symmetric matrix, the trace of the submatrix that we obtain by deleting rows $i_1, i_2, \dots, i_m$ and columns $j_1, j_2, \dots, j_m$ is equal to the trace of the submatrix obtained by deleting rows $j_1, j_2, \dots, j_m$ and columns $i_1, i_2, \dots, i_m$. Thus, we only have to consider the elements $i(i+j)$ for $j>0$ and take two times the final result. In particular, when only one row and one column are deleted, the elements that will be in the main diagonal of some submatrix are those of the form $i(i+j)$ for $i=1,\dots,n-1$ and $j=1$. When two rows and two columns are deleted, the elements to be considered in $T(n,m)$ are the previous elements and those of the form $i(i+j)$ for $i=1,\dots,n-2$ and $j=2$. In general, when m rows and m columns are deleted, we have to consider all the elements that were contemplated when $m-1$ rows and $m-1$ columns were deleted, and those of the form $i(i+j)$ for $i=1,\dots,n-m$ and $j=m$.

As we just noted, when we delete m rows and m columns, the elements that are in the trace of some submatrix are those of the form $i(i+j)$ with $j=1,2,\dots,m$. So consider an element $i(i+j)$. This element has $i-1$ rows above and $n-i$ below. Moreover, it has $i+j-1$ columns at the left and $n-(i+j)$ columns at the right. Suppose we delete k columns at the left of $i(i+j)$ and $m-(i+j)$ at the right. Now the element is in column $i+j-k$. So as to be in the main diagonal of a submatrix, it should be that $j-k\le 0$. Moreover, we have to delete $k-j$ rows above the element $i(i+j)$ to ensure that the element is in the main diagonal of a submatrix.

Then the share of $T(n,m)$ that corresponds to these elements is

$$2\sum_{j=1}^{m}\left[\sum_{i=1}^{n-j} i(i+j)\left(\sum_{k=j}^{m}\binom{i+j-1}{k}\binom{n-(i+j)}{m-k}\binom{i-1}{k-j}\binom{n-i}{m-(k-j)}\right)\right]. \qquad \square$$

Lemma 3. *For $n\in\mathbb{N}$ and $m\in 1,2,\dots,n-1$, it holds that*

$$T(n,m)=\binom{n+1}{m}^2\left(\sum_{i=1}^{n-m} i^2\right).$$

Proof. The proof was provided to us by Doron Zeilberger. It is organized in five claims.

Claim 1. *The term $T(n,m)$ can be written as*

$$\sum_{i,j,k} i(i+j)\binom{i+j-1}{k}\binom{n-(i+j)}{m-k}\binom{n-i}{m-k+j}\binom{i-1}{k-j},$$

where the summation range is over all triples (i,j,k), with the convention that the binomial coefficient $\binom{r}{s}$ is zero if it is not the case that $0\le s\le r$.

Proof. In the proof of the last lemma, we found an expression for $T(n,m)$ using the symmetry of the matrix $\mathbb{U}$. If we do not use the symmetry, we obtain the equivalent expression

$$\begin{aligned}T(n,m)=&\sum_{i=1}^{n}\left[i^2\sum_{k=0}^{m}\left(\binom{i-1}{m-k}\binom{n-i}{k}\right)^2\right]\\&+\sum_{j=1}^{m}\left[\sum_{i=1}^{n-j} i(i+j)\left(\sum_{k=j}^{m}\binom{i+j-1}{k}\binom{n-(i+j)}{m-k}\binom{n-i}{m-k+j}\binom{i-1}{k-j}\right)\right]\\&+\sum_{i=1}^{m}\left[\sum_{j=1}^{n-i} j(i+j)\left(\sum_{k=i}^{m}\binom{i+j-1}{k}\binom{n-(i+j)}{m-k}\binom{n-j}{m-k+i}\binom{j-1}{k-i}\right)\right].\end{aligned}$$

Note that for each $j = 1, \dots, m$, the range for i is $1 \le i \le n - j$, and for each $i = 1, \dots, m$, the range for j is $1 \le j \le n - i$. Thus, we can write these conditions as $1 \le i \le n$, $1 \le j \le n$, and $1 \le i + j \le n$. Now consider the sum

$$\sum_{i,j,k} i(i+j) \binom{i+j-1}{k} \binom{n-(i+j)}{m-k} \binom{n-i}{m-k+j} \binom{i-1}{k-j}.$$

The implicit range for each variable is $j \le k \le m$, $1 \le i \le n$, $1 \le j \le n$, and $1 \le i + j \le n$. This implies that both sums are equal. ◁

Claim 2. *The sum of Claim 1 equals*

$$\sum_{a=1}^{n} a \sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)} \binom{a-1}{k} \binom{n-a}{m-k} \sum_{i=a-k}^{a-k+m} i \binom{n-i}{m-k+a-i} \binom{i-1}{k-a+i}.$$

Proof. Writing $a = i + j$ (and leaving i as a discrete variable, but letting $j = a - i$), the sum of the last claim is equal to

$$\sum_{a,k,i} ia \binom{a-1}{k} \binom{n-a}{m-k} \binom{n-i}{m-k+a-i} \binom{i-1}{k-a+i}.$$

Note that summation range of each variable is defined as follows:

(a) For a, $1 \le a \le n$.

(b) For k, $0 \le k \le m$, $0 \le m - k + a - i \le n - i$, and $0 \le k \le a - 1$. This implies that $\max(0, a - (n - m)) \le k \le \min(a - 1, m)$.

(c) For i, $1 \le i \le n$, $0 \le m - k + a - i$, and $0 \le k - a + i$. This implies that $a - k \le i \le m - k + a$.

Then the last sum equals the iterated summation

$$\sum_{a=1}^{n} \sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)} \sum_{i=a-k}^{a-k+m} ia \binom{a-1}{k} \binom{n-a}{m-k} \binom{n-i}{m-k+a-i} \binom{i-1}{k-a+i},$$

which is equivalent to

$$\sum_{a=1}^{n} a \sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)} \binom{a-1}{k} \binom{n-a}{m-k} \sum_{i=a-k}^{a-k+m} i \binom{n-i}{m-k+a-i} \binom{i-1}{k-a+i}.$$ ◁

Claim 3. *The innermost sum is*

$$\sum_{i=a-k}^{a-k+m} i \binom{n-i}{m-k+a-i} \binom{i-1}{k-a+i} = (a-k) \binom{n+1}{m}.$$

Proof. First note that $i\binom{i-1}{k-a+i} = (a-k)\binom{i}{a-k}$. Then we have

$$\sum_{i=a-k}^{a-k+m} i\binom{n-i}{m-k+a-i}\binom{i-1}{k-a+i} = (a-k)\sum_{i=a-k}^{a-k+m}\binom{n-i}{m-k+a-i}\binom{i}{a-k}.$$

Now notice that

$$\sum_{i=a-k}^{a-k+m}\binom{n-i}{m-k+a-i}\binom{i}{a-k} = \sum_{i=0}^{m}\binom{n-(a-k+i)}{m-i}\binom{a-k+i}{a-k}.$$

Since $\binom{a-k+i}{a-k} = \binom{a-k+i}{i}$, the last sum can be written as

$$\sum_{i=0}^{m}\binom{n-(a-k+i)}{m-i}\binom{a-k+i}{i},$$

which is equal to

$$\sum_{i=0}^{m}\binom{(n-m-a+k)+m-i}{m-i}\binom{a-k+i}{i}.$$

Finally, we use the Vandermonde–Chu identity (Sprugnoli 2012, p. 54):

$$\sum_{k=0}^{n}\binom{x+k}{k}\binom{y+n-k}{n-k} = \binom{x+y+n+1}{n}.$$

Defining $x = a-k$ and $y = (n-m-a+k)$, we have

$$\sum_{i=0}^{m}\binom{(n-m-a+k)+m-i}{m-i}\binom{a-k+i}{i} = \binom{x+y+n+1}{n} = \binom{n+1}{m}. \qquad \triangleleft$$

Claim 4. *For the middle sum it holds that*

$$\sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)}(a-k)\binom{a-1}{k}\binom{n-a}{m-k} = a\binom{n-1}{m} - (a-1)\binom{n-2}{m-1}.$$

Proof. First, we divide the sum:

$$\sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)}(a-k)\binom{a-1}{k}\binom{n-a}{m-k}$$
$$= a\sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)}\binom{a-1}{k}\binom{n-a}{m-k} - \sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)}k\binom{a-1}{k}\binom{n-a}{m-k}.$$

We use the Vandermonde–Chu identity (Sprugnoli 2012, p. 53):

$$\sum_{k=0}^{n}\binom{x}{k}\binom{y}{n-k} = \binom{x+y}{n}.$$

The first sum is[6]

$$a \sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)} \binom{a-1}{k}\binom{n-a}{m-k} = a\binom{n-1}{m}.$$

If we replace $k\binom{a-1}{k} = (a-1)\binom{a-2}{k-1}$ in the second sum, we have

$$(a-1) \sum_{k=\max(0,a-(n-m))}^{\min(a-1,m)} \binom{a-2}{k-1}\binom{n-a}{m-k},$$

which is equal to

$$(a-1)\sum_{k=0}^{m}\binom{a-2}{m-1-k}\binom{n-a}{k}.$$

By the Vandermonde–Chu identity, the sum is

$$(a-1)\binom{n-2}{m-1}. \qquad \triangleleft$$

Claim 5. *Finally, we have*

$$T(n,m) = \binom{n+1}{m}^2\left(\sum_{i=1}^{n-m} i^2\right).$$

Proof. Since the last claims, we know that

$$T(n,m) = \binom{n+1}{m}\left(\binom{n-1}{m}\left(\sum_{a=1}^{n} a^2\right) - \binom{n-2}{m-1}\left(\sum_{a=1}^{n} a(a-1)\right)\right).$$

Then compute

$$\begin{aligned}
&\binom{n+1}{m}\left(\binom{n-1}{m}\left(\sum_{a=1}^{n} a^2\right) - \binom{n-2}{m-1}\left(\sum_{a=1}^{n} a(a-1)\right)\right)\\
&\quad= \binom{n+1}{m}\left(\frac{(n-1)!}{m!(n-m-1)!}\frac{n(n+1)(2n+1)}{6}\right.\\
&\qquad\qquad\left. - \frac{(n-2)!}{(m-1)!(n-m-1)!}\frac{(n-1)n(n+1)}{3}\right)\\
&\quad= \binom{n+1}{m}\left(\frac{(n+1)!}{m!(n-m-1)!}\frac{(2n+1)}{6} - \frac{(n+1)!}{m!(n-m-1)!}\frac{m}{3}\right)\\
&\quad= \binom{n+1}{m}\frac{(n+1)!}{m!(n-m-1)!}\left(\frac{2n+1}{6} - \frac{m}{3}\right)
\end{aligned}$$

[6]Note that $\max(0, a-(n-m)) = 0$. Indeed, if $(a-(n-m)) > 0$, we have $n-a-m-k < 0$ and, thus, $\binom{n-a}{m-k} = 0$. Also, we can write the sum up to $k = m$, because for $k = a, a+1, \dots, m$, $\binom{a-1}{k} = 0$.

$$= \binom{n+1}{m} \frac{(n+1)!}{m!(n-m-1)!} \frac{2n-2m+1}{6}$$
$$= \binom{n+1}{m} \frac{(n+1)!}{m!(n-m+1)!} \frac{(n-m)(n-m+1)(2n-2m+1)}{6}$$
$$= \binom{n+1}{m}^2 \frac{(n-m)(n-m+1)(2(n-m)+1)}{6}$$
$$= \binom{n+1}{m}^2 \sum_{i=1}^{n-m} i^2. \qquad \triangleleft$$

This completes the proof of Lemma 3. □

Finally, we obtain the formula for $\mathcal{U}_m$. We know that

$$\mathcal{U}_m = \frac{T(n,m)}{\binom{n}{m}\binom{n}{m}(n-m)}.$$

First note that

$$\binom{n+1}{m}^2 = \left[\frac{n+1}{n-m+1}\right]^2 \binom{n}{m}^2.$$

Then, by replacing the last expression in $\mathcal{U}_m$, we obtain

$$\mathcal{U}_m = \frac{(n+1)^2}{(n-m+1)^2(n-m)} \frac{(n-m)(n-m+1)(2(n-m)+1)}{6}.$$

By simplifying the last equation, we prove the result:

$$\mathcal{U}_m = \frac{(n+1)^2(2(n-m)+1)}{6(n-m+1)}.$$

Note that $\mathcal{U}_m$ increases with n, the number of agents. This means that if there are more agents in the market, the expected utility of waiting when a fixed number of agents leave the market at $t = 0$, increases.

The next result shows that $\mathcal{U}_m$ decreases with m, a property that will be used in the next section. Then the expected utility of waiting and match at $t = 1$ decreases as more agents leave early.

Corollary 3. *Let $\mathcal{U}_m$ be expected utility of a hospital that decides to wait for the second period when m pairs of hospital–doctor leave the market at $t = 0$. Then for $n \in \mathbb{N}$ and $m = 0, 1, 2, \ldots, n-1$, we have*

$$\mathcal{U}_m - \mathcal{U}_{m+1} = \frac{(n+1)^2}{6(n-m)(n-m+1)}.$$

Note that $\mathcal{U}_{m+1} - \mathcal{U}_m$ represents the negative externality imposed on the rest of the market by one pair of hospital–doctor that decides to go early, when m agents

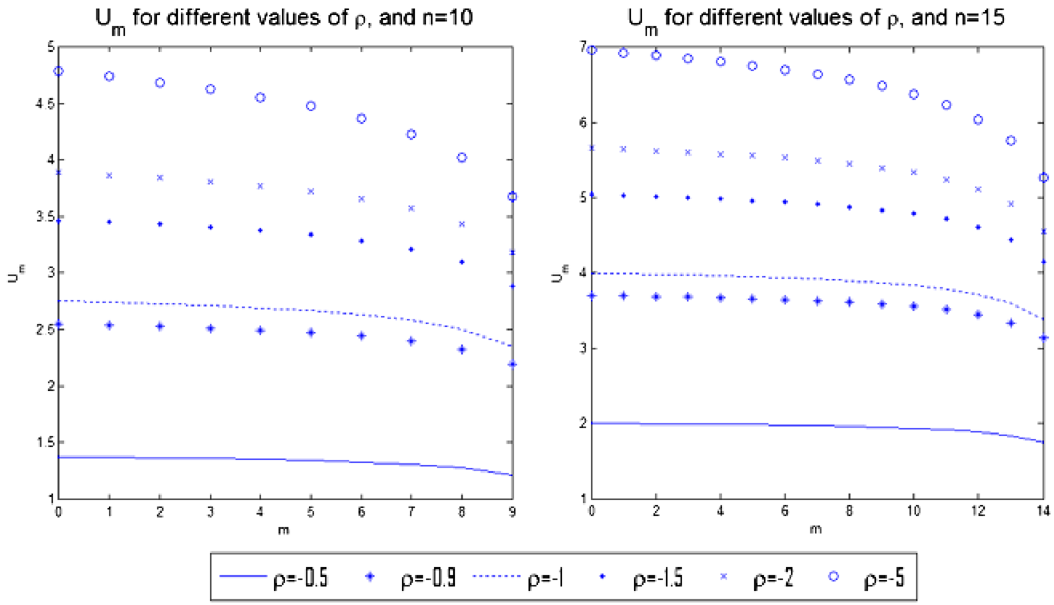

FIGURE 2. The $\mathcal{U}_m$ for different degrees of complementarity of payoffs.

have already decided to match at $t = 0$. Since $\mathcal{U}_m - \mathcal{U}_{m+1}$ increases when m becomes larger, the negative externality imposed by one more pair going early increases (in absolute value) as more agents have decided to go early. Moreover, when the number of agents (that is, n) increases, the negative externality decreases. However, since $\lim_{n\to\infty} \mathcal{U}_{m+1} - \mathcal{U}_m = \frac{1}{6}$, it does not converge to zero as the market size goes to infinity. Thus, the negative externality becomes neutral when n tends to infinity because it does not depend on the number of agents who have previously decided to match early.

Finally, we use the algorithm described in Proposition 1 to study whether even modest degrees of complementarity of payoffs would continue to imply that $\mathcal{U}_m$ is decreasing in m. Specifically, we assume a CES specification for the payoffs

$$f(\pi(h), \pi(d)) = (\pi(h)^\rho + \pi(d)^\rho)^{1/\rho},$$

with $\rho < 1$, and we compute $\mathcal{U}_m$ with $m \in \{0, \dots, n-1\}$ for $n = 10$ and $n = 15$, and for different values of ρ. As Figure 2 shows, $\mathcal{U}_m$ is decreasing in m for all the studied values of ρ, which means that our result holds even for small degrees of complementarity.

6. Proof of Theorem 2

Recall that the best-response function of the game, β^n, is defined by (1) in the following way. Given identical thresholds $\delta_{-h} = \delta$ for all hospitals other than h, $\beta^n(\delta)$ is given by the equation

$$\frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n} ij = \beta^n(\delta)\mathbf{E}_{s_{-h}}\mathcal{U}_m,$$

where s_{-h} is such that $s_{\tilde{h}} = 0$ if $\delta_{\tilde{h}} \leq \delta$ and $s_{\tilde{h}} = 1$ if $\delta_{\tilde{h}} > \delta$ for all $\tilde{h} \neq h$.

Note that

$$\frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n} ij = \frac{(n+1)^2}{4}.$$

When all hospitals other than h have the same threshold δ, the probability that m hospitals make early offers is the probability that m hospitals have discount factors less than or equal to δ, and that $n-m$ hospitals have discount factors higher than δ. Since discount factors are drawn independently from a uniform distribution on $[0,1]$, the probability that m hospitals leave at $t=0$ is given by $\delta^m(1-\delta)^{n-1-m}\binom{n-1}{m}$. Therefore,

$$\mathbf{E}_{s_{-h}}\mathcal{U}_m = \sum_{m=0}^{n-1}\delta^m(1-\delta)^{n-1-m}\binom{n-1}{m}\mathcal{U}_m.$$

Then β^n is defined by

$$\beta^n(\delta) = \frac{(n+1)^2}{4\big[\sum_{m=0}^{n-1}\delta^m(1-\delta)^{n-1-m}\binom{n-1}{m}\mathcal{U}_m\big]}.$$

The symmetric BNE of our game are the fixed points of the best-response function β^n. Since we know that $\mathcal{U}_{n-1} = \frac{1}{4}(n+1)^2$ from Lemma 1, then $\beta^n(1) = 1$ for all n. Thus, full unraveling is a BNE for all n. In this section, we investigate the existence of other fixed points. In particular, Lemma 4 gives a simple formula for β^n. Lemma 6 shows that β^n is a convex and increasing function of δ and $\beta^n(0) > \frac{3}{4}$. Thus, β^n may have, at most, one fixed point different from $\delta = 1$. Moreover, if it exists, the fixed point is higher than $\frac{3}{4}$. Lemma 6 proves that $\delta = 1$ is the unique fixed point of β^n for all $n \leq 10$, and if $n > 10$, β^n has exactly two fixed points. Finally, Lemma 8 studies the behavior of β^n when n tends to infinity.

It is worthwhile noting that the threshold at the BNE s^n defined in Theorem 2 decreases as more agents are present in the market. This means that the probability that a hospital makes early offers decreases as the number of agents increases. The intuition of this result is straightforward since, as we noted before, the incentives to make early offers when a fixed number of agents leave the market at $t=0$ decreases with n.

Lemma 4. *We have*

$$\beta^n(\delta) = \frac{3}{2}\left(2 - \frac{1}{n(n+1)}\sum_{m=1}^{n} m\delta^{n-m}\right)^{-1}.$$

First we will prove the following lemma, which will be useful in the proof of Lemma 4.

Lemma 5. *For any $n \in \mathbb{N}$ and $\delta \in \mathbb{R}$, it holds that*

$$\sum_{m=0}^{n}\frac{(1-\delta)^{n-m}\delta^m\binom{n}{m}}{n-m+2} = \sum_{m=0}^{n}\frac{(m+1)\delta^{n-m}}{(n+1)(n+2)}.$$

Proof.[7] Consider the following polynomials of degree n:

$$p(\delta) = \sum_{m=0}^{n} \frac{(1-\delta)^{n-m}\delta^m \binom{n}{m}}{n-m+2}$$

$$q(\delta) = \sum_{m=0}^{n} \frac{(m+1)\delta^{n-m}}{(n+1)(n+2)}.$$

We want to prove that $p = q$, and to this end, we will show that all the derivatives of p and q are equal at $\delta = 0$. Denote by $p^{(k)}$ and $q^{(k)}$ the kth derivative of p and q, respectively. It is straightforward to show that

$$q^{(k)}(\delta) = \sum_{m=0}^{n-k} \frac{1}{(n+1)(n+2)}(m+1)(n-m)(n-m-1)\cdots(n-m-k+1)\delta^{n-m-k}$$

for $k = 1, 2, \dots, n$.

Then

$$q^{(k)}(\delta) = \sum_{m=0}^{n-k} \frac{1}{(n+1)(n+2)}(m+1)\frac{(n-m)!}{(n-m-k)!}\delta^{n-m-k}.$$

When we evaluate at $\delta = 0$, we have

$$q^{(k)}(0) = \frac{(n-k+1)k!}{(n+1)(n+2)}.$$

To compute the kth derivative of p, consider the functions

$$g_1(\delta) = (1-\delta)^{n-m}$$

$$g_2(\delta) = \delta^m.$$

Then

$$g_1^{(i)}(\delta) = \frac{(n-m)!}{(n-m-i)!}(-1)^i(1-\delta)^{n-m-i}$$

$$g_2^{(k-i)}(\delta) = \frac{m!}{(m-k+i)!}\delta^{m-(k-i)}.$$

By the general Leibniz rule, we have, for $k = 1, 2, \dots, n$,

$$(g_1g_2)^{(k)}(\delta) = \sum_{i=0}^{k} \binom{k}{i} \frac{(n-m)!}{(n-m-i)!}\frac{m!}{(m-k+i)!}(-1)^i(1-\delta)^{n-m-i}\delta^{m-(k-i)}.$$

If $m-k \geq 0$, then $m-(k-i) \geq 0$ for all i and, thus, $(g_1g_2)^{(k)}(0) = 0$ Then supposing $m-k \leq 0$, we have

$$(g_1g_2)^{(k)}(0) = \binom{k}{k-m}\frac{(n-m)!}{(n-k)!}m!(-1)^{k-m}.$$

[7] We are very grateful to Andrés Sambarino for helpful comments on this proof.

Thus, the kth derivative of p is

$$p^{(k)}(0)=\sum_{m=0}^{n}\binom{n}{m}\frac{1}{(n+2-m)}\binom{k}{k-m}\frac{(n-m)!}{(n-k)!}(-1)^{k-m}m!.$$

As we just noted, when $m\geq k$, $p^{(k)}(0)=0$, then we can write the previous sum from $m=0$ to $m=k$.

We want to prove that $p^{(k)}(0)=q^{(k)}(0)$ for all $k=1,2,\dots,n$; that is,

$$\sum_{m=0}^{k}\binom{n}{m}\frac{1}{(n+2-m)}\binom{k}{k-m}\frac{(n-m)!}{(n-k)!}(-1)^{k-m}m!=\frac{(n-k+1)!k!}{(n+1)(n+2)}.$$

Note that

$$\binom{n}{m}\frac{(n+1)(n+2)}{(n+m-2)}=\binom{n+2}{m}(n+1-m)$$

$$\binom{k}{k-m}\frac{(n-m)!}{(n-k)!}m!\frac{1}{(n-k+1)k!}(n+1-m)=\binom{n+1-m}{n-k+1}.$$

Thus, we have to prove that

$$(-1)^k\sum_{m=0}^{k}\binom{n+2}{m}\binom{n+1-m}{n-k+1}(-1)^m=1.$$

To finish the proof, we use the binomial identity (Riordan 1968, p. 8)

$$\sum_{k=0}^{n}(-1)^k\binom{n}{k}\binom{x-k}{r}=\binom{x-n}{r-n}=\binom{x-n}{x-r}.$$

Thus,

$$(-1)^k\sum_{m=0}^{k}\binom{n+2}{m}\binom{n+1-m}{n-k+1}(-1)^m=(-1)^k\binom{-1}{k}.$$

Finally, by the negation rule, we have $\binom{-1}{k}=(-1)^k\binom{1+k-1}{k}=(-1)^k$, and then

$$(-1)^k\sum_{m=0}^{k}\binom{n+2}{m}\binom{n+1-m}{n-k+1}(-1)^m=(-1)^{2k}=1.$$

□

Proof of Lemma 4. We know that

$$\beta^n(\delta)=\frac{(n+1)^2}{4\left[\sum_{m=0}^{n-1}(1-\delta)^{n-1-m}\delta^m\binom{n-1}{m}\mathcal{U}_m\right]}$$

$$\mathcal{U}_m=\frac{(n+1)^2(2(n-m)+1)}{6(n-m+1)}$$

for $m = 0, \dots, n-1$.

We will use the two identities

$$\frac{2(n-m)+1}{n-m+1} = 2 - \frac{1}{n-m+1},$$

$$\sum_{m=0}^{n-1} (1-\delta)^{n-1-m} \delta^m \binom{n-1}{m} = 1.$$

Then, by substituting $\mathcal{U}_m$ and since the last identities, we have

$$\beta^n(\delta) = \frac{3}{2\Big(2 - \sum_{m=0}^{n-1} \frac{(1-\delta)^{n-1-m}\delta^m \binom{n-1}{m}}{n-m+1}\Big)}.$$

By the previous lemma, we can write β^n as

$$\beta^n(\delta) = \frac{3}{2\big(2 - \sum_{m=0}^{n-1} \frac{(m+1)}{n(n+1)} \delta^{n-1-m}\big)},$$

which is equivalent to

$$\beta^n(\delta) = \frac{3}{2\big(2 - \frac{1}{n(n+1)} \sum_{m=1}^{n} m\delta^{n-m}\big)}.$$

□

The following lemma gives more information on the nature of β^n.

Lemma 6. *We have*

$$\beta^n(\delta) = \begin{cases} \frac{3}{2[2 - \frac{1}{n(n+1)}(\frac{\delta^{n+1} - (n+1)\delta + n}{(1-\delta)^2})]} & \text{if } \delta \in (0,1) \\ 1 & \text{if } \delta = 1. \end{cases}$$

Further, the following statements hold:

(a) *The function β^n is increasing for each n.*

(b) *The function $\beta^n(0) > \frac{3}{4}$ and $\beta^n(1) = 1$ for all n.*

(c) *The function β^n is convex.*

(d) *The function β^n has, at most, two fixed points: $\delta = 1$ is a fixed point of β^n for all $n \in \mathbb{N}$ and it may have another fixed point that, if it exists, is higher than $\frac{3}{4}$.*

Proof. When $\delta = 1$, we have

$$\beta^n(1) = \frac{3}{2\big(2 - \frac{1}{n(n+1)} \sum_{m=1}^{n} m\big)} = \frac{3}{2\big(2 - \frac{1}{n(n+1)} \frac{n(n+1)}{2}\big)} = 1.$$

Then suppose $\delta \in (0,1)$ and note that

$$\sum_{m=1}^{n} m\delta^{n-m} = \delta^n \sum_{m=1}^{n} m\left(\frac{1}{\delta}\right)^m.$$

We use the following identity, which holds for $x \neq 1$:

$$\sum_{m=1}^{n} m x^m = \frac{x(1-(n+1)x^n + n x^{n+1})}{(1-x)^2}.$$

Finally, for $\delta \in (0,1)$, we have

$$\sum_{m=1}^{n} m \delta^{n-m} = \frac{\delta^{n+1} - (n+1)\delta + n}{(1-\delta)^2}.$$

To prove part (a), note that for all $\delta \in [0,1]$,

$$\left(\sum_{m=1}^{n} m \delta^{n-m}\right)' = \sum_{m=1}^{n-1} m(n-m)\delta^{n-m-1} \geq 0.$$

Then the expression $\sum_{m=1}^{n} m\delta^{n-m}$ increases with δ and, thus, β^n is increasing.

For part (b), notice that $\beta^n(0) = (3(n+1))/(2(2n+1)) > \frac{3}{4}$ for all $n \geq 1$.

To prove part (c), note that $\beta^n(\delta) = f(g(\delta))$, where $g(\delta) = \sum_{m=1}^{n} m\delta^{n-m}$ and $f(x) = 3/(2(2-x/(n(n+1))))$. Since f is an increasing and convex function, and g is a convex function, then β^n is a convex function.

To prove part (d), we know that $\delta = 1$ is a fixed point of β^n. Since $\beta^n(0) > \frac{3}{4}$, β^n is convex, and β^n is increasing, β^n crosses the line $y = x$ at, at most, one point different from $\delta = 1$. Thus, if it exists, the second fixed point is higher than $\frac{3}{4}$. □

Lemma 7. *Consider the best-response function β^n. Then the following statements hold:*

(a) For each $\delta \in [0,1]$, it holds that $\beta^n(\delta) \geq \beta^{n+1}(\delta)$.

(b) For all $n \leq 10$, $\delta = 1$ is the unique fixed point of β^n.

(c) For all $n > 10$, β^n has two and only two fixed points.

Proof. (a) We will show that

$$\frac{\delta^{n+1} - (n+1)\delta + n}{n(n+1)} \geq \frac{\delta^{n+2} - (n+2)\delta + (n+1)}{(n+1)(n+2)}$$

or, equivalently,

$$\frac{\delta^{n+1}}{n} - \frac{\delta}{n} - \frac{\delta^{n+2}}{n+2} + \frac{1}{n+2} \geq 0.$$

Define

$$h^n(\delta) = \frac{\delta^{n+1}}{n} - \frac{\delta}{n} - \frac{\delta^{n+2}}{n+2} + \frac{1}{n+2}.$$

It is straightforward to prove that h^n is such that $h^n(0) = 1/(n+2)$, $h^n(1) = 0$, and h^n is decreasing in $[0,1)$. So, we have that $h^n(\delta) \geq 0$.

For (b) and (c), we have to study the solutions in [0, 1] of the equation

$$\beta^n(\delta) = \frac{3}{2\left(2 - \frac{1}{n(n+1)} \sum_{m=1}^{n} m\delta^{n-m}\right)} = \delta.$$

Since for each $\delta \in [0,1]$ it holds that $\beta^n(\delta) \geq \beta^{n+1}(\delta)$, if β^n has two fixed points for some n_0, then β^n has two fixed points for all n such that $n \geq n_0$. We know that $\delta = 1$ is one solution of the equation and there may be, at most, one more solution in [0, 1]. The equation is equivalent to

$$p_n(\delta) = \delta^n + 2\delta^{n-1} + 3\delta^{n-2} + \cdots + (n-1)\delta^2 + (-2n^2 - n)\delta + \frac{3n(n+1)}{2} = 0.$$

As we noted, $\delta = 1$ is a root of p_n. We also know that $p_n(0) = 3n(n+1)/2 > 0$ and that p_n has, at most, one more root. Then we will prove that for some n_0, $p'_{n_0}(1) > 0$, which implies that for all $n \geq n_0$, p_n has two fixed points in [0, 1]. Then compute

$$p'_n(1) = \left[\sum_{i=1}^{n-1} i(n-i+1)\right] + (-2n^2 - n) = \frac{n(n+1)(n-10)}{6}.$$

Thus, for all n such that $0 \leq n \leq 10$, $p'_n(1) \leq 0$ and, for all $n > 10$, $p'_n(1) > 0$. This finishes the proof. □

6.1 *Behavior as $n \to \infty$*

Lemma 8. *For each $\delta \in [0,1]$,*

$$\lim_{n\to\infty} \beta^n(\delta) = \begin{cases} \frac{3}{4} & \text{if } \delta \in [0,1) \\ 1 & \text{if } \delta = 1. \end{cases}$$

Proof. For $\delta = 1$ we know that $\beta^n(1) = 1$ for all n. Assume $\delta < 1$. Then, by Lemma 6, it is enough to show that

$$\lim_{n\to+\infty} \frac{1}{n(n+1)}\left[\frac{\delta^{n+1} - (n+1)\delta + n}{(1-\delta)^2}\right] = 0.$$

The last expression is equivalent to

$$\frac{1}{(1-\delta)^2}\left[\frac{\delta^{n+1}}{n+1} - \frac{\delta}{n} + \frac{1}{n+1}\right].$$

Finally, it is straightforward to show that the limit of the last expression when n tends to infinity is zero. □

Note that Lemma 8 implies that the best-response function β^n converges to a discontinuous function as $n \to \infty$.

Finally, note that in any symmetric BNE, the expected number of hospitals that go early is given by

$$\sum_{m=0}^{n} m(1-\delta^*)^{n-m}(\delta^*)^m \binom{n}{m},$$

where δ^* is a fixed point of β^n.

The last expression equals $n\delta^*$. As we noted before, β^n has, at most, two fixed points, each one higher than $\frac{3}{4}$. Thus, in any symmetric BNE, the expected number of hospitals that go early is at least $\frac{3}{4}n$.

7. Proof of Theorem 1

When all agents share the prior that different δ_h are drawn independently from a distribution over $[0, 1]$ with c.d.f. F, the best-response function is given by $F(\beta^n(x))$. Since β^n is an increasing function and $F(x) \geq x$, we have that $\beta^n(F(x)) \geq \beta^n(x)$. Finally, note that $F(1) = 1$ and that $\beta^n(1) = 1$. Then Theorem 1 follows directly from Theorem 2.

8. Proof of Theorem 4

Theorem 3 follows from Theorem 4, so we present the proof of Theorem 4 before that of Theorem 3.

When both sides of the market are strategic, the game is analyzed in the same way as in the previous sections. The difference is that now the probability that m agents leave early is the probability that the minimum between the hospitals and the doctors that play at $t = 0$ equals m. Then the expected value of $\mathcal{U}_m$ involves the probability distribution of the minimum of two independent binomial random variables, one with parameters $(\delta, n-1)$ and the other with parameters (δ, n).[8]

We introduce some additional notation. Let x_m be the probability that a binomial random variable with parameters $(\delta, n-1)$ equals m, let y_m be the probability that a binomial random variable with parameters (δ, n) equals m, and let h_m be the probability that the minimum of two independent such random variables equals m. Denote by G and H the cumulative distribution function of a binomial random variable with parameters $(\delta, n-1)$ and (δ, n), respectively, and let $\bar{G} = 1 - G$ and $\bar{H} = 1 - H$.

Therefore, the best-response function is defined by

$$\tilde{\beta}^n(\delta) = \frac{(n+1)^2}{4[\sum_{m=0}^{n-1} h_m \mathcal{U}_m]}.$$

We use the results of the previous sections to find a lower and upper bound for $\tilde{\beta}^n$. We first prove some properties of $\tilde{\beta}^n$. In particular, Lemma 9 shows that $\delta = 1$ is a fixed

[8]Note that the best-response function is defined from the perspective of a single player. Assume it is a hospital. So the probability that $\mathcal{U}_m$ happens is the probability that the minimum between a binomial random variable with parameters $(\delta, n-1)$, representing the $(n-1)$ hospitals, and an independent binomial random variable with parameters (δ, n), representing all n doctors, equals m.

point of $\tilde{\beta}^n$ for all n, that $\tilde{\beta}^n$ is an increasing function of δ, and that $\tilde{\beta}^n(\delta) \geq \tilde{\beta}^{n+1}(\delta)$ for all n. Lemma 10 demonstrates that for each $\epsilon > 0$, there exists n_0 such that for all $n \geq n_0$ and $\delta \in [0, 1]$, it holds that

$$\beta^n(\delta) \geq \tilde{\beta}^n(\delta) \geq \frac{\frac{3}{4}}{\left(\frac{3}{2} + \epsilon\right)}.$$

Then $\lim_{n\to\infty} \tilde{\beta}^n(\delta) \geq \frac{1}{2}$, and since $\tilde{\beta}$ decreases when n increases, we conclude that for all n,

$$\beta^n(\delta) \geq \tilde{\beta}^n(\delta) \geq \tfrac{1}{2}.$$

Finally, given that β^n has two fixed points when $n > 10$, we can conclude that $\tilde{\beta}^n$ has, at least, two fixed points for $n > 10$.

Lemma 9. *Consider the best-response function $\tilde{\beta}^n$ as defined before. Then the following statements hold:*

(a) We have $\tilde{\beta}^n(1) = 1$ for all n.

(b) The function $\tilde{\beta}^n$ is an increasing function of δ.

(c) For each $\delta \in [0, 1]$, $\tilde{\beta}^n(\delta) \geq \tilde{\beta}^{n+1}(\delta)$ for all n.

Proof. (a) Given that the cumulative distribution function of the minimum of two independent random variables with c.d.f. G and H is $1 - (1 - G)(1 - H)$, we have

$$\begin{aligned} h_m &= 1 - (1 - G(m))(1 - H(m)) - \big(1 - (1 - G(m-1))(1 - H(m-1))\big) \\ &= H(m) - H(m-1) + G(m) - G(m-1) + H(m-1)G(m-1) - H(m)G(m) \\ &= x_m + y_m + H(m-1)(G(m-1) - G(m)) + G(m)(H(m-1) - H(m)) \\ &= x_m \bar{H}(m-1) + y_m \bar{G}(m). \end{aligned}$$

Thus,

$$\tilde{\beta}^n(\delta) = \frac{(n+1)^2}{4[\sum_{m=0}^{n-1}(x_m \bar{H}(m-1) + y_m \bar{G}(m))\mathcal{U}_m]}.$$

When we compute $\tilde{\beta}^n(1)$, we obtain

$$\tilde{\beta}^n(1) = \frac{(n+1)^2}{4[(\bar{H}(n-2))\mathcal{U}_{n-1}]}.$$

Given that $\mathcal{U}_{n-1} = \frac{1}{4}(n-1)^2$ and, for $\delta = 1$, $\bar{H}(n-2) = 1$, we have that $\tilde{\beta}^n(1) = 1$.

(b) Now let $\hat{G}$ and $\hat{H}$ be the cumulative distribution function of a binomial random variable with parameters $(\hat{\delta}, n-1)$ and $(\hat{\delta}, n)$, respectively, with $\hat{\delta} > \delta$. We know that

$\hat{G}(m) \leq G(m)$ and $\hat{H}(m) \leq H(m)$ for $m \in \{0, \ldots, n-1\}$; then $1-(1-\hat{G}(m))(1-\hat{H}(m)) \leq 1-(1-G(m))(1-H(m))$. Let $\hat{h}_m$ be the probability that the minimum of two independent binomial random variables, one with parameters $(\hat{\delta}, n-1)$ and the other $(\hat{\delta}, n)$, equals m. Then, since $\mathcal{U}_m$ decreases with m, we have that

$$\sum_{m=0}^{n-1} \hat{h}_m \mathcal{U}_m \leq \sum_{m=0}^{n-1} h_m \mathcal{U}_m.$$

Therefore, $\tilde{\beta}^n$ is an increasing function of δ.

(c) We know that $\mathcal{U}_m = (n+1)^2(2(n-m)+1)/(6(n-m+1))$. Then the best-response function can be written as

$$\tilde{\beta}^n(\delta) = \frac{3}{2\left[1+\sum_{m=0}^{n-1} \frac{n-m}{n-m+1} h_m\right]}.$$

Using a change of variable, $k = n-m$, we obtain

$$\sum_{m=0}^{n-1} \frac{n-m}{n-m+1} h_m = \sum_{k=1}^{n} \frac{k}{k+1} h_{n-k} = \sum_{k=0}^{n} \frac{k}{k+1} h_{n-k}.$$

Consider two binomial random variables X_1^n and X_2^{n+1}. Each random variable is defined on the same sample space, i.e., the space of an infinite number of Bernoulli trials. For X_1^n and X_2^{n+1}, we count the number of successes in the first n and $n+1$ such trials, respectively. The sample spaces for X_1^n and X_2^{n+1} are independent.

Now, for each n, there are also the random variables Y_1^n and Y_2^{n+1} counting the number of failures. Note that $X_1^n + Y_1^n = n$ and $X_2^{n+1} + Y_2^{n+1} = n+1$.

Let r_k be the probability that $\max\{Y_1^n, Y_2^{n+1} - 1\} = k$. Observe that $h_{n-k} = r_k$. So we have that

$$\sum_{m=0}^{n-1} \frac{n-m}{n-m+1} h_m = \sum_{k=0}^{n} \frac{k}{k+1} r_k.$$

Since we have defined these random variables on the same sample space, it is true that

$$\{Y_1^n \geq x\} \subseteq \{Y_1^{n+1} \geq x\}$$

and

$$\{Y_2^{n+1} - 1 \geq x\} \subseteq \{Y_2^{n+2} - 1 \geq x\}$$

for any x, because any time that we have at least x failures in the first n or $n+1$ Bernoulli trials, we have at least x failures in the first $n+1$ or $n+2$ Bernoulli trials (past failures cannot be undone).

By the same token,

$$\{\max\{Y_1^n, Y_2^{n+1} - 1\} \geq x\} \subseteq \{\max\{Y_1^{n+1}, Y_2^{n+2} - 1\} \geq x\},$$

so that the probability distribution (r_k) increases in the sense of first-order stochastic dominance (it actually increases in a stronger sense).

The function $k \mapsto k/(k+1)$ is monotone increasing. Thus, the sum

$$\sum_{k=0}^{n} \frac{k}{k+1} r_k$$

is increasing in n, as it is the expected value of a monotone increasing function and the probability law is monotone increasing in n. $\square$

Consider the function

$$\alpha^n(\delta) = \frac{(n+1)^2}{4[\sum_{m=0}^{n-1} j_m \mathcal{U}_m]},$$

where j_m is the probability that the minimum of two binomial independent random variables with parameters $(\delta, n-1)$ equals m. We will prove that there exists n_0 such that for all $n \geq n_0$, we have $\alpha^n(\delta) \geq \frac{3}{4}/(\frac{3}{2}+\epsilon)$. Then, after showing that $\tilde{\beta}^n(\delta) \geq \alpha^n(\delta)$ for all δ, we deduce a lower bound of the best-response function.

Lemma 10. *Let $\epsilon > 0$. Then there exists n_0 such that for all $n \geq n_0$, the function α^n defined previously satisfies*

$$\alpha^n(\delta) \geq \frac{\frac{3}{4}}{\left(\frac{3}{2}+\epsilon\right)}.$$

Proof. Since the cumulative distribution function of the minimum of two binomial independent random variables with parameters $(\delta, n-1)$ is $1-(1-G)^2$, we have

$$\begin{aligned} j_m &= \left(1-(1-G(m))^2\right) - \left(1-(1-G(m-1))^2\right) \\ &= (1-G(m-1))^2 - (1-G(m))^2 \\ &= 2(G(m)-G(m-1)) + G(m-1)^2 - G(m)^2 \\ &= 2x_m + (G(m-1)-G(m))(G(m)+G(m-1)) \\ &= x_m(2-G(m-1)-G(m)) \\ &= x_m(\bar{G}(m-1)+\bar{G}(m)) \\ &\leq 2x_m \bar{G}(m-1). \end{aligned}$$

Then

$$\sum_{m=0}^{n-1} j_m \mathcal{U}_m = \sum_{m=0}^{n-1} x_m(\bar{G}(m-1)+\bar{G}(m))\mathcal{U}_m \leq \sum_{m=0}^{n-1} 2x_m \bar{G}(m-1)\mathcal{U}_m.$$

The median of a binomial distribution with parameter $(\delta, n-1)$ lies within the interval $[\lfloor (n-1)\delta \rfloor, \lceil (n-1)\delta \rceil]$. Moreover, if $(n-1)\delta$ is an integer, the median is $(n-1)\delta$. So if $(n-1)\delta$ is an integer, we have that $\bar{G}((n-1)\delta) = \Pr[x_m \geq (n-1)\delta + 1] \leq \frac{1}{2}$. Otherwise, if

$(n-1)\delta$ is not an integer, $\bar{G}(\lfloor(n-1)\delta\rfloor) = \Pr[x_m > \lfloor(n-1)\delta\rfloor] = \Pr[x_m \geq \lceil(n-1)\delta\rceil] \leq \frac{1}{2}$. Thus, if $m \geq \lfloor(n-1)\delta\rfloor + 1$, we have that $\bar{G}(m-1) \leq \bar{G}(\lfloor(n-1)\delta\rfloor) \leq \frac{1}{2}$. Then

$$\begin{aligned}
\sum_{m=0}^{n-1} j_m \mathcal{U}_m &\leq 2\left[\sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \mathcal{U}_m x_m \bar{G}(m-1) + \sum_{m=\lfloor(n-1)\delta\rfloor+1}^{n-1} \mathcal{U}_m x_m \bar{G}(m-1)\right] \\
&\leq 2\left[\sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \mathcal{U}_m x_m + \frac{1}{2} \sum_{m=\lfloor(n-1)\delta\rfloor+1}^{n-1} \mathcal{U}_m x_m\right] \\
&= \sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \mathcal{U}_m x_m + \sum_{m=0}^{n-1} \mathcal{U}_m x_m = g(\delta) + \sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \mathcal{U}_m x_m,
\end{aligned}$$

where $g(\delta) = \sum_{m=0}^{n-1} \mathcal{U}_m x_m$.

Now recall that

$$\mathcal{U}_m = \frac{(n+1)^2(2(n-m)+1)}{6(n-m+1)}.$$

So we obtain that

$$\begin{aligned}
\sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \mathcal{U}_m x_m &= \frac{(n+1)^2}{6} \sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \frac{(2(n-m)+1)}{n-m+1} x_m \\
&= \frac{(n+1)^2}{6} \sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \left(1 + \frac{(n-m)}{n-m+1}\right) x_m \\
&\leq \frac{(n+1)^2}{6} \sum_{m=0}^{\lfloor(n-1)\delta\rfloor} 2x_m \\
&\leq \frac{(n+1)^2}{6},
\end{aligned}$$

where, in the last inequality, we use that $G(\lfloor(n-1)\delta\rfloor) \leq \frac{1}{2}$.

Therefore,

$$\frac{\sum_{m=0}^{\lfloor(n-1)\delta\rfloor} \mathcal{U}_m x_m}{g(\delta)} \leq \frac{\frac{(n+1)^2}{6}}{g(\delta)} = \left(\sum_{m=0}^{n-1} \frac{2(n-m)+1}{n-m+1} x_m\right)^{-1}. \tag{2}$$

Now let $\epsilon > 0$. Choose $\rho_0, \rho_1 \in (0,1)$ such that[9]

$$\frac{1}{1+\rho_0\rho_1} < \frac{1}{2} + \epsilon.$$

[9]Note that ρ_0 and ρ_1 exist since $l(x) = 1/(1+x)$ is a continuous and decreasing function in $[0,1]$ with $l(0) = 1$ and $l(1) = \frac{1}{2}$.

Let n be large enough such that

$$\Pr\left(\tilde{M} \leq n - \frac{\rho_0}{1-\rho_0}\right) \geq \rho_1,$$

where $\tilde{M}$ is a binomial random variable with parameters $(n-1, \delta)$.

Clearly, the value of n that satisfies the last inequality depends on δ. Moreover, for higher values of δ, we need to consider higher values of n. Then assume that $\delta \leq \frac{1}{2}$ and take n large enough such that the inequality holds. In the last step of the proof, we extend the result for all values of δ.

Now $m \leq n - \rho_0/(1-\rho_0)$ if and only if $\rho_0 \leq (1-\rho_0)(n-m)$ if and only if

$$\rho_0 \leq \frac{n-m}{n-m+1}.$$

Noting that $\sum_{m=0}^{n-1}(2(n-m)+1)/(n-m+1)x_m$ is the expectation of the random variable

$$\left(\frac{2(n-\tilde{M})+1}{n-\tilde{M}+1}\right),$$

then we have

$$\sum_{m=0}^{n-1} \frac{2(n-m)+1}{n-m+1} x_m = \mathbb{E}_{\tilde{M}}\left(\frac{2(n-\tilde{M})+1}{n-\tilde{M}+1}\right) = \mathbb{E}_{\tilde{M}} 1 + \mathbb{E}_{\tilde{M}}\left(\frac{n-\tilde{M}}{n-\tilde{M}+1}\right).$$

Now note that

$$\begin{aligned}
\mathbb{E}_{\tilde{M}}\left(\frac{n-\tilde{M}}{n-\tilde{M}+1}\right) &= \sum_{m=0}^{n-1}\left(\frac{n-m}{n-m+1}\right) x_m \\
&\geq \sum_{m=0}^{\lfloor n-\rho_0/(1-\rho_0)\rfloor}\left(\frac{n-m}{n-m+1}\right) x_m \\
&\geq \rho_0 \sum_{m=0}^{\lfloor n-\rho_0/(1-\rho_0)\rfloor} x_m \\
&= \rho_0 \Pr\left(\tilde{M} \leq n - \frac{\rho_0}{1-\rho_0}\right) \\
&\geq \rho_0 \rho_1.
\end{aligned}$$

Thus,

$$\sum_{m=0}^{n-1} \frac{2(n-m)+1}{n-m+1} x_m \geq 1 + \rho_0 \rho_1.$$

Now, using (2) and the definition of ρ_0 and ρ_1, we obtain that

$$\frac{\sum_{m=0}^{\lfloor (n-1)\delta \rfloor} \mathcal{U}_m x_m}{g(\delta)} \leq \frac{1}{1+\rho_0\rho_1} < \frac{1}{2} + \epsilon.$$

Then

$$\sum_{m=0}^{\lfloor (n-1)\delta \rfloor} \mathcal{U}_m x_m < \left(\frac{1}{2}+\epsilon\right) g(\delta),$$

which implies that

$$\sum_{m=0}^{n-1} j_m \mathcal{U}_m \leq \left(\frac{3}{2}+\epsilon\right) g(\delta).$$

Finally, note that

$$\begin{aligned}\alpha^n(\delta) &= \frac{(n+1)^2}{4[\sum_{m=0}^{n-1} j_m \mathcal{U}_m]} \\ &\geq \frac{(n+1)^2}{4g(\delta)} \frac{1}{\left(\frac{3}{2}+\epsilon\right)} \\ &= \beta^n(\delta) \frac{1}{\left(\frac{3}{2}+\epsilon\right)}.\end{aligned}$$

Therefore, there exists n_0 such that for all $n \geq n_0$,

$$\alpha^n(\delta) \geq \beta^n(\delta) \frac{1}{\left(\frac{3}{2}+\epsilon\right)} \geq \frac{\frac{3}{4}}{\left(\frac{3}{2}+\epsilon\right)}$$

for all $\delta \leq \frac{1}{2}$.

By the same argument used in the previous lemma, it can be easily shown that $\alpha^n(\delta)$ is an increasing function of δ. So if $\delta > \frac{1}{2}$, then

$$\alpha^n(\delta) \geq \alpha^n\left(\frac{1}{2}\right) \geq \frac{\frac{3}{4}}{\left(\frac{3}{2}+\epsilon\right)}. \qquad \square$$

Lemma 11. *Let $\epsilon > 0$. Then there exists n_0 such that for all $n \geq n_0$,*

$$\beta^n(\delta) \geq \tilde{\beta}^n(\delta) \geq \frac{\frac{3}{4}}{\left(\frac{3}{2}+\epsilon\right)}.$$

Proof. First note that $1-(1-G(m))(1-H(m)) \geq 1-(1-G(m)) = G(m)$, and since $\mathcal{U}_m$ is decreasing in m, we have

$$\sum_{m=0}^{n-1} \mathcal{U}_m h_m \geq \sum_{m=0}^{n-1} \mathcal{U}_m x_m.$$

Then

$$\tilde{\beta}^n(\delta) = \frac{(n+1)^2}{4[\sum_{m=0}^{n-1} \mathcal{U}_m h_m]} \leq \frac{(n+1)^2}{4[\sum_{m=0}^{n-1} \mathcal{U}_m x_m]} = \beta^n(\delta).$$

Now recall that H is the cumulative distribution function of a binomial random variable with parameters (δ, n). Given that $H(m) \leq G(m)$ for each $m \in \{0, 1, \dots, n-1\}$, we have $1-(1-G(m))(1-H(m)) \leq 1-(1-G(m))^2$. Thus,

$$\sum_{m=0}^{n-1} h_m \mathcal{U}_m \leq \sum_{m=0}^{n-1} j_m \mathcal{U}_m$$

and then $\tilde{\beta}^n(\delta) \geq \alpha^n(\delta)$.

By the last lemma we know that there exists n_0 such that for all $n \geq n_0$,

$$\alpha^n(\delta) \geq \frac{\frac{3}{4}}{\left(\frac{3}{2}+\epsilon\right)}.$$

Thus,

$$\beta^n(\delta) \geq \tilde{\beta}^n(\delta) \geq \frac{\frac{3}{4}}{\left(\frac{3}{2}+\epsilon\right)}.$$

The lower bond $(\frac{3}{2}+\epsilon)$ is arbitrarily close to $\frac{3}{2}$. Then, for each δ, we have that

$$\lim_{n\to\infty} \tilde{\beta}^n(\delta) \geq \tfrac{1}{2}.$$

Since by Lemma 9, $\tilde{\beta}^n$ decreases when n increases, we have that for all n,

$$\beta^n \geq \tilde{\beta}^n \geq \tfrac{1}{2}. \qquad \square$$

Finally, note that the following statements hold:

(a) The term $\tilde{\beta}^n(\delta)$ is an increasing function of δ.

(b) We have $\tilde{\beta}^n(1) = 1$.

(c) Given that we are assuming the uniform distribution of discount factors, we know that β^n has exactly two fixed points if $n > 10$.

(d) We have $\beta^n \geq \tilde{\beta}^n \geq \frac{1}{2}$.

Then $\tilde{\beta}^n$ has, at least, two fixed point: $\delta = 1$ and another stable and higher than $\frac{1}{2}$. Thus, the expected number of agents who go early is at least $\frac{1}{2}n$.

9. Proof of Theorem 3

Theorem 3 follows from Theorem 4 by observing that $\tilde{\beta}^n(F(\delta)) \geq \tilde{\beta}^n(\delta)$ and by employing the same argument used in Section 7.

References

Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin E. Roth (2001), "The market for federal judicial law clerks." *University of Chicago Law Review*, 68, 793–902. [1, 3]

Bulow, Jeremy and Jonathan Levin (2006), "Matching and price competition." *American Economic Review*, 96, 652–668. [12]

Damiano, Ettore, Hao Li, and Wing Suen (2005), "Unravelling of dynamic sorting." *Review of Economic Studies*, 72, 1057–1076. [5, 12]

Diamond, Douglas W. and Philip H. Dybvig (1983), "Bank runs, deposit insurance, and liquidity." *Journal of Political Economy*, 91, 401–419. [2, 13]

Du, Songzi and Yair Livne (2014), "Rigidity of transfers and unraveling in matching markets." Unpublished paper. [5]

Echenique, Federico and Aaron Edlin (2004), "Mixed equilibria are unstable in games of strategic complements." *Journal of Economic Theory*, 118, 61–79. [9]

Frankel, David M., Stephen Morris, and Ady Pauzner (2003), "Equilibrium selection in global games with strategic complementarities." *Journal of Economic Theory*, 108, 1–44. [4]

Fréchette, Guillaume R., Alvin E. Roth, and M. Utku Ünver (2007), "Unraveling yields inefficient matchings: Evidence from post-season college football bowls." *RAND Journal of Economics*, 38, 967–982. [2]

Hałaburda, Hanna (2010), "Unravelling in two-sided matching markets and similarity of preferences." *Games and Economic Behavior*, 69, 365–393. [5]

Kagel, John H. and Alvin E. Roth (2000), "The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment." *Quarterly Journal of Economics*, 115, 201–235. [4]

Kreps, David M. and Evan Porteus (1978), "Temporal resolution of uncertainty and dynamic choice theory." *Econometrica*, 46, 185–200. [2]

Li, Hao and Sherwin Rosen (1998), "Unraveling in matching markets." *American Economic Review*, 88, 371–387. [5, 12]

Li, Hao and Wing Suen (2000), "Risk sharing, sorting, and early contracting." *Journal of Political Economy*, 108, 1058–1091. [5]

Li, Hao and Wing Suen (2004), "Self-fulfilling early-contracting rush." *International Economic Review*, 45, 301–324. [5]

Mongell, Susan and Alvin E. Roth (1991), "Sorority rush as a two-sided matching mechanism." *American Economic Review*, 81, 441–464. [2]

Niederle, Muriel and Alvin E. Roth (2003), "Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match." *Journal of Political Economy*, 111, 1342–1352. [2]

Niederle, Muriel and Alvin E. Roth (2004), "The gastroenterology fellowship match: How it failed and why it could succeed once again." *Gastroenterology*, 127, 658–666. [2]

Niederle, Muriel, Alvin E. Roth, and M. Utku Ünver (2013), "Unraveling results from comparable demand and supply: An experimental investigation." *Games*, 4, 243–282. [5]

Riordan, John (1968), *Combinatorial Identities*. Wiley, New York. [26]

Roth, Alvin E. (1984), "Stability and polarization of interests in job matching." *Econometrica*, 52, 47–58. [1, 4, 7, 13]

Roth, Alvin E. (1991), "A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom." *American Economic Review*, 81, 415–440. [4]

Roth, Alvin E. (2002), "The economist as engineer: Game theory, experimentation, and computation as tools for design economics." *Econometrica*, 70, 1341–1378. [4]

Roth, Alvin E. (2012), "Unravelling in college football." http://marketdesigner.blogspot.mx/2012/09/unravelling-in-college-football.html (accessed 01/06/2013). [2]

Roth, Alvin E. (2013), "Supreme Court clerks—Some hiring already for 2014." http://marketdesigner.blogspot.mx/2013/01/supreme-court-clerks-some-hiring.html (accessed 01/06/2013). [1, 2]

Roth, Alvin E. and Marilda A. Oliveira Sotomayor (1990), *Two-Sided Matching: A Study in Game-Theoretic Modelling and Analysis*. Cambridge University Press, Cambridge. [1, 4]

Roth, Alvin E. and Xiaolin Xing (1994), "Jumping the gun: Imperfections and institutions related to the timing of market transactions." *American Economic Review*, 84, 992–1044. [13]

Shapley, Lloyd S. and Martin Shubik (1972), "The assignment game I: The core." *International Journal of Game Theory*, 1, 111–130. [5]

Sprugnoli, Renzo (2012), "Riordan array proofs of identities in Gould's book." Technical report. [20]

Wetz, Robert V., Charles B. Seelig, Georges Khoueiry, and Kera F. Weiserbs (2010), "Out-of-match residency offers: The possible extent and implications of prematching in graduate medical education." *Journal of Graduate Medical Education*, 2, 327–333. [4]

Submitted 2014-4-13. Final version accepted 2014-11-3. Available online 2014-11-7.