



Unraveling over time

Sandro Ambuehl^{a,*}, Vivienne Groves^b

^a University of Zurich, Department of Economics, Blüemlisalpstrasse 10, 8006 Zurich, Switzerland

^b Stanford University, Graduate School of Business, Knight Management Center, 655 Knight Way, Stanford, CA, 94305, United States of America

ARTICLE INFO

Article history:

Received 17 August 2018

Available online 10 March 2020

Keywords:

Unraveling

Two-sided matching

Continuous time game

Law clerks

ABSTRACT

Unraveling, the excessively early matching of future workers to employers, leads to hiring decisions based on severely incomplete information. We provide a model of unraveling in a one-to-one matching market without transfers. Its distinguishing feature is the gradual arrival of information about students that occurs over an extended period of time during which matches can be made. In equilibrium, the market spreads thinly over that period and employers' attractiveness to students is uncorrelated with their worker's productivity. Our model connects previous, seemingly unrelated models of unraveling and it highlights the implications of the two-period assumption on which they rely. Our setting permits the analysis of information timing policies. These are effective only if they provide a sudden and sufficiently large surge in information. Our main application is in the market for clerks to U.S. Federal Appellate Courts, a significant input into the efficiency of the justice system.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Entry-level labor markets frequently suffer from the excessively early matching of workers to employers, a phenomenon known as *unraveling* (Roth and Xing, 1994). Faced with a cohort of prospective employees all graduating at similar points in time, some employers attempt to preempt competition by hiring promising students early. By doing so, these employers diminish the pool of applicants available for hire after graduation, prompting others to hire promising students early, too. The resulting mutual undercutting frequently leads to hiring employees up to multiple years before employment starts, to the dismay of many participants on either side of the market.

In this paper we ask: what are the equilibrium properties of two-sided matching markets without transfers if unraveling is allowed to occur without intervention? How severe will unraveling become? What are the efficiency properties of unraveled markets? What policies can be used to alleviate unraveling, especially if centralized clearinghouses are no option, or do not eliminate unraveling? We model a market that may take place over an extended period of time during which about students arises gradually. Our setting contrasts with the previous literature, which almost exclusively considers two-period models. It allows us to connect earlier, seemingly unrelated models of unraveling (Halaburda, 2010; Du and Livne, 2013), to examine the robustness of their implications, and to examine policies that manipulate the timing of information release.

Our main application is in the market for U.S. federal appellate court clerks, a significant input into the efficiency of the justice system.¹ Judges depend heavily on their law clerks (Avery et al., 2001); their role includes everything from the

* Corresponding author.

E-mail addresses: sandro.ambuehl@econ.uzh.ch (S. Ambuehl), vivienne.groves@gmail.com (V. Groves).

¹ Other examples of unraveling in entry-level labor markets include clinical psychology internships, dental residencies, optometry residencies and Japanese university graduates (Roth and Xing, 1994). More recently, unraveling has also been observed in the market for entry-level private equity bankers

management of disputes in district courts to the drafting of opinions at the U.S. Supreme Court (Bonica et al., 2017). Next to the non-negotiable salaries and the gradual arrival of information, our model captures a third characteristic of this market: given the elite nature of federal judicial clerkship positions and the power differential between students and federal judges, students feel pressured to accept any offer they receive (see Section 2).

Our theoretical analysis yields four results, each driven by the fact that employers hire students as soon as the risk of being preempted by another employer no longer justifies waiting for more information to arrive.

First, any equilibrium is characterized by the same set of widely dispersed points in time at which students are hired. The temporal spread of the market, which increases with the number of employers, implies that most hiring decisions are made based on severely incomplete information. Because equilibrium hiring times are unique, however, this informational inefficiency cannot be addressed by ‘soft’ interventions such as appeals to postpone hiring to some later time—there is no informationally more efficient equilibrium the market could coordinate on.

Second, unraveling entirely erodes the advantage that more desirable employers would have if the market was concentrated in time. Each employer is matched with a student of the same expected skill level as all other employers. We assess the efficiency and welfare properties of the unraveling equilibria in comparison to the full-information stable match, which, in our setting, is positively assortative. We find that while unraveling necessarily involves an informational loss, it may decrease or increase total market production, depending on whether student skills and employer desirability are complements or substitutes. Moreover, only sufficiently desirable employers would prefer the full-information stable match; less desirable employers obtain higher payoffs in the unraveled equilibrium. Unraveling is not, as it is sometimes portrayed,² a social dilemma game in which participants coordinate on a Pareto-inefficient equilibrium.

Third, we impose the restriction that matches may be made only at two predetermined points in time, thus reducing our model to a two-period setting. We find that if the first period occurs early, and thus entails little information about students, our model mirrors Halaburda (2010). By contrast, if the first period occurs late, and thus entails much information about students, it resembles Du and Livne (2013). Importantly, we also show that the mechanisms proposed to eliminate unraveling in Halaburda (2010) are specific to the two-period setting; they are ineffective in our temporally extended market.

Fourth, we examine information timing policies. We show that withholding information about students and releasing it later affects unraveling only if a sufficiently large amount of information is issued at a single point in time. Otherwise, information timing policies may be ineffective (if they do not cause an informational discontinuity) or may exacerbate unraveling (if they cause an informational discontinuity that is not sufficiently large to curb unraveling).

Our main contribution to the theoretical literature on unraveling in matching markets consists in relaxing the two-period assumption that pervades existing work. We show that the two settings (two periods or extended market) lead to qualitatively different predictions and policy implications. Unlike two-period models, our setting also permits the study of information timing policies.³ Other papers on unraveling identify separate drivers and moderators of unraveling. Halaburda (2010) shows that unraveling intensifies as participants’ preferences become more similar, Du and Livne (2013) study the effect of introducing transfer payments, Niederle et al. (2013) examine how unraveling depends on imbalances in the demand and supply of applicants, Fainmesser (2013) studies how unraveling depends on the structure of a network of informal connections, and Vohra (2019) examines how unraveling is affected by the opportunity to change employers after being hired.⁴ This paper also relates to Anderson et al. (2017) which provides a general model of timing games.⁵

Our work also contributes to the study of the market for law clerks. Next to the descriptive research in Wald (1990) and Avery et al. (2001, 2007), Haruvy et al. (2006) use laboratory experiments and computational methods to investigate proposed reforms in this market.

The remainder of this paper is organized as follows. Section 2 presents a brief overview of the market for law clerks as it relates to our model. Section 3 introduces the model. Section 4 contains the theoretical analysis. Finally, Section 5 concludes. All mathematical proofs are in Appendix A.

(Alden, 2014), for technology jobs (Ante, 2012), for faculty in philosophy departments (Kueppers, 2016), for postdoctoral positions in theoretical physics (Yu, 2017), for college athletes (Popper, 2014), for professional soccer players (Hautmann, 2013), and in several medical subspecialties such as the epilepsy and clinical neurophysiology fellowships (Vidaurre and Campbell, 2017).

² See, e.g., Alden (2014).

³ Ostrovsky and Schwarz (2010) study optimal information design in a two-period model of unraveling.

⁴ Additional related work includes Echenique and Pereyra (2016) and Damiano et al. (2005). A separate literature on matching *with* transfers that starts with Li and Rosen (1998) includes Suen (2000), Li and Suen (2000, 2004), Du and Livne (2013), and Ishii et al. (2018).

⁵ Our model does not fall into the class of games analyzed in Anderson et al. (2017). The reason is that unlike in our paper, in Anderson et al. (2017), a player’s payoff may only depend on his stopping time and stopping quantile, but may not depend on the time at which other players have stopped. Moreover, while employers in our paper learn about student’s skills, employers in the unraveling application in Anderson et al. (2017) learn about their own needs.

2. The market for clerks to federal appellate judges

Each year, students in law schools across the U.S. seek one of a small number of clerkship positions with a federal appellate judge. These are the most prestigious of judicial clerkships available to new J.D. graduates, and are the toughest to obtain.⁶

From the viewpoint of judges, finding good law clerks is vital. A judge's staff of clerks is a major component of their productivity, potentially making "the difference between a bad year and a wonderful one" (Kozinski, 1991).⁷ Judges thus compete fiercely for the best students in a cohort, often asking them to sign a contract before they may consider offers from other judges. As a result, the market for law clerks has suffered from unraveling for multiple decades. Both students and judges frequently complain about the nature of that market. Repeated attempts at reform have failed. Most recently, the Federal Law Clerk Hiring plan, instituted in 2005 and officially abandoned in November 2013, was an attempt to coordinate the hiring process to begin no earlier than an agreed upon date (Online System for Clerkship Application and Review, 2013).

Our model matches three characteristics of this particular market, the first two of which are also present in many other entry-level labor markets. First, salaries are fixed and not negotiable, as stipulated by the Judicial Salary Plan.⁸

Second, information about students arrives gradually. Time to completion in a J.D. program is three years and includes a large number of exams and evaluations, which progressively offer information about students.⁹

Third, students feel pressured to accept any offer they might receive (Haruvy et al., 2006; Avery et al., 2001). Law career counselors reason that "one can never tell how far the wrath of a judge scorned will extend and exactly what the consequences will be, for you or for others" (Strauss, 2015). They exhort students to avoid declining any offer "merely as a matter of strategy with the expectation of receiving an offer from another judge" lest they jeopardize their chances for an offer from another judge (Strauss, 2015).¹⁰ Systematic survey evidence confirms the prevalence of this norm. In the year 2000, the most recent year for which such data are available, 73% of the students accepted the first offer they received (Avery et al., 2001).¹¹

In order to obtain detailed information about the timing of this market, we have collected data from the "Clerkship Notification Blog" (McDaniel et al., 12/08/2013), a public internet forum coordinated by law students for federal law clerk positions starting in the year 2009.¹² Our data covers statements concerning the start or finish of hiring of 168 judges (a judge may hire several clerks in a given year).¹³

Fig. 1 plots, for each judge, the first and last points in time an entry was posted about a judge hiring a clerk or completing the hiring process. The dashed vertical line indicates date of the Hiring Plan, which was still in effect in 2008. The market is widely dispersed over a time period of about half a year. There is no lull in hiring at any point in time between April and December.¹⁴ Similar dispersion has also been documented in the market for gastroenterology fellowships (Niederle et al., 2006) and has been reported in other markets (Kueppers, 2016; Yu, 2017). Existing two-period models preclude this characteristic feature of unraveled markets by design.

⁶ Clerkship positions at federal appellate courts represent a small fraction of all judicial clerkship positions; about 10% of federal judges are at federal appellate courts (Administrative Office of the United States Courts, 2013).

⁷ Judges' decisions on a broad range of matters such as whether to file a separate opinion or to dissent in a case are based upon the support they anticipate from their clerks. Moreover, law clerks are essentially the only persons a judge can talk to in depth about a case (Wald, 1990).

⁸ While allowing for flexible salaries sometimes solves the unraveling problem in theory (Du and Livne, 2013), in practice there are many markets with flexible salaries that nonetheless unravel. The market for entry-level employees with J.D. degrees at elite law firms is an example (Ginsburg and Wolf, 2003). For each recent cohort, there was a single salary (\$160k in 2011) offered by the vast majority of elite law firms (Collins, 2012), rendering that market into one of *de facto* fixed salaries.

⁹ Judge Kozinski (1991) articulates the fundamental tradeoff: "Time ... can buy more information, reducing uncertainty. But the cost of delay may be the loss of opportunity, as some candidates receive offers from judges who believe they can make reasoned judgments on the basis of less information." He also lists specific pieces of information that are relevant to a judge's hiring decision if they are available.

¹⁰ Similarly, the Yale Law School's Career Development Office urges applicants *not* to ask questions regarding the judge's policies on hours of work, vacation, or post-clerkship jobs lest they diminish their chances of receiving an offer (Yale Law School Career Development Office, 2008).

¹¹ There is evidence that students can be more strategic in their decisions which judges to apply to, and in accepting or avoiding phone calls (Avery et al., 2001). Our model will abstract from such behavior.

¹² Official records on the points in time at which judges hire clerks do not exist, since judges who hire earlier than the Hiring Plan date are breaking the very rule this plan was meant to uphold. The blog largely consists of time-stamped entries such as "Ikuta has hired one HLS" or "Noonan is done", meaning "Judge Ikuta has hired one Harvard Law Student", and "Judge Noonan has finished hiring for this season". While blog entries are anonymous, they are often seconded and mistakes based on hearsay are frequently corrected. Often, they are written by current law clerks divulging inside information. See Appendix C for details.

¹³ These 168 judges comprise just over 80% of the 239 judges currently active in the districts that we analyze.

¹⁴ Notably, judges hiring early do so at a loss of information, such as the release of grades (around June and July) or students' performance in their first position as an intern at a law firm in the summer between second and third year. While our data cover one single year, dispersion in hiring times persists past 2008. Tobias (2013) details the various states of hiring that different courts are in for October 2013, which fall into six categories ranging from "finished hiring" to "have recently started or not yet begun" and "remain unclear." A small number of blog postings from the hiring seasons in more recent years are available on a different forum (www.top-law-schools.com, 2013-2015) and suggest that hirings are spread over similarly long intervals of time.




Fig. 1. Distribution of hiring dates in the 2008 market for judicial clerks at Federal Appellate Courts. Each horizontal bar corresponds to a judge. It stretches from the date the judge is reported have started hiring to the date the judge is reported to have finished hiring. Judges are ordered by the start date of their hiring activity. If data on hiring end date is unavailable, the length of a bar is set at 2 days for visibility.

3. Model assumptions

Setting. We study a two-sided, one-to-one matching market without transfers. There are $J \geq 2$ employers and I students. The number of students exceeds the number of employers, $I > J$. Each employer can hire exactly one student. Each student i is identified with a skill level $q_i(t)$ that evolves stochastically over the time interval $[0, T]$. We think of T as the start date of employment. The processes $q_i(t)$ may represent the accumulation of skills as well as gradually uncovered information about pre-existing differences in ability. The properties of these processes matter only through the conditional expectations they imply, which we will specify precisely below. If each $q_i(t)$ is an independent standard Brownian motion with $q_i(0) = 0$, for instance, all requirements will be satisfied.

While the continuity of time in our model is a convenient technical assumption, Appendix B.2 shows, for the case of two employers, that our substantive results are unchanged in a discrete-time approximation.

Objective functions. Maximizing production is the employers' sole concern. Employer j 's production is proportional to the matched student's skill level at time T , $q_i(T)$. Hence, employers thinking of hiring before T face a prediction problem. Employer j 's production is given by $v_j = \alpha_j q_i(T)$, with $\alpha_j > 0$ for all j . The vector of parameters $(\alpha_j)_j$ allows us to capture complementarity or substitutability between students and employers.¹⁵ Employers not matched with a student receive payoff z , which we assume to be less than the *ex-ante* expected payoff that can be obtained from hiring the J -best student at time T . This assumption ensures that all employers will hire a student.

We will consider the case in which students have common preferences over employers, which allows us to make statements that would otherwise be vacuous (for instance regarding the relation between employer rank and equilibrium payoffs). Our results apply equally if student preferences are heterogeneous, except in isolated cases that we will highlight. Specifically, each student prefers employer 1 to employer 2 to employer 3, and so on. We say that employer 1 is *more desirable* than employer 2 who is in turn more desirable than employer 3, and so forth.

Actions and matching. At each time $t \in [0, T]$ employers must decide whether to remain in the market or exit the market by hiring a student. All matchings are binding for both sides. Students choose between employers if they receive multiple job offers at the same time. Otherwise, they must accept any job offer they receive. The assumption concerning simultaneous offers is not vacuous because equilibrium incentives lead employers to act at the same points in time.¹⁶ While reluctance to reject offers is arguably an idiosyncratic feature of the market for law clerks, one can view it as an approximation to markets in which students are greatly concerned about matching to one of a set of *elite employers* and are less concerned about which of the elite employers they match to (see Appendix B.3).

Formally, at any moment in time $t \in [0, T]$ an employer can either stay in the market by choosing action s ("stay") or exit by choosing one of the actions in $E = \{e_1, \dots, e_I\}$ ("exit"). If, at time t , an employer is the only one to play some action $e \in E$, he is automatically assigned the student whose skill level is, at time t , highest out of the students who are still unmatched. If more than one employer attempts to exit at the same time t , students are matched to employers

¹⁵ If α_j is larger for more desirable employers, student skill and employer desirability are complements. They are substitutes if α_j decreases with employer desirability.

¹⁶ We study the effects of relaxing the no-rejections assumption in the case of two employers in Appendix B.3.

assortatively.¹⁷ An employer's choice of e_k indicates his willingness to accept any student ranked k or higher amongst those who are available at that instance. Implicitly, this matching procedure embodies the assumption that a student may reject an offer by an employer if the student simultaneously receives an offer by a more desirable employer.¹⁸

Information. Employers know the stochastic process by which students' skill levels evolve. They also observe all other employers' previous actions. They do not, however, observe the current realization of the student skill levels $q_i(t)$ at any $t \in (0, T)$. The assumption can be viewed as an approximation to the case in which employers observe the rank of students, but not magnitude of their skill.¹⁹ Critically, employers can infer what skill level to expect when they exit the game at any time t , conditional on the points in time $t_1 \leq \dots \leq t_k$ at which previous employers have exited. With this assumption, the model is a game of complete and perfect information.²⁰

Tractability is the main justification for our assumption that employers do not observe time t student skill levels $q_i(t)$. Moreover, our informational assumptions are a generalization of those in Halaburda (2010).²¹

Conditional expectation functions. We use the expression $x_k(t|t_1, \dots, t_{k-1})$ to denote the skill level an employer expects from hiring the best available student at time t given the previous hire times of other employers $t_1, \dots, t_{k-1}$. We call x_k a *conditional expectation function*. Observe that, for each k , the function x_k has k arguments. The conditional expectation functions also cover the case in which multiple employers exit the market at the same point in time. For instance, if no employer has exited so far, and both employer 1 and employer 2 exit at t , then the more desirable employer 1 expects student skill level $x_1(t)$ and the less desirable employer 2 expects student skill level $x_2(t|t)$.

Our main assumptions regarding the structure of information arrival concern the family of conditional expectation functions $(x_j(t|t_1, \dots, t_{j-1}))_{1 \leq j \leq J}$. We require each conditional expectation $x_{k+1}(t|t_1, \dots, t_k)$ is

- (i) jointly continuous in all arguments,
- (ii) strictly increasing in t , and
- (iii) weakly decreasing in $t_1, \dots, t_k$,

for all $0 \leq k \leq J - 1$ and all $t \geq t_k \geq \dots \geq t_1$.

Continuity captures the intuition that information arrives gradually. It does not hold, for instance, if students' skill levels jump with positive probability at some known point in time. $x_{k+1}(t|t_1, \dots, t_k)$ is increasing in t because predicting less far into the future is easier—the closer t is to the start date of employment, T , the higher the probability that the student who is deemed best at time t will still be ranked highly at T . The fact that $x_{k+1}(t|t_1, \dots, t_k)$ is decreasing in $t_1, \dots, t_k$ is closely related. The earlier an employer's rivals hire, the more likely the hired students have dropped in rank by time t , and hence the better the pool of remaining students from which the employer can choose.

For our additional results, we also make the following, natural assumptions. First, unraveling is inefficient from the employers' point of view. Intuitively, if hirings are made early, then there is a positive probability that some employer will be matched to a student whose rank at T is lower than J . Such a student would not be hired at all if all employers postponed hiring until T . Second, the expected payoff of an employer who exits at some given time t is lower the more employers have already exited. Intuitively, the more students have already been hired, the less likely highly skilled students are still available for hire. Third, we assume that all students are identical at $t = 0$. That latter assumption is inessential for our results, but simplifies presentation. In fact, the analysis of subgame-perfect equilibria requires that strategies specify actions for the case in which no students have been hired up to any time t . Such subgames correspond to the case in which students are differentiated from the start. Continuing the list of properties of the conditional expectation functions, these additional requirements read as follows.

- (iv) For all $t_1 \leq \dots \leq t_j$ with $t_k < T$ for some k , we have $\sum_k x_k(t_k|t_1, \dots, t_{k-1}) < \sum_k x_k(T|T, \dots, T)$.
- (v) For all $0 \leq k \leq J - 1$, $t_1 \leq \dots \leq t_k$ with $t_1 > 0$, and $t \geq t_k$, we have $x_{k+1}(t|t_1, \dots, t_k) < x_k(t|t_1, \dots, t_{k-1})$.
- (vi) For all $0 \leq j \leq J$, we have $x_j(0|0, \dots, 0) = 0$.

It is far from obvious that a family of stochastic processes $(q_1(t), \dots, q_I(t))$ exists whose conditional expectation functions satisfy all properties (i) to (vi). Appendix B.1 proves that a collection of independent standard Brownian motions with $q_i(0) = 0$ for all i indeed satisfies all requirements.

¹⁷ In the case of heterogenous student preferences, we assume that the matching occurs according to an ex-post pairwise stable mechanism.

¹⁸ This matching corresponds to a serial dictatorship mechanism in which priority order is determined by the participating employers' desirability.

¹⁹ Our informational assumptions are not equivalent to observable rank information to the extent that rank information would permit inference about the magnitude of the difference in student skill levels. (For example, if a student is ranked at the top for a long time, it is likely that his skills exceed that of the second-ranked student by a large magnitude.)

²⁰ We study the effects of relaxing this assumption in Appendix B.4.

²¹ In that model, I students who are undifferentiated in the first period are randomly assigned one of I fixed values in the second period. In our setup, students are undifferentiated at $t = 0$, expected skills conditional on rank j at time T are given by $x_j(T|T, \dots, T)$, and ranks are assigned randomly.




Fig. 2. The model with two employers. The horizontal axis plots time t . $x_1(t)$ is the student skill level an employer expects to obtain if he hires a student at time t , given that no other student has been hired before. $x_2(T|t)$ is the student skill level an employer expects to obtain if he hires a student at T , given that one student has previously been hired at time t . In every equilibrium of the model, the first hire is made at t_1^* , and the second hire is made at T .

Histories, strategies, and equilibrium. A *history*, H_t , of the game up to time t describes what action has been taken by which player at what time up to, but excluding, time t .²² The history of the game is common knowledge. We let $\mathcal{H}$ denote the set of histories and we use the symbol $\emptyset$ to denote the empty history.

A *pure strategy* of employer j is a function $\sigma_j : \mathcal{H} \rightarrow E \cup \{s\}$ that maps a history H_t into an action taken at time t . We write $\sigma_j(H_t) = a$ to refer to a strategy in which employer j plays action $a \in \{s, e_1, e_2, \dots, e_J\}$ at time t after history H_t with certainty. We require that any interval of time during which an employer plays some action $e \in E$ is closed on the left. This condition, which we call *right-continuity*, ensures that each strategy profile maps into exactly one vector of payoffs, as we show explicitly in Appendix A.1 While our results are unchanged even if we permit mixed strategies (see Appendix A.3), we simplify the exposition by focusing on pure strategy profiles.

We derive our main result using subgame perfect Nash equilibrium as a solution concept. We also consider the limit of subgame perfect ϵ -equilibria for ϵ approaching 0.²³ Moreover, Appendix B.2 shows, for the case $J = 2$, that a discrete-time approximation in which the length of each time interval tends towards zero coincides with our continuous-time results derived using ϵ -equilibria.

4. Analysis

We begin the analysis by building intuition for the case of two employers. Subsection 4.2 fully characterizes the set of equilibria with J employers. We then use that characterization to derive the behavioral characteristics (Subsection 4.3) and the efficiency and welfare properties of unraveled markets (Subsection 4.4). Subsection 4.5 compares our model to the two-period models by Halaburda (2010) and Du and Livne (2013). Finally, Subsection 4.6 examines information timing policies.

4.1. Heuristic solution with two employers

Consider a market with two employers (called employer 1 and employer 2), and assume $x_2(T|T) > 0$.²⁴ We first argue that, in any subgame-perfect equilibrium, employers may exit at only one of two points in time that we call *hiring times*. We then argue that a subgame-perfect equilibrium in fact exists.

To characterize the hiring times, observe that the last student must be hired at time T : employers have no reason to exit the market early if there is no competition, but they do benefit from information that arises over time ($x_2(t'|t)$ is increasing in t'). Hence, if one employer has left the market, the other employer will wait until T and will hire the best of the remaining students at that time.

When will the first student be hired? Consider Fig. 2. It plots the time at which the first student is hired on the horizontal axis, and the student skill level an employer can expect to obtain on the vertical axis. The employer who exits first can expect that the student he hired will have skill level $x_1(t)$ once employment starts (at time T). The conditional expectation $x_1(t)$ is increasing in the point in time t at which that employer exits—the closer t is to T , the higher is the likelihood that the student who was the most skilled at t will still be the most skilled at T .

Crucially, the point in time t at which the first student is hired also affects the skill level $x_2(T|t)$ that the employer who hires at T will obtain. The student who is best at some early time t is unlikely to be the best student at T . Hence, if the

²² In conjunction with the remaining model assumptions, it will suffice if players know which employer has hired a student and when.

²³ For any $\epsilon > 0$, a strategy profile σ is an ϵ -equilibrium if no player can gain more than ϵ by unilaterally deviating from his strategy.

²⁴ This assumption ensures that there is no rush at $t = 0$. See Subsection 4.3.

first match occurs early, it is likely that the employer who exits at T will obtain the best student. If the first student is hired late, however, the employer who exits at T will likely have to settle for the second-best student. ($x_2(T|t)$ is decreasing in t .)

The above considerations also explain why the two curves $x_1(t)$ and $x_2(T|t)$ cross. Hiring the best student at a very early time is similar to hiring a student at random. This is worse than hiring the best remaining student—who has a high chance of being the best student overall—at T . Hence, for early t , we have $x_1(t) < x_2(T|t)$.²⁵ At some late time t , by contrast, the best student at that time is likely to remain the best student until T . Thus, being the first employer to hire at some late time t is better than hiring the best remaining student at T (likely the second-best student). Hence, for late t , we have $x_1(t) > x_2(T|t)$.²⁶ The conditional expectation functions, which are continuous, must therefore cross at some intermediate time that we denote t_1^* . (Formally, t_1^* satisfies $x_1(t_1^*) = x_2(T|t_1^*)$.)

It is clear that the first match in any equilibrium must happen at t_1^* . It will not occur before t_1^* because at such early points in time, waiting beats exiting. Neither will it occur after t_1^* —if both employers postponed hiring until some time t strictly after t_1^* , exiting at t would beat waiting, allowing one of the employers to profitably preempt the other.

Indeed, there is a subgame perfect equilibrium in which the first exit occurs at t_1^* . It involves the following strategies. Both employers play s at any time before t_1^* . At t_1^* , both employers play e_1 at t_1^* (causing the more desirable employer to exit at that time, and the less desirable employer to remain in the game). After t_1^* , if one of the employers has already exited, the remaining employer plays s at any time $t < T$, and plays e_1 at time T . If both employers are still in the market at some $t > t_1^*$, they both play e_1 at that time. With these strategies, no employer has an incentive to unilaterally deviate: exiting before t_1^* is not profitable because at these early points in time, waiting beats exiting. Moreover, if an employer fails to play e_1 at t_1^* neither employer's expected payoff will change; a reversal of the order of exit is the only possible consequence. Failure to play e_1 at $t > t_1^*$ if both employers are still in the market will leave the employer with $x_2(T|t)$, which is lower than $x_1(t)$. Finally, no employer has an incentive to deviate from remaining in the market until time T if the opponent exited at some time t' , because $x_2(t'|t)$ is increasing in t' for any t .

Hence, unraveling causes complete equalization of expected student skills across the two employers (because $x_1(t_1^*) = x_2(T|t_1^*)$). This is true even though both employers try to hire the same student at the same point in time, and even though employer 1 is more desirable than employer 2. As we show next, the mechanisms just explained generalize to the case of arbitrarily many employers.

4.2. Equilibrium characterization

Even with an arbitrary number of employers, equilibrium strategies within each subgame are characterized by the points in time at which employers exit the market. These *subgame hiring times* are characterized by employers weakly preferring exit over staying in the game. If k employers have exited in a history H_t , we will denote the subgame hiring times by $t_{k+1}(H_t), \dots, t_J(H_t)$. We define them through four properties.

Definition 1. (Subgame hiring times) Consider a history H_t with past hiring times $h = (t_1, \dots, t_k)$ that satisfies $t_1 \leq \dots \leq t_k \leq t$. Then $t_{k+1} \leq \dots \leq t_J$ are *subgame hiring times* for H_t if they satisfy $t_{k+1} \geq t$ and

- (i) $t_J = T$.
- (ii) For all $j \in \{k+1, \dots, J-1\}$:

$$x_j(t_j|h, t_{k+1}, \dots, t_{j-1}) \geq x_{j+1}(t_{j+1}|h, t_{k+1}, \dots, t_j).$$
- (iii) If $t < t_j < t_{j+1}$ for some $j \in \{k+1, \dots, J-1\}$, then $x_j(t_j|h, t_{k+1}, \dots, t_{j-1}) = x_{j+1}(t_{j+1}|h, t_{k+1}, \dots, t_j)$.
- (iv) If $t_j = t_{j+1}$ for some $j \in \{k+1, \dots, J-1\}$, then $t_{j+1} = t$.

Property (i) requires that the last student is hired at T . The reason is that once all but one employers have exited, the remaining employer will optimally wait for the arrival of all pertinent information without a risk of preemption. Property (ii) demands that exiting at a later subgame hiring time is no more profitable than exiting at an earlier subgame hiring time. Otherwise, employers would have an incentive to postpone hiring. Property (iii) requires that if all future exits from the market occur at subgame hiring times, each employer is indifferent between exiting at any of the (strictly) future subgame hiring times. Finally, property (iv) concerns *rushes*. It states that hiring times must not coincide, except at the beginning of a subgame. In other words, rushes may never be known to occur in the (strict) future—if they were, less desirable employers would have an incentive to hire the best available student just before the rush is known to take place.

The following proposition relies on the monotonicity of the conditional expectation functions and shows that for each history there is exactly one vector of subgame hiring times.

Proposition 1. For all $k \in \{0, \dots, J-1\}$ and for each history H_t in which k employers have exited, there exists a unique vector of subgame hiring times $t_{k+1}(H_t) \leq \dots \leq t_J(H_t)$.

²⁵ Formally, this is true because $0 = x_1(0) < x_2(T|T) < x_2(T|0)$, and due to continuity.

²⁶ Formally, this is true due to $x_1(T) > x_2(T|T)$ and due to continuity.

Proposition 1 allows us to construct a strategy profile $\sigma = (\sigma_1, \dots, \sigma_J)$ which we will show constitutes a subgame perfect Nash equilibrium. For each history H_t with previous hiring times $h = (t_1, \dots, t_k)$ that satisfy $t_1 \leq \dots \leq t_k \leq t$, we impose the following properties on σ_j .

First, employer j may not exit at any time that is not a subgame hiring time (property (i)). Second, employer j must exit with any student who is at least J -best at T (property (ii)). Third, employer j must attempt to exit whenever doing so is strictly more profitable than remaining in the market (property (iii)). By implication, employer j participates in every rush in which that employer has a chance of employing a student. Importantly, this property also ensures that at least one employer exits at each subgame hiring time: any attempt to postpone hiring by $\epsilon > 0$ will be preempted by some other employer. Fourth, employer j may or may not attempt to hire a student at any other subgame hiring time (property (iv)). Formally, σ_j thus satisfies:

- (i) $\sigma_j(H_t) = s$ if $t < t_{k+1}(H_t)$.
- (ii) $\sigma_j(H_T) = e_j$.
- (iii) If there exists $m \geq 1$ such that $t_{k+1}(H_t) = t_{k+m}(H_t) = t$ and $x(t_{k+m}(H_t)|h, t_{k+1}(H_t), \dots, t_{k+m-1}(H_t)) > x(T|h, t_{k+1}(H_t), \dots, t_{J-1}(H_t))$, then $\sigma_j(H_t) = e_m$.
- (iv) $\sigma_j(H_t) = s$ or $\sigma_j(H_t) = e_1$ at any other subgame hiring time.

We define the full strategy profile σ by resolving the ambiguity in part (iv): At each time t that coincides with $l \geq 1$ subgame hiring times, the $(l+1)$ most desirable remaining employers play action e_l . The crucial role of the $(l+1)$ st employer is to ensure that none of the l employers who do exit have an incentive to postpone exiting.

The next proposition shows that subgame hiring times fully characterize equilibrium play. This proposition generalizes to the case of mixed strategies, appropriately defined, as formalized in Appendix A.3.

Proposition 2.

- (i) σ is an SPE profile.
- (ii) In each SPE profile, students are hired at subgame hiring times in each subgame.

The uniqueness of equilibrium hiring times has a sobering policy implication. It shows that ‘soft’ interventions, such as appeals to wait for a longer time before hiring students, cannot lead to coordination on some ‘better’ equilibrium.²⁷ Policies are effective only if they alter the conditional expectation functions, as we explain below.

4.3. Equilibrium behavior

Proposition 2 allows us to answer our introductory question: What are the equilibrium properties of a market in which unraveling occurs without intervention? Our first answer is that hiring activity will be dispersed over an extended amount of time. Indeed, in any equilibrium one employer exits at each subgame hiring time $t_1(\emptyset), \dots, t_J(\emptyset)$. We call these values the *equilibrium hiring times*, and denote them by $t_1^*, \dots, t_J^*$. Unless there is a rush at $t = 0$, we have $t_1^* < \dots < t_J^*$, as summarized in part (i) of Corollary 1 below. A sufficient condition for the absence of a rush at $t = 0$ is $x_j(T|T, \dots, T) > 0$.^{28,29}

Corollary 1.

- (i) If $x_j(T|T, \dots, T) > 0$, equilibrium hiring times are dispersed: $0 < t_1^* < \dots < t_J^* = T$.
- (ii) If the number of employers increases to $\tilde{J} > J$ with associated equilibrium hiring times $\tilde{t}_1^*, \dots, \tilde{t}_{\tilde{J}}^*$, then $\tilde{t}_j^* < t_j^*$ for all $j = 1, \dots, J$.
- (iii) Any order of employers exiting the market is consistent with a limit of subgame perfect pure strategy ϵ -equilibria.

To understand the extent of dispersion in hiring times (if there is no initial rush), we bound the time of the first match from above. If the first student were hired as soon as the expected payoff from hiring the currently best student were equal to being randomly assigned one of the J most highly skilled students at T , we would still observe a later first hiring

²⁷ This result may explain why in the market for law clerks all past attempts at coordinating hiring to a specified point in time have failed (see Section 2).

²⁸ The reason is that hiring at $t = 0$ yields payoff 0 whereas hiring at T yields expected student skills of at least $x_j(T|T, \dots, T)$. Hence, if $x_j(T|T, \dots, T) > 0$, then an employer would want to deviate from any strategy that prescribes hiring at $t = 0$.

²⁹ To understand the case in which students are differentiated already initially, we let the game start at some $\underline{t} \in (0, T)$. The construction of σ shows that in such a subgame, a rush occurs at $\underline{t}$, after which the game continues in a largely unchanged manner. Specifically, there is some number $\hat{j} \in \{1, \dots, J\}$ such that $t_1^* = \dots = t_{\hat{j}}^* = \underline{t} < t_{\hat{j}+1}^* < \dots < t_J^* = T$. By property (iii) of the definition of σ , all employers must participate in the rush at $t = \underline{t}$, and the $\hat{j}$ best students at $\underline{t}$ are assortatively matched to employers. Amongst employers who hire during the rush, more desirable employers obtain students with higher expected skills; all employers who hire after $\underline{t}$ obtain students with the same, lower, expected skills.

time than we do in equilibrium. This follows from the fact that the first student is hired as soon as doing so yields the equilibrium expected student skill level, and because the equilibrium expected student skill level is lower than that from the random matching.

Reasoning similarly, we can show that competitive pressure leads to earlier hiring times, as stated in part (ii) of the corollary. If a larger number of employers compete for the same set of students, then the j th student is hired at an earlier point in time for all j . Intuitively, the larger J , the lower the student skills an employer can expect from a random assignment to one of the J most highly skilled students at T . This expectation is an upper bound on $x_1(t_1^*)$. Because $x_1(t)$ is increasing in t , a larger number of employers J therefore implies an earlier first match time. The monotonicity of the conditional expectation functions then implies that all other hiring times must be earlier as well.

A third equilibrium property we can characterize concerns the order in which employers exit.³⁰ By construction of the subgame hiring times, in any equilibrium all employers obtain a student with the same expected skill level, irrespective of their own desirability (with the possible exception of employers who participate in a rush at $t = 0$). Accordingly, employers will be indifferent between exiting at different equilibrium hiring times. By implication, any order of exit is consistent with a limit of subgame perfect pure strategy ϵ -equilibria, for ϵ approaching zero. It is also the result one obtains from modeling the continuous-time game as the limit of discrete-time games, as Appendix B.2 shows for the case of two employers. In the case of subgame-perfect equilibrium (in which ϵ equals zero, rather than approaches it), however, the technical requirement that strategies be right-continuous implies that the most desirable employer exits first, followed by the second-most desirable employer, and so on.³¹ Because that prediction about the order of exit arises from a technical rather than from a substantive assumption, and because it is not robust to the use of two alternative equilibrium concepts, we regard the anything-goes result as the main prediction of our model regarding the order of exit.³²

4.4. Efficiency and welfare in equilibrium

As argued above, in equilibrium, each employer obtains a student of the same expected quality, regardless of his own desirability or (equilibrium) time of exit. Unraveling therefore implies an extreme level of redistribution among employers, compared to the full information assortative match (which is the unique *ex post* pairwise stable match in our setting).

While unraveling entails an informational inefficiency, it is *not* Pareto inferior to a positive assortative matching at T . In the full information assortative match, the least desirable employer, employer J , obtains the J -best student. In the unraveling equilibrium in which employer J hires last, there is a chance that employer J obtains a better student than the J -best. Accordingly, employer J is strictly better off with unraveling.^{33,34}

Relatedly, while matches in an unraveling equilibrium are made based on partial information, an assortative match at T is not necessarily more economically efficient than the unraveling equilibrium. An assortative match maximizes market output if employer desirability and student skills are complements (α_j decreases with j), but it minimizes market output if they are substitutes (α_j increases with j). In the latter case, the greater allocative efficiency that arises from the less assortative match in the unraveling equilibrium may outweigh its informational inefficiency. Whether the substitutes or the complements case better describes the market for law clerks is an open empirical question.³⁵

We summarize the foregoing discussion in Corollary 2.

Corollary 2.

- (i) Let Σ denote the set of limits of ϵ -equilibria for $\epsilon \rightarrow 0$. If $x_J(T|T, \dots, T) > 0$, then for any $\sigma, \sigma' \in \Sigma$, and for each j , employer j 's expected payoff in equilibrium σ and in equilibrium σ' are equal.
- (ii) There is no Pareto-relation between the full-information assortative match and the unraveling equilibrium σ .

³⁰ Trivially, no such statement can be made in case of heterogeneous student preferences. All the foregoing results apply equally if students have heterogeneous preferences and the strategy profile σ requires all employers to play e_1 at any subgame hiring time.

³¹ Specifically, by property (iii) of the equilibrium strategy σ_j , an employer must attempt to exit the game whenever doing so is strictly more profitable than the continuation strategy. By right-continuity, therefore, each employer must also exit when doing so is weakly more profitable. Thus, at each point in time, the most desirable remaining employer must exit at the next subgame hiring time. While ϵ -equilibrium does not share this implication, a more desirable employer in an ϵ equilibrium nonetheless has a competitive advantage over less desirable one: Whenever desired, a more desirable employer can exit at the same point in time as a less desirable employer, and thus preempt him.

³² Moreover, if students have heterogeneous preferences, the order in which employers hire students will reflect the preference of whichever student happens to be the most highly skilled at a given equilibrium hiring time.

³³ Formally, in any unraveling equilibrium, employer J obtains $x_J(T|t_1^*, \dots, t_{J-1}^*)$ and in the full information assortative match he obtains $x_J(T|T, \dots, T)$. The latter payoff is lower because $x_J(T|t_1, \dots, t_{J-1})$ is decreasing in $t_1, \dots, t_{J-1}$.

³⁴ This result depends on the assumption that students have similar preferences. For instance, if student preferences across employers were such that from the full information assortative match each employer expected the average skill level, all employers would be worse off in the unraveling equilibrium, in expectation. (Recall that heterogeneous student preferences would leave the equilibrium hiring times unchanged.)

³⁵ The case of substitutes, however, is not merely a theoretical curiosity. Davis (2017) examines the matching of teachers to schools with Teach For America. He examines the change in the matching procedure from a first-offer mechanism, which induced some unraveling, to a deferred-acceptance algorithm. He finds that while teacher retention increased (indicating that the new match outcomes better account for teacher preferences), it lowers performance in the teachers' first year.

- (iii) If student skills and employer desirability are complements, production in a positive assortative match at T exceeds that from any unraveling equilibrium. If student skills and employer desirability are substitutes, production in a positive assortative match at T may fall short of that in the unraveling equilibrium.

4.5. Relation to two-period models

Most existing work on unraveling in matching markets restricts transactions to two points in time. Imposing such a restriction in our model substantially changes equilibrium outcomes. Formally, we impose the restriction that employers can exit the market only at times t_0 and T , for some $t_0 \in [0, T)$.

If t_0 is early, there is little information about students at t_0 , and pure strategy equilibria exist. In any such equilibrium, all employers above some threshold hire at T ; only less desirable employers hire at t_0 . The reason is that given the small amounts of information at t_0 , hiring at that point in time provides an employer with a student he expects to be average. Therefore, hiring at t_0 is profitable only if entering into competition with more desirable employers at T would guarantee a student with expected skills below average.³⁶ By implication, expected equilibrium payoffs are higher for more desirable employers. These mechanisms resemble Halaburda (2010).³⁷

If t_0 is late, then the amount of information available in the first period is sizable. In this case, the game has a matching-pennies structure: less desirable employers prefer to avoid competition with stronger employers, but stronger employers prefer direct competition with less desirable ones.³⁸ Hence, any equilibrium will be in mixed strategies. Moreover, we show that in any such equilibrium, expected equilibrium payoffs are higher for more desirable employers. In this case, our model resembles Du and Livne (2013). Formally, we prove the following proposition, letting π_j denote the skill level employer j expects to obtain in equilibrium.³⁹

Proposition 3.

- (i) If $t_0 = 0$, there exists a pure strategy equilibrium profile, and every such profile is in threshold strategies: there is an employer $\hat{j} \in \{1, \dots, J\}$ such that employers $1, \dots, \hat{j}$ exit at T and all other employers exit at t_0 . Expected equilibrium payoffs satisfy $\pi_1 > \dots > \pi_{\hat{j}-1} \geq \pi_{\hat{j}} = \dots = \pi_J$.
- (ii) If t_0 is sufficiently close to T , then no pure strategy equilibrium exists. In any mixed strategy equilibrium, expected equilibrium payoffs satisfy $\pi_1 > \dots > \pi_J$.

Why does the two-period restriction change the structure of equilibria so dramatically? The market unravels because hiring early allows weaker employers to eschew competition with stronger employers. Two-period models severely restrict this possibility—most employers will have to hire simultaneously with others. Accordingly, many properties that would emerge if all employers hired at the same time (such as more desirable employers obtaining higher payoffs) partially survive in two-period models, but are completely eroded in our continuous-time model, in which employers face far more opportunities to avoid competition.

Crucially, relaxing the two-period restriction also affects policy implications. A common response to unraveling in matching markets is the introduction of a centralized clearinghouse close to the start date of employment (time T). In Halaburda (2010), if a market unravels in spite of an ex-post stable matching mechanism in the second period, there always exists a (pairwise unstable) mechanism that prevents unraveling. In stark contrast, if unraveling occurs in a market described by our model, no clearinghouse run at T will prevent unraveling. The reason is that as soon as two or more employers participate in a clearinghouse, one of them faces a positive probability of not matching with the best available student—regardless of the matching mechanism. Because information arises gradually, and because there is an extended period of time during which matches can be made, such an employer can increase his payoff by hiring the best available student at some time slightly before T .⁴⁰

³⁶ Accordingly, if $x_j(T|T, \dots, T) > 0$, then there will be no unraveling, even though absent the two-period restriction, significant unraveling is predicted. In the case that student skills are independent Brownian motions with $q_i(0) = 0$, $I > 2J$ is a sufficient condition for $x_j(T|T, \dots, T) > 0$.

³⁷ Halaburda (2010) allows students to reject offers. Under this assumption, sufficiently undesirable employers will join the sufficiently desirable employers in hiring late. The reason is that students will reject early offers from sufficiently undesirable employers, hoping for better offers in the second period. If students are sufficiently risk averse (either because they suffer much disutility from remaining unmatched, or because there is a large excess number of students), all early offers will be accepted.

³⁸ If weaker employers hire in the first period, stronger employers will also prefer hiring in the first period in order to avoid being preempted, whereas if weaker employers hire in the second period, stronger employers prefer hiring in the second period in order to reap the informational gains from postponing hiring.

³⁹ In our unrestricted model, we obtain different implications regarding the order of hiring depending on whether we consider subgame perfect equilibrium or the limit of ϵ -equilibria for ϵ approaching zero (see Corollary 1). The two-period restriction renders the model into a finite game. Accordingly, the two equilibrium concepts yield the same solution.

⁴⁰ There is empirical evidence that centralized clearinghouses are not always sufficient to prevent unraveling, even if they implement the ex-post pairwise stable match. The Canadian market for lawyers, for instance, exhibits severe unraveling even though a centralized clearinghouse is in place (Roth and Xing, 1994). Additionally, Wetz et al. (2010) study the market for hospital interns in the year 2007, which was cleared using a pairwise stable matching algorithm,




Fig. 3. The effects of withholding information about students in the case of two employers. In panel (A), the withholding of information starts early so that the first student is hired at $\bar{t}$. In panel (B), the withholding of information starts late. The first student is hired at $\underline{t}$ and so the policy exacerbates unraveling.

This result is due to the gradual arrival of information rather than due to our assumption that students are unable to reject offers. The attempt to prevent unraveling through the mechanisms proposed in Halaburda (2010) will not succeed as they imply blocking pairs—pairs consisting of a student and an employer who would prefer matching to each other over their respective mechanism-assigned match. Because information arrives continuously, these participants can find each other with very high probability and sign a contract at some time shortly before T , thus causing unraveling.

4.6. Information timing policies

In contrast to two-period models, our model allows for a detailed study of policies that alter the timing of the release of information about students. We define an *information timing policy* as a strictly increasing function $\gamma : [0, T] \rightarrow [0, T]$ with $\gamma(0) = 0$ and $\gamma(T) = T$, which we interpret as follows. Information that arises at t is made available to employers at $\gamma(t)$ in the sense that if an employer plays e_1 at time $\gamma(t)$, this employer is assigned the student who was top-ranked at t .

We first examine the case of a continuous function γ . In this case we can construct the set of equilibria as before, by finding points in time where the expected payoff from exiting the market equals the expected payoff from staying in the market—continuity guarantees their existence. Because equilibrium hiring times are defined by the amount of information available rather than by how much time is left until employment starts, the policy has no effect on any employers' payoff, or on total production.

Proposition 4. *If $\gamma : [0, T] \rightarrow [0, T]$ is continuous, then in any subgame perfect equilibrium, students are hired at times $\gamma(t_1^*), \dots, \gamma(t_1^*)$, and the equilibrium payoff of each employer is the same as without the policy.*

Hence, we consider informational discontinuities. We focus on the case in which information arriving between some time $\underline{t}$ and some later time $\bar{t} < T$ is withheld, and is released at $\bar{t}$. Formally, $\gamma(t) = t$ if $t \notin [\underline{t}, \bar{t}]$, and $\gamma(t) = \underline{t}$ if $t \in [\underline{t}, \bar{t}]$. Hence, if employers play actions $e \in E$ during $[\underline{t}, \bar{t}]$, they are assigned the students they would have been assigned if they all had played the these actions at time $\underline{t}$ —except that it is the earlier judges rather than the more desirable judges who are then assigned the better students. Therefore, any employer can do at least as well by exiting at $\underline{t}$ as by exiting at some time $t \in (\underline{t}, \bar{t})$ within the withholding period.

The effect of the policy depends on the values of $\underline{t}$ and $\bar{t}$. Plainly, the policy does not eliminate any equilibria if $[\underline{t}, \bar{t}]$ does not contain any equilibrium hiring time. Otherwise, the policy may mitigate or exacerbate unraveling, as is best explained in the case of two employers.

Consider the graphs in Fig. 3, which parallel Fig. 2. In both panels, information is released at the same time $\bar{t}$. If neither employer hires a student before $\bar{t}$, then exiting at $\bar{t}$ is strictly better than remaining in the market ($x_1(\bar{t}) > x_2(T|\bar{t})$). Thus, both employers will attempt to exit at that time. The more desirable employer 1 will be successful and will obtain a student of expected skill level $x_1(\bar{t})$. The less desirable employer 2 will hire the best remaining student at T and will expect a student with lower skill level $x_2(T|\bar{t})$. Whether employer 2 can do better by exiting before $\bar{t}$ depends on the start time $\underline{t}$ of the information withholding period.

In panel (A), $\underline{t}$ is sufficiently early that $x_1(\underline{t}) < x_2(T|\bar{t})$. Hence, employer 2 prefers hiring the best remaining student at T to exiting the market at any time at or before $\bar{t}$. In this case, the policy mitigates unraveling. In panel (B), $\underline{t}$ is late, so that $x_1(\underline{t}) > x_2(T|\bar{t})$. In this case, employer 2 prefers exiting at $\underline{t}$ to hiring the best of the remaining students at T . Since $\underline{t} < t_1^*$, however, the policy exacerbates unraveling. Exacerbation occurs because the information policy places employers in direct

and find that 15.7% of the postgraduate positions were offered outside the match. Relatedly, Haruvy et al. (2006) conclude that unraveling in the market for law clerks is not due to lack of pairwise stability, but due to the ease with which binding contracts are forged.

competition with each other at time $\bar{t}$. While more desirable employers benefit from the policy, less desirable employers lose. Therefore, the policy provides additional incentives for less desirable employers to hire early.⁴¹

These arguments generalize to the case of J employers, as shown in Proposition 5 below.⁴² In any case, if the informational discontinuity delays the hiring activity of multiple employers, the information release will lead to a rush at time $\bar{t}$. In practice, rushes may lead to market congestion, which can be addressed with policies such as centralized clearinghouses.

Proposition 5. Consider the information timing policy γ with $\gamma(t) = t$ if $t \notin [\underline{t}, \bar{t}]$, and $\gamma(t) = \underline{t}$ if $t \in [\underline{t}, \bar{t}]$.

- (i) If $0 < x_J(T|T, \dots, T)$, then, for sufficiently low $\underline{t}$, no students are hired before $\bar{t}$.⁴³
- (ii) There are values $\underline{t}$ and $\bar{t}$ with $\underline{t} < \bar{t}$ such that under policy γ the first student is hired strictly before t_1^* .

Historically, a policy of withholding information about students had successfully curtailed unraveling in the market for medical residency internships in 1945 (Roth, 1984). The policy was implemented boldly. Medical colleges agreed to release neither academic transcripts nor letters of recommendation before an agreed-upon date.⁴⁴ In other markets, it may be practically difficult to withhold information for the required amount of time, and to prevent leakage of information. Indeed, such an information withholding policy faced very limited success in the market for law clerks (Wald, 1990).

In addition to withholding information, informational discontinuities may be introduced by limiting the production of information about students, for instance by switching from letter grades to a pass/fail system.⁴⁵ Another possibility is a nationwide test designed to provide a sudden surge in information at a pre-specified point in time about students' abilities.

5. Conclusion

Unraveling plagues many entry-level labor markets. To better understand the phenomenon, we have introduced a model of a market that may take place over an extended period of time during which information about students arrives gradually. We have examined the equilibrium properties of this market if unraveling is allowed to occur without intervention.

We find that equilibrium hiring times are widely dispersed, consistent with empirical observations. The first student is hired so early that his employer would be better off if all employers waited until all information is available and then hired a random student from the pool of those whose skills are high enough to be employed at all. Competitive pressure intensifies unraveling; students are hired even earlier as the number of employers grows.

In equilibrium, unraveling completely equalizes the expected student skills across employers. Nonetheless, employers who, in the full information assortative match, would be matched to one of the less highly skilled students prefer the unraveled market—unraveling is not Pareto-inefficient relative to the full information assortative match. In fact, if student skills and employer desirability are substitutes, unraveling may even lead to higher market output than the full information assortative match.

Our model contrasts with, and informs, previous models of unraveling, which are set in two periods. We show that they can be understood as extremes on a continuum in which the temporal position of first period determines the qualitative structure of the equilibrium. Our model also highlights the sensitivity of the previous research's mechanism design implications on the two-period assumption.

Finally, our extended-market formulation allows us to study information timing policies (which two-period models, by design, cannot). We have identified informational discontinuities as a necessary condition for their effectiveness.

Further research might profitably extend the study of unraveling over time by focusing on students. For instance, our model abstracts from human capital investment by students, which, in principle, might depend on the efficiency of the match and on the timing of the market.⁴⁶ It would also be interesting to test the main prediction of our model empirically, namely that the date at which a student is hired does not predict his productivity over the course of employment. Perhaps the most significant contribution could be made by studying unraveling over time when workers' wages are up for negotiation.

⁴¹ These incentives are weaker if students' preferences are more heterogeneous.

⁴² The proposition is a direct implication of the subgame-perfect strategy profile constructed in Section 4.2.

⁴³ The condition $0 < x_J(T|T, \dots, T)$ is necessary. If it is not satisfied, the least desirable employer prefers hiring at 0 to hiring at T even if information is withheld in the interval $[0, T)$.

⁴⁴ The matching that resulted after the release of information was not pairwise stable, which led to a different set of issues that ultimately led to the adoption of what is known today as the National Resident Matching Program. See Roth (1984), and Roth and Peranson (1999).

⁴⁵ Yale Law School has used pass/fail grading for first year students in those years in which unraveling was particularly extreme (Al Roth, personal communication). Ostrovsky and Schwarz (2010) show how limiting information about students in a particular fashion may prevent unraveling.

⁴⁶ In the only related empirical study, using the market for college graduates in Japan, Okudaira (2017) finds that a guideline revision that successfully delayed the timing of job searches had no positive effect on students' human capital investment.

Acknowledgments

We thank Muriel Niederle, Al Roth, and Andy Skrzypacz for their invaluable advice, and Mohammed Akbarpour, Tim Bresnahan, Itay Fainmesser, John Hatfield, Paul J. Healy, Matthew Mitchell, Joshua Mollner, Mikhail Panov, Florian Scheuer, and Pavel Zryumov for helpful comments and suggestions. Financial support from Stanford University's Department of Economics and Graduate School of Business are gratefully acknowledged.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.geb.2020.02.009>.

References

- Administrative Office of the United States Courts, 2013. Key OSCAR statistics. https://oscar.uscourts.gov/_assets/FY_2013_Key_OSCAR_Statistics.
- Alden, William, 2014. A mad scramble for young bankers. *N.Y. Times*.
- Anderson, Axel, Smith, Lones, Park, Andreas, 2017. Rushes in large timing games. *Econometrica* 85 (3), 871–913.
- Ante, Spencer E., 2012. Revenge of the nerds: tech firms scour college campuses for talent. *Wall St. J.*
- Avery, Christopher, Jolls, Christine, Posner, Richard A., Roth, Alvin E., 2001. The market for federal judicial law clerks. *Univ. Chic. Law Rev.* 68 (3), 793–902.
- Avery, Christopher, Jolls, Christine, Posner, Richard A., Roth, Alvin E., 2007. The new market for federal judicial law clerks. *Univ. Chic. Law Rev.* 74 (2), 447–486.
- Bonica, Adam, Chilton, Adam S., Goldin, Jacob, Rozema, Kyle, Sen, Maya, 2017. The political ideologies of law clerks. *Am. Law Econ. Rev.* 19 (1), 96–128.
- Collins, Judith N., 2012. Salaries for new lawyers: an update on where we are and how we got here. *Nat. Assoc. Law Place. Bull.*, 1–4.
- Damiano, Ettore, Hao, Li, Suen, Wing, 2005. Unraveling of dynamic sorting. *Rev. Econ. Stud.* 72, 1057–1076.
- Davis, Jonathan M.V., 2017. The Short and Long Run Impacts of Centralized Clearinghouses: Evidence from Matching Teach for America Teachers to Schools. Unpublished.
- Du, Songzi, Livne, Yair, 2013. Rigidity of Transfers and Unraveling in Matching Markets. Mimeo.
- Echenique, Frederico, Pereyra, Juan Sebastian, 2016. Strategic complementarities and unraveling in matching markets. *Theor. Econ.* 11.
- Fainmesser, Itay P., 2013. Social networks and unraveling in labor markets. *J. Econ. Theory* 131 (148), 64–103.
- Ginsburg, Tom, Wolf, Jeffrey A., 2003. The market for elite law firm associates. *Fla. State Univ. Law Rev.* 31, 909.
- Halaburda, Hanna, 2010. Unraveling in two-sided matching markets and similarities of preferences. *Games Econ. Behav.* 69, 365–393.
- Haruvy, Ernan, Roth, Alvin E., Utku Unver, M., 2006. The dynamics of law clerk matching: an experimental and computational investigation of proposals for reform of the market. *J. Econ. Dyn. Control* 30, 457–486.
- Hautmann, Thomas, 2013. 20-months-old baby becomes world's youngest pro footballer. *Fox Soccer Blog*.
- Ishii, Yuhta, Oery, Aniko, Vigier, Adrien, 2018. Competing for Talent. Cowles Foundation Discussion Paper, 3019.
- Kozinski, Alex, 1991. Confessions of a bad apple. *Yale Law J.* 100 (6), 1707–1730.
- Kueppers, Courtney, 2016. A disciplinary association aims to rein in a 'chaotic' hiring calendar. *Chron. High. Educ.*
- Li, Hao, Rosen, Sherwin, 1998. Unraveling in matching markets. *Am. Econ. Rev.* 88 (3), 371–387.
- Li, Hao, Suen, Wing, 2000. Risk sharing, sorting, and early contracting. *J. Polit. Econ.* 108 (5), 1058–1091.
- Li, Hao, Suen, Wing, 2004. Self-fulfilling early-contracting rush. *Int. Econ. Rev.* 45 (1), 301–324.
- McDaniel, Katherine, Drummond, Dylan, Walia, Mani, 12/08/2013. Clerkship notification blog. <http://cnb09-10.blogspot.com>.
- Niederle, Muriel, Roth, Alvin E., Utku Unver, M., 2013. Unraveling results from comparative demand and supply: an experimental investigation. *Games* 4, 243–282.
- Niederle, Muriel, Proctor, Deborah D., Roth, Alvin E., 2006. What will be needed for the new gastroenterology fellowship match to succeed? *Gastroenterology* 130, 218–224.
- Okudaira, Hiroko, 2017. Regulating the Timing of Job Search: Evidence from New College Graduates in Japan. Unpublished.
- Online System for Clerkship Application and Review, 2013. OSCAR de-links from the hiring plan. https://oscar.uscourts.gov/blog_post/_1/21/OSCAR_De-links_from_the_Hiring_Plan.
- Ostrovsky, Michael, Schwarz, Michael, 2010. Information disclosure and unraveling in matching markets. *Am. Econ. J. Microecon.* 2 (2), 34–63.
- Popper, Nathaniel, 2014. Committing to play for a college, then starting 9th grade. *N.Y. Times*.
- Roth, Alvin E., 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *J. Polit. Econ.* 92 (6), 991–1016.
- Roth, Alvin E., Peranson, Elliot, 1999. The redesign of the matching market for American physicians: some engineering aspects of economic design. *Am. Econ. Rev.* 89 (4), 748–780.
- Roth, Alvin E., Xing, Xiaolin, 1994. Jumping the gun: imperfections and institutions related to the timing of market transactions. *Am. Econ. Rev.* 84 (4), 992–1044.
- Strauss, Debra M., 2015. Top ten tips to the judicial clerkship interview. <http://abovethelaw.com/career-files/ten-tips-to-the-judicial-clerkship-interview/>.
- Suen, Wing, 2000. A competitive theory of equilibrium and disequilibrium unraveling in a two-sided matching. *Rand J. Econ.* 31 (1), 101–120.
- Tobias, Carl, 2013. Tips for capturing 2014 federal court clerkships. *Vanderbilt Law Review En Banc* 66, 145–152.
- Vidaurre, Jorge, Campbell, Julia, 2017. The case for an epilepsy and clinical neurophysiology match. *Pediatr. Neurol.* 72, 5–6.
- Vohra, Akhil, 2019. Unraveling in the Presence of a Secondary Market. Unpublished.
- Wald, Patricia M., 1990. Selecting law clerks. *Mich. Law Rev.* 89 (1), 152–163.
- Wetz, Robert V., Seelig, Charles B., Khouri, Georges, Weiserbs, Kera F., 2010. Out-of-match residency offers: the possible extent and implications of prematching in graduate medical education. *J. Grad. Med. Educ.* 2 (3), 327–333.
- www.top-law-schools.com, 2013–2015. Clerkship application thread. <http://www.top-law-schools.com/forums/viewtopic.php?f=34&t=202465>, <http://www.top-law-schools.com/forums/viewtopic.php?t=214750>, <http://www.top-law-schools.com/forums/viewtopic.php?f=34&t=248994>.
- Yale Law School Career Development Office, 2008. Judicial clerkships in the U.S. http://www.law.yale.edu/documents/pdf/CDO_Public/cdo-JCpublic2008.pdf.
- Yu, Felix, 2017. HEP Postdoc Rumor Mill. <https://sites.google.com/site/postdocrumor/>.