Unraveling and Judge Productivity in the Market for Federal Judicial Law Clerks: Evidence and Proposal*

Daniel L. Chen† YingHua He‡ Takuro Yamashita§

March 24, 2024

(preliminary and incomplete)

First draft: June, 2019.

Abstract

We study the market for federal judicial law clerks, which suffers from market unraveling for many years with multiple unsuccessful reforms. In the empirical part, based on a unique dataset on judicial production, we find that (i) preventing unraveling would benefit overall judicial production, but (ii) its effect is heterogeneous: only high-productivity judges would be better off, while low-productivity judges could be much less or even worse off. Motivated by this empirical finding, we develop a theoretical model of hiring, and propose a novel form of reform, which explicitly allows early hiring of some judges in an appropriately regulated manner.

1 Introduction

Consider a hiring market of new graduates. From the informational perspective, it is socially desirable if hiring starts after students mostly finish their study so that they

*We thank Yusuke Yasuda, Takuma Kamada, Masaki Aoyagi and seminar participants at Osaka for valuable comments and discussion. Daniel Chen acknowledges IAST funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EUR-0010, European Research Council (Grant No. 614708), and the research foundation TSE-Partnership. Takuro Yamashita acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 714693), ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program), and JSPS KAKENHI (JP23K01311).

†Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France. daniel.chen@iast.fr

‡Rice University, Houston, Texas, USA. yinghua.he@rice.edu

§Osaka School of International Public Policy, Osaka University, Osaka, Japan. yamashita.takuro.osipp@osaka-u.ac.jp

are endowed with informative certificates and transcripts. In some markets, however, potential employers move early and contact much younger students, a phenomenon called *unravelling*.[1]

One of the main reasons why unravelling happens is due to employers' competition: If an employer anticipates that, by waiting until students finish their study, some other employers are likely to get all the "good-quality" students, then this employer may want to contact younger students whose quality is still relatively uncertain. There is certainly a risk that this student's quality turns out to not be satisfactory, but a small chance that this student may be good can potentially be enough to take that risk.[2]

Unravelling is socially undesirable, especially from the informational-efficiency viewpoint. Compared to the "ideal" case where hiring happens after all the workers' quality information is publicly revealed, in an unravelled case, some early-moving employers may hire a worker whose quality is bad (because at the time of (early) hiring the quality was not yet known), and instead, some good-quality workers may be left unemployed.

Given that unravelling is prevalent in many hiring markets despite its problems,[3] it is important to better understand this phenomenon: In which markets is unravelling more relevant? How severe is the informational inefficiency due to unravelling? Are there effective policies which prevent / mitigate unravelling?

In this paper, we study this unravelling problem, *both* empirically and theoretically. In the empirical part (Section 2), we look at Federal Judicial Clerkship Market, where judges of Federal Courts of Appeal hire top new graduates of law schools as law clerks. As explained later more in detail, unravelling is a long-standing problem in this market. Some judges contact students at the beginning of their second year, and multiple attempts of reform agreements all failed.

Using the variation of the hiring timing in the reform years and other years, based on the unique dataset about published decisions in US Circuit Courts, we estimate the size of the effect of unravelling on judges' judicial production. We find that the effect is overall positive and significant. That is, preventing unravelling would be valuable in improving judges' judicial production (which we interpret as the consequence of improved information efficiency in hiring). We also uncover heterogeneity of that effect across judges: those categorized as high-productivity judges would be more positively affected if unravelling is prevented, while those categorized as low-productivity could be less positively or even negatively affected. This heterogeneity may explain why the past reform agreements failed: Although preventing unravelling makes some judges better off, other judges are *hurt* by that, that is, it is not Pareto improving.

The theoretical part (Section 3) builds a game model where judges decide their hiring timing. We show that, our empirical finding that preventing unravelling is

[1] See Roth and Xing (1994) for a more detailed and generic description of an unravelling problem.

[2] Of course, for this logic to work, additional conditions must be satisfied. For example, this younger student who gets an early offer may prefer waiting rather than taking this early offer, if he thinks that he will have a high chance of being hired by better employers by doing so. This is why unravelling is more relevant in some hiring markets but less in others. In Section 3, we investigate those conditions theoretically.

[3] See Roth (2018).

overall improving judicial productivity but hurting lower-productivity judges can well be explained in a simple model where (i) talented students are rare, (ii) the number of clerkship positions is much smaller than the number of applicants, and (iii) agents in each side has the same preference over agents in the other side. These conditions basically mean that the market is competitive in *both* sides. As also pointed out in the literature,[4] this two-sided competitiveness is important in inducing unravelling: on the judges' side, less popular judges would have an incentive to move early at the risk of informational inefficiency, as it is very unlikely that they get talented students by waiting; and on the students' side, accepting such early proposals would be better than waiting, as it is very unlikely that they get any offer by waiting.

The novelty of our paper relative to the theoretical literature of unravelling is that we obtain some insights for a successful reform in preventing unravelling based on the repeated nature of this clerkship market. Indeed, the set of hiring judges are mostly fixed over time, while the law clerkship in US Circuit Courts is usually once in a lifetime for law-school students, making it appropriate to model it as a repeated game among judges with "short-lived" students. We obtain two key results: first, even if judges are arbitrarily patient, the "first-best" informational efficiency where every judge waits until students finish their study is not implementable. That is, unravelling is still a relevant problem even in the repeated-game framework (consistent with the reality where all past reform attempts failed). Nevertheless, the *second-best informational efficiency* may be attainable by allowing some judges' early offers in an appropriately regulated way. This novel idea of permitting some early offers has not yet been attempted in practice,[5] but we believe that it can be helpful in designing a *sustainable* rule in hiring and mitigating uncontrolled unravelling, and thereby in improving informational efficiency.

1.1 Literature

The problem of unravelling in hiring has been discussed in the literature of two-sided matching.[6] Its prevalence in many real hiring markets has been discussed, for example, by Roth and Xing (1994) and Roth (2018) (and the references therein). As one of the representative instances, federal judicial clerkship markets has been studied by Avery et al. (2001), Avery et al. (2007), and Fréchette et al. (2007). Although it is natural to think that unravelling is socially undesirable in view of the informational efficiency, no prior research has examined its impact quantitatively. Our empirical contribution

[4]For example, Li and Rosen (1998), Li and Suen (2000) and Suen (2000) argue that workers' risk aversion can result in unravelling. Halaburda (2010) shows that similarity of preference rankings in each side can imply unravelling.

[5]Asymmetric treatment of employers in the hiring timing may be criticized as "unfair". However, such an asymmetric hiring rule exists in other industries, such as in Major League Baseball Draft and National Basketball Association Draft. Though quite different context, it suggests the general principle that the fairness concern could be mitigated with appropriate "rationale". Successful cases in other industries could be insightful in its practical implementation.

[6]See Roth (2018) for a more complete list of the papers about unravelling in matching markets.

is to quantify the cost of unravelling in judicial production, and moreover, to find its heterogeneity among judges. The former implies that preventing unravelling would be socially valuable, while the latter points to a potential difficulty in sustaining the informational efficiency.

On the theoretical side, several papers study the source of unravelling, its implied inefficiency, and some possible ideas to prevent it. For example, Niederle and Roth (2009) consider the role of exploding offers, Li and Rosen (1998), Li and Suen (2000) and Suen (2000) point to workers' risk aversion, Damiano et al. (2005) concern search cost, and Halaburda (2010) studies the effect of similarity of preference rankings in each side.[7] Our theoretical model builds mostly on those prior research, highlighting the importance of competitiveness of the market in both sides and participants' similarity in their preference. The novelty of our work relative to the theoretical literature is that we provide some insights for a successful reform based on the repeated nature of the judicial clerkship market.

Although it is new in the hiring context, the idea of permitting some "selfishness" in a regulated way in a repeated-game equilibrium has been explored in the context of sustaining bidders' collusion in repeated auction. For example, Aoyagi (2003) and Skrzypacz and Hopenhayn (2004) study how "bid rotation" may be designed in a history-dependent way in order to better sustain bidders' collusion. Our proposed mechanism has some similarity to the idea of bid rotation in that who should be allowed to make an early offer should be determined in a history-dependent way, in order to better sustain the "second-best" informational efficiency. On the other hand, it is worthwhile to note that, while high-enough patience is typically enough to sustain the "first-best" collusion in the context of repeated auction based on Fudenberg et al. (1994), in our hiring context, the first-best informational efficiency is not implementable even if the players are arbitrarily patient.

2 Empirical Analysis: Federal Judicial Clerkship Market

2.1 Institutional Background and Data

Every year, judges of Federal Courts of Appeal hire top new graduates of law schools as law clerks. These law clerks assist judges on a range of tasks, such as in research issues, drafting opinions, and making legal determinations. Clerkship positions are highly sought after, leading to prestigious professional opportunities, which naturally makes the market very competitive. To give an idea, there are roughly around 170 judges in our analysis, each of whom opens four positions of judicial clerks each year, and hence, there are several hundreds of positions in total. There are many more

[7] Relatedly, Chiu and Weng (2019) show that less popular schools may prefer certain modification to the standard deferred-acceptance mechanism so that they can "steal" good applicants who otherwise would be matched to more popular schools.

applicants (law school students) applying for these positions, so that each position receives a hundred (least popular) to thousands (most popular) of applications. As a consequence, only a small fraction of applicants end up with clerkship positions.

The competition is also harsh on the judges' side. Some positions are much more popular than others, and less-popular judges are more desperate in getting potentially talented students. Due to this two-sided competitiveness of the market, it is often (mostly) unravelled. While the National Federal Judges Law Clerk Hiring Plan recommends when judges may receive applications and when they may contact, interview, and hire clerks, generally many do not follow this schedule and hire law students quite early, in some time periods, as early as right after the first year of law school. Sometimes judges ask a candidate to provide a prompt answer to the question, "Will you accept an offer?" even prior to scheduling an interview. It goes without saying that job offers are expected to be accepted on the spot. To defer would be a sign of disrespect that can stigmatize the year-long relationship.[8]

Unravelling has been well-recognized as a major and long-standing problem in this market. At least four reform agreements were attempted in the past: 1983, 1986, 1990, and 2005. Although the concrete contents of the agreements differ, they share the common goal of mitigating unravelling by uniformly regulating the earliest date at which law students could be hired.[9] However, strikingly, these past reforms all failed, by the market's promptly unravelling. Even at some of the reform years, some Circuits were noted already as "cheating".[10] Then, it is followed by many others' deviations from the agreement in later years.

Even though unravelling has long been recognized as a serious problem in this market, not much has been known about the actual magnitude of the social cost of unravelling. Our main empirical contribution is to quantify this cost of unravelling on judges' judicial production, based on a unique dataset on all 380,000 published decisions (over a million judge votes) in U.S. Circuit Courts since 1891. We have the full citation network between the cases. We also have detailed metadata for each case, from which we use in particular the court, publication date, and authoring judge. Sourced from Bloomberg Law, these cases are verified against other existing datasets, such as the Songer Database, the Federal Judicial Center's Administrative Office of the U.S. Courts

[8] Many other anecdotal evidences of unravelling can be found in other papers in the literature, such as in Roth (2018). Avery et al. (2001) and Avery et al. (2007) conduct large-scale surveys both for judges and students, and observe many instances of early exploding offers.

[9] More specifically, Reform 1 occurred in March 1983. The Judicial Conference requested that judges not consider applications before September 15 of the students' third year of law school. We consider the reform as being effective for clerks starting in Summer/Fall 1983. Reform 2 occurred in 1986. For the 1986 season, federal appellate judges were asked to not consider student applications before April 1. We consider offers being made in 1986, for work that commenced in Fall 1987 and opinions published in 1988. Reform 3 occurred in 1990. Commencing in 1990, no job offers were allowed before May 1st of the applicant's second year. We consider the reform as being effective for clerks starting work in Summer/Fall 1991 and opinions published in 1991 or 1992. Reform 4 occurred in 2005. In Fall 2005, Sept 15 is the date for scheduling of interviews and Sept 22 the date for offers. We consider the reform as being effective for clerks starting work in Summer/Fall 2005.

[10] See Avery et al. (2001), Avery et al. (2007), and Fréchette et al. (2007).

dataset, and Lexis Nexis. Each case in our dataset is linked to whether it is appealed to the U.S. Supreme Court as well as whether it is reversed on appeal. Additionally, we have access to extensive biographical information on federal circuit judges from the Federal Judicial Center. This aspect of the dataset includes the appointment date with which we construct a measure for judicial experience.

The empirical findings suggest that past reforms are good for average judges but hurt some. This implies that those reforms are not Pareto improving. We present judge-level panel regressions, which aggregates the published decisions to the author-year level. A year is defined as September to August, taking into account that a clerk can arrive between June and October.

2.2 Overall Effect of Unravelling on Judicial Production

In order to see the effect of unravelling on judicial production, we basically compare the judicial production at the year where the reform is in effect, and that of the other "non-reform" years. Some anecdotes suggest that some judges may deviate even in the reform years, while others may follow the agreement for multiple years since the reform implementation. However, it is reasonable to assume that more judges conform the agreement in the reform years than in the non-reform years, and if so, our estimate provides a *lower bound* of the overall effect of the reform.

Regarding the measure of judicial production, we consider four variables: the number of judicial case publications of each judge, some citation measures of those publications, and whether those cases are reversed in the Supreme Court. The first measure is a sort of a "quantity" measure, while the others are more of "quality" measures. The case publications are produced roughly as follows.[11] For each case brought to a Circuit Court, judges are randomly assigned as a panel of three within the corresponding Circuit. If there is an oral argument, they meet right afterwards and have a preliminary vote of 3-0 or 2-1. Then, the panel writes the opinion to the case, and decides whether to *publish* it (in *Federal Reporter*). The Judicial Conference of the United States decided in 1964 that Circuit Courts should publish only opinions that are "of general precedential value". Thus, much like in academia, we interpret the number of judicial case publications of a judge as a measure of the amount of his effective labor (and better law clerks could be helpful in increasing the number of case publications either by directly helping the publication process, or indirectly by working on other tasks and hence letting the judge more focused on the publication). This "quantity" measure is complemented by other "quality" measures: Regrading the citation measures, we consider "within" citations (which are from other publications of the same Circuit) and "outside" citations (which are from other Circuits). The case reversion is when the decision of the Circuit Court is reversed by that of the Supreme Court (though quite rare).

We regress each of these variables on a number of control variables, including the indicator variable for the reform year. More specifically, judge i's production in circuit

[11] See also Cohen (2023) for the case publications.

Table 1: Overall Effect

Dependent variable (yearly)	Cases Published	Citations (within)	Citations (outside)	Citations (total)	Cases Reversed
	(1)	(2)	(3)	(4)	(5)
Reform	0.606***	7.671*	-3.263*	4.407	0.000389
	(0.231)	(4.550)	(1.893)	(6.032)	(0.0102)
Mean	13.9	180	79.7	259	0.123
N	7853	7853	7853	7853	7853
R-sq	0.756	0.689	0.689	0.702	0.202

Notes: Dependent variable calculated as the yearly total for cases worked on during a market year (September to August). Reversed refers to cases worked on in a year eventually reversed by the Supreme Court. Standard errors clustered at the judge level are in parentheses. ***: significant at the 1% level, *: significant at the 10% level.

c in year t is given by:

$$q_{i,c,t} = \beta \text{Reform}_t + \alpha_i + \text{OtherControls}_{i,c,t} + \epsilon_{i,c,t},$$

where $q_{i,c,t}$ is one of the productivity measures discussed above; Reform_t is a binary dummy which equals to one if year t is a reform year; α_i is judge i's fixed effect, and $\text{OtherControls}_{i,c,t}$ include i's years of experience and its square, Circuit-specific time trends, and case types handled by judge i in year t. The coefficient β captures the overall effect of a reform. Aggregating the decision-level data to judge-level, the sample size becomes 7853 (see Table 1).

The OLS regression results are summarized in Table 1.

In the reform year, the number of case publications in the reform year is on average higher than that in the non-reform year by 0.606. This corresponds to 4% of the average production across judges (13.9 cases). We interpret this as the effect of mitigating unravelling and hence improving the informational efficiency in hiring. We also observe positive but insignificant effects on within-citations and total citations, and also on the number of reversed cases (and a negative effect on outside-citations). The reform has no significant effect on the quality of case publications in terms of total citations and cases reversed. To understand these effects better, we now consider heterogeneity in the effect of reforms.

2.3 Heterogeneity in the Effect of Unravelling

The previous regression concerns the overall effect of unravelling on judicial production. Here, we observe that the effects are quite heterogeneous across different types of judges, in particular, across different productivity categories of them. To see this, we categorize each judge either as a "high-productivity" judge or a "low-productiviety" judge, as follows: For each reform r (recall there have been four reforms), we use data on the three years prior to that reform year (not included), and estimate judge fixed effects from a regression controlling for years of experience, experience squared, and circuit fixed effects. Then we define above-median judges as "high productivity for that reform r". Hence, the identification of high and low productivity judges vary

Table 2: Effect of Reform by Judge Productivity

Dependent variable (yearly)	Cases Published	Citations (within)	Citations (outside)	Citations (total)	Cases Reversed
	(1)	(2)	(3)	(4)	(5)
Reform x Low Productivity	-3.916***	-31.75***	-40.14***	-71.89***	-0.0753***
	(0.498)	(9.745)	(4.671)	(13.58)	(0.0245)
Reform x High Productivity	4.492***	71.17***	7.687	78.86***	-0.00860
	(0.621)	(12.64)	(5.951)	(17.24)	(0.0349)
Mean	18.3	253	120	373	0.184
N	1939	1939	1939	1939	1939
R-sq	0.836	0.773	0.765	0.783	0.334

Notes: Dependent variable calculated as the yearly total for cases worked on during a market year (September to August). Reversed refers to cases worked on in a year eventually reversed by the Supreme Court. Standard errors clustered at the judge level are in parentheses. ***: siginficant at the 1% level. **: significant at the 5% level.

by reform r, although whether a judge is classified as high productivity is positively correlated across reforms.

Now judge i's production in circuit c in year t is given by:

$$\begin{aligned} q_{i,c,t} &= \beta_{low}(\text{Reform}_t \times \text{Low}_{ir}) + \beta_{high}(\text{Reform}_t \times \text{High}_{ir}) \\ &\quad +\alpha_i + \text{OtherControls}_{i,c,t} + \epsilon_{i,c,t}. \end{aligned}$$

That is, compared to the equation in the last section, the effect of a reform is allowed to be different between high-productivity judges and low-productivity judges. In this regression, we limit the analysis to three years prior to and after each reform, reducing the sample size to 1939 (see Table 2). Standard errors are clustered at the judge level.

The OLS regression results are summarized in Table 2.

The results are striking. First, the effect of a reform is positive and significant for high-productivity judges. Moreover, this effect is much stronger than the overall effect: The number of case publications increases by 4.492, which is 24.5% of the average production across judges (18.3 cases). The within-citation and total citation measures are also positive and now significant at the 1% level (recall that they are much less statistically significant in the last section). Conversely, the effect of a reform is rather *negative* for low-productivity judges: The number of case publications decreases by 3.916, which is 21.4% of the average production across judges. The outside-citation measure also decreases significantly (at the 1% level), and the other citation measures are also significant. We find significant effects on the number of cases reversed in the Supreme Court.

These patterns explain well why past reforms were unsuccessful. Even though a reform tends to improve the overall productivity of judges, it actually benefits a particular category of judges (high-productivity judges) significantly, while it rather *hurts* the other category (low-productivity judges). This unfair effect of a reform would naturally make those disadvantaged judges less willing to obey the rule. By deviating from the agreement, those judges could potentially hire talented law clerks, at the cost of informational efficiency.

3 Theory

The purpose of this section is two-fold. First, we provide a game model of judges' competing for hiring talented students. We show that, under certain conditions corresponding to both sides' competitiveness and preference similarity, unravelling is unavoidable in the equilibrium of the game. The conditions provide some insights as to which markets are more vulnerable to unravelling than others, in line with the existing insights in the literature.

Next, based on the property that this hiring game is *repeated* every year with a mostly fixed set of judges and a varying set of students across time, we consider a repeated game model of long-lived judges with short-lived students. We show two main results of this section: First, even in this repeated-game perspective, unravelling is still a relevant problem. As opposed to the standard "folk-theorem" type of results in the applied repeated-game literature, in our setting, the *first-best* informational efficiency is impossible even if judges are arbitrarily patient. Second, nevertheless, the repeated-game perspective introduces a new possibility of sustaining the *second-best* informational efficiency, which can be much better than mere repetition of static unravelling equilibrium. This second-best equilibrium construction suggests some insights for a sustainable agreement for mitigating unravelling.

3.1 Baseline Model

In what follows, we use the generic terminology of "firms" and "workers" instead of judges and law-school students, so that it is easier to understand our model in the context of general labor markets.

Let J be the number of firms and I be the number of workers, where $J < I$. Matching is one-to-one: each firm has one vacant position, and each worker can work only for one firm.[12] There is no monetary transfer, so that it is a non-transferable-utility environment.[13]

We assume simple preferences on both sides. First, every worker $i = 1, \ldots, I$ has the same preference over the firms: $u(1) > u(2) > \ldots > u(J) > \underline{u}$, where $u(j)$ denotes a worker's payoff of being hired by firm j, and $\underline{u}$ denotes his payoff of being unmatched.[14]

Every worker i is either talented ($\theta_i = 1$) or not ($\theta_i = 0$), which is unknown at the beginning of the game (the information structure is explained more in detail in the next paragraph). Every firm $j = 1, \ldots, J$ has the same preference over the workers:

[12] This is a simplifying assumption relative to the reality where matching is often many-to-one. We believe that some of our insights would be robust in those more general settings (under appropriate assumptions), but it is left for future research.

[13] This is consistent with the judicial clerkship market, where wages are fixed. More generally, our results here would be insightful for labor markets where wage bargaining is not so flexible.

[14] The assumption that being unmatched yields the lowest payoff is consistent with the judicial clerkship market, where any clerkship position in US Circuit Court judges is (much) more preferable than no clerkship.

let each firm j's payoff be θ_i if he hires worker i. Let $\underline{v} < 0$ be the firm's payoff of being unmatched, though this does not play much role in the following analysis.

As mentioned above, θ_i is unknown to anyone including i himself at the beginning of the game. Instead, a noisy signal about θ_i, denoted by $s_i \in \{L, H\}$ with $0 < L < H < 1$, is publicly available at Stage 1 (e.g., s_i corresponds to a law-school student i's transcript at the end of his first year of study). Only at Stage 2, θ_i becomes publicly known. Let the unconditional probability of $s_i = H$ be $p \in (0, 1)$ (and $1 - p$ for $s_i = L$), and the conditional probability of $\theta_i = 1$ given s_i be s_i (and $1 - s_i$ for $\theta_i = 0$). Thus, the unconditional probability that $\theta_i = 1$ is $q \equiv pH + (1 - p)L$. Assume that $L = \gamma p < H$ for some $\gamma > 0$, which implies $L = O(p)$ and $q = O(p)$. We also assume mutual independence across workers so that each i's joint distribution of (θ_i, s_i) does not vary with $(\theta_{-i}, s_{-i}) = (\theta_{i'}, s_{i'})_{i' \neq i}$.

3.2 Hiring Game

Here, we analyze the following simple hiring game. At Stage 1, only $s = (s_1, \ldots, s_I)$ is publicly available. Each firm simultaneously decides whether to make an offer at this stage, and if so, to whom. We assume that each firm can make only one offer. Thus, if a firm makes it at this stage, he cannot make any offer at $t = 2$. This reflects the property that making an offer is costly: search, negotiation with the candidate or within the firm, etc.

Each worker who is offered at $t = 1$ decides which one to take (or none). Any matched worker and firm leave the game.

At Stage 2, now $\theta = (\theta_1, \ldots, \theta_I)$ becomes publicly available. Each firm who hasn't made an offer at $t = 1$ makes an offer at $t = 2$ to one of the remaining workers. Each worker who is offered at $t = 2$ decides which one to take.

We interpret any matching realized at Stage 1 as occurrence of unravelling. If worker i matched with a firm at Stage 1, it is possible that his θ_i may turn out to be 0 ex post, while another worker i' with $\theta_{i'} = 1$ may be left unemployed at Stage 2. Ideally, if all firms can wait until Stage 2 to make offers, then such a "reversal" would never happen.

Indeed, the main social cost of unravelling in this model is due to *informational inefficiency*. Let us define the measure of informational efficiency as the total θ_i among hired. The *first-best* informational efficiency corresponds to $\sum_{i=1}^{I} \theta_i$ if $\sum_{i=1}^{I} \theta_i \leq J$ (i.e., all talented workers are hired); and J if $\sum_{i=1}^{I} \theta_i > J$ (i.e., all vacancies are filled by talented workers). Compared to this first-best outcome, unravelling may imply a smaller number of hired talented workers.

Let us first observe that unravelling is unavoidable under certain conditions.

Theorem 1. With sufficiently low p and $\underline{u}$ (i.e., there exist some threshold values of p and $\underline{u}$, and if p and $\underline{u}$ are below those thresholds), it is not an equilibrium that all firms hire at Stage 2.

Proof. Let all firms hire at Stage 2. For firm J, his expected payoff is:

$$\begin{aligned}
\Pr(\#(\theta_i = 1) \geq J) &= 1 - \sum_{k=0}^{J-1} C(I,k) q^k (1-q)^{I-k} \\
&= 1 - [1 - O(q^2)] \\
&= O(p^2),
\end{aligned}$$

where the second equality is by Taylor expansion of $\sum_{k=0}^{J-1} C(I,k)q^k(1-q)^{I-k}$ around $q = 0$ and the last equality is because $q = O(p)$.

If J deviates and hires at $t = 1$, assuming that the offered worker would accept it, J's payoff would be:

$$(1-p)^I L + (1-(1-p)^I)H = L + Ip(H-L) + O(p^2) = O(p),$$

where the first equality is by Taylor expansion of $(1-p)^I$ around $p = 0$ and the last equality is because $L + Ip(H-L) = O(p)$; therefore, J would be strictly better by this deviation for sufficiently small p, under the assumption that the offered worker would accept his offer.

Indeed, the worker would accept it given any signal at $t = 1$, as:

$$u(J) > Hu(1) + (1-H)\frac{\sum_{j=1}^{J} u(j) + (I-J)\underline{u}}{I}$$

if $\underline{u}$ is sufficiently low. □

The result uncovers that unravelling would be more relevant if (i) talented workers are more precious and hence firms are more desperate in getting them, and (ii) clerkship positions are more prestigious and hence workers are more desperate in being hired. This sort of two-sided competitiveness is important for their incentives of earlier matching.

The result also suggests that it is relatively unpopular firms who have more incentives of making early offers than the popular firms. This is because they would have much less chance of getting talented workers (as they are rare) if they wait until Stage 2. Moving early is a bet of getting a talented worker.

Given that it is *not* an equilibrium that all hire at Stage 2, a natural question is then what the equilibrium is. For simplicity, assume $J = 2 (< I)$. This make the computation of the equilibrium easier, and also makes it clear that firm 1 is the popular firm and firm 2 is the non-popular firm. Also, assume sufficiently small p and $\underline{u}$ as in the previous result. Let $a_j \in \{1, 2\}$ be j's timing of making an offer (at Stage 1 or 2), and let $v_j(a)$ be j's expected payoff given $a = (a_1, a_2)$. The expected payoffs are summarized in the following table:

	$a_2 = 1$	$a_2 = 2$
$a_1 = 1$	$(L + Ip(H-L), L)$	$(L + Ip(H-L), (I-1)L)$
$a_1 = 2$	$((I-1)L, L + Ip(H-L))$	$(Iq, 0)$

where the first (second) number denotes firm 1's (2's) expected payoff given each a, and all terms are up to $O(p^2)$.[15]

Recall the second empirical result in Section 2, showing that low-productivity judges are hurt by the agreement that all judges wait. Consistently with this finding, we can show that firm 2's (unravelled) equilibrium payoff is higher than when they all wait until Stage 2, which uncovers the fundamental difficulty of sustaining the "all must wait" type of reform agreements attempted in the context of the judicial clerkship market.

Proposition 1. In any equilibrium of the game, firm 2's expected payoff is strictly higher than $v_2(2,2)(= O(p^2))$.

Proof. Substituting $L = \gamma p$, and rearranging further, the payoff table can be written as follows (again, all terms up to $O(p^2)$):

	$a_2 = 1$	$a_2 = 2$
$a_1 = 1$	$((\gamma + IH)p, \gamma p)$	$((\gamma + IH)p, (I-1)\gamma p)$
$a_1 = 2$	$((I-1)\gamma p, (\gamma + IH)p)$	$((\gamma + H)Ip, 0)$

It suffices to show that, in any equilibrium, there is at least one action profile other than $a = (2,2)$ played with a strictly positive probability that is non-vanishing in p, because then firm 2's expected payoff is $O(p)$, and hence strictly greater than $v_2(2,2) = O(p^2)$ with sufficiently small p.

First, if $a_2 = 1$ is a strict best response in an equilibrium (and hence it is played with probability one), then we are done. If $a_2 = 1$ is not a strict best response in some equilibrium, then, letting α be the probability that $a_1 = 1$ is played in that equilibrium, we must have:

$$\alpha(I-1)\gamma p + O(p^2) \geq \alpha\gamma p + (1-\alpha)(\gamma + IH)p + O(p^2),$$

implying that $\alpha \geq \frac{\gamma + IH}{(I-1)\gamma + IH} \geq \frac{1}{I-1}$ (up to $O(p)$). $\square$

The exact shape of the equilibrium depends on the specificity of the parameters, but the next result shows that an equilibrium could be in mixed strategies. This means that actual hiring timing we see in data may be somewhat random, and hence it may be difficult to see unambiguous patterns in the equilibrium behavior.

Moreover, in case of a mixed-strategy equilibrium, its attained informational efficiency is much lower than under the first-best outcome.

Proposition 2. Assume $J = 2$, sufficiently small $p, \underline{u}$ as in the previous result, and $L + Ip(H-L) > (I-1)L$ (equivalently, $IH > (I-2)\gamma$). Then, the unique Bayesian Nash equilibrium is in mixed strategies. Moreover, in that equilibrium, $a = (1,1)$ is to be played with a probability non-vanishing in p, which implies that the loss in the informational efficiency relative to the first-best outcome is non-vanishing in p.

[15] Specifically, $v_1(2,2) = 1-(1-q)^I = Iq + O(q^2) = Iq + O(p^2)$, $v_2(2,2) = 1-(1-q)^I - Iq(1-q)^{I-1} = O(p^2)$; $v_1(1,2) = v_2(2,1) = \Pr(\#(s_i = H) \geq 1)H + \Pr(\#(s_i = H) = 0)L = L + Ip(H-L) + O(p^2)$, $v_2(1,2) = v_1(2,1) = \Pr(\#(s_i = H) \leq 1)(1-(1-L)^{I-1}) + O(p^2) = (I-1)L + O(p^2)$; $v_1(1,1) = \Pr(\#(s_i = H) \geq 1)H + \Pr(\#(s_i = H) = 0)L = L + Ip(H-L) + O(p^2)$, $v_2(1,1) = \Pr(\#(s_i = H) \leq 1)L + \Pr(\#(s_i = H) \geq 2)H = L + O(p^2)$.

Proof. For the first part, it suffices to show that any pure-strategy combination cannot be an equilibrium, which we omit for brevity.

For the second part, by the previous proposition, we already know that $a_1 = 1$ is played with a probability non-vanishing in p. Thus, it suffices to show that $a_2 = 1$ too is played with a probability non-vanishing in p.

Indeed, letting α be the probability that firm 2 plays $a_2 = 1$ in this mixed-strategy equilibrium, firm 1's indifference condition implies:

$$\begin{aligned} & (\gamma + IH)p + O(p^2) = \alpha(I-1)\gamma p + (1-\alpha)(\gamma + H)Ip + O(p^2) \\ \Leftrightarrow \quad & \alpha = \frac{\gamma(I-1)}{IH - \gamma} + O(p). \end{aligned}$$

□

3.3 Repeated-game perspective

As in the judicial clerkship market, for hiring markets of new graduates, it is usually the case that workers participate in the market only once in their life-time, and in this sense, they are "short-lived" players. On the contrary, firms often participate in the market regularly. Thus, it seems natural to view this hiring practice as a repeated game among long-lived firms with short-lived workers. As we show below, this perspective provides some insights for sustainable improvement of informational efficiency.

Although the basic intuition should be robust in terms of the number of firms, let us continue to assume $J = 2(< I)$ for simplicity, in order to limit the number of cases to consider. For each year $t = 0, 1, \ldots$, firms and workers play the two-stage game as in the previous subsection. The same two firms operate across all t, and hence they maximize their own discounted payoff sums across t with a common discount factor δ; while workers playing at t only live in year t (and new workers of the same size I will come at $t+1$), and hence they maximize own myopic payoffs. Firm 1 is always the popular firm, and firm 2 is always the non-popular firm (i.e., $u(1) > u(2) > \underline{u}$, as in the previous subsection).

We first show that, even from the repeated-game perspective, unravelling is still a severe problem: With an arbitrary discount factor $\delta \in (0, 1)$, it is *not* an equilibrium in the repeated game where both firms hire at Stage 2 at every t. That is, the reforms of the form implemented in the past in the judicial clerkship market would not work, no matter how patient judges are. This is in a stark contrast to other applied repeated game settings (such as repeated auctions as discussed in Section 1), where δ close to one often allows the possibility that the "most efficient" outcome (in the sense of the corresponding context) is attainable thanks to the folk-theorem logic (Fudenberg et al. (1994)). In our case, the first-best informational efficiency necessarily makes the non-popular firm *a loser*, which is the key source of this impossibility.

Proposition 3. With sufficiently small $p, \underline{u}$, for any $\delta \in (0, 1)$, it is impossible to make all firms hire at $t = 2$ with probability one at every year t.

Proof. If it were possible in some equilibrium, then the discounted sum of payoffs of firm 2 is $\frac{1}{1-\delta}O(p^2)$.

If firm 2 deviates at any t, then his continuation payoff comprises $L + Ip(H - L) + O(p^2)$ at t, and (at least) his minmax payoff from $t+1$ on. Firm 2 can always guarantee, by playing $a_2 = 1$, at least the payoff of $L + O(p^2)$. Thus, his continuation payoff from t on is:

$$L + Ip(H - L) + \frac{\delta}{1-\delta}L + O(p^2) = \frac{L}{1-\delta} + O(p^2) = \frac{1}{1-\delta}\gamma p + O(p^2),$$

which is strictly greater than $\frac{1}{1-\delta}O(p^2)$ with sufficiently small p. □

Although the first-best informational efficiency is impossible, we show that some improvement would be possible relative to the unregulated unravelling equilibrium.

Let us first introduce some additional definitions. If only one firm hires at Stage 1 (and the other hires at Stage 2), then we call the outcome *second-best informationally efficient*, in the following sense.

Proposition 4. Consider any outcome of the stage game where both firms hire at Stage 1. Higher informational efficiency is achieved if only one of them hires at $t = 1$.

Proof. Recall that the informational efficiency is measured by the total number of hired talented workers, which coincides with the total firm payoff in the current specification. Thus, it suffices to show that the total payoff given $a = (1, 1)$ is lower than that given $a = (1, 2)$ (and $a = (2, 1)$, but recall that the total payoff under $a = (2, 1)$ equals that under $a = (1, 2)$: see Footnote 4).

If firm 1 hires at Stage 1, then regardless of whether firm 2 hires at Stage 1 or 2, firm 1 hires one of the workers with the highest s_i, and hence his expected payoff does not depend on firm 2's behavior either. For firm 2, given firm 1's hiring at Stage 1, it is better to wait until Stage 2 and hires a worker with $\theta_i = 1$ if any; if none, then firm 2's payoff is 0 regardless of whether he hires at Stage 1 or 2. Therefore, firm 2's expected payoff is strictly better off by waiting until Stage 2. To conclude, the total expected payoff is strictly higher under $a = (1, 2)$ than under $a = (1, 1)$. □

Note that, when p is small, the difference between the first-best and second-best information efficient outcomes is tiny, as it is given by $O(p^2)$.

3.3.1 Random random early-mover schemes

Given the impossibility of the first-best informational efficiency (with any δ), a natural alternative is to investigate the possibility of achieving (at least) the second-best informational efficiency. In the rest of this section, to simplify the argument, we treat the game as if only (two) firms are the players, whose stage-game payoffs are given as in the table in page 11.[16]

We consider the following *random early-mover scheme*:

[16]Workers are only implicitly treated as a factor which induces that payoff table.

1. Begin at $t = 1$ with the play of $a = (2, 2)$ with probability 1.

2. At each $t \geq 2$:

   (a) If at least one of the players chose an action different from the prescription at some $\tau < t$, then the players play the static mixed equilibrium.

   (b) If $a = (2, 2)$ was the prescription at $t - 1$ and both players indeed played it at $t - 1$, then the players play $a = (2, 2)$ again with probability β_2, while they play $a = (2, 1)$ with probability $1 - \beta_2$.

   (c) If $a = (2, 1)$ was the prescription at $t - 1$ and both players indeed played it at $t - 1$, then the players play $a = (2, 1)$ again with probability β_1, while they play $a = (2, 2)$ with probability $1 - \beta_1$.

Hence, each random early-mover scheme is identified by $\beta = (\beta_1, \beta_2)$. We say that a scheme is *implementable* if it is indeed a (subgame-perfect) equilibrium that both firms play according to the scheme's proposal. Note that its implementability implies that the induced average payoffs of the firms Pareto dominate that of the static mixed equilibrium, because each player can always revert to the repetition of the static mixed equilibrium if he prefers.

Moreover, by construction, a random early-mover scheme achieves the second-best informational efficiency, as long as it is implementable. Therefore, it is inferior to the first-best outcome only up to $O(p^2)$. Compare this with the static mixed-strategy equilibrium. There, $a = (1, 1)$ can be played with a significant probability (Proposition 2), and therefore, its expected informational efficiency is much lower than the first- and second-best informationally efficient outcomes.

Obviously, there is no β with which a random early-mover scheme is implementable if δ is too low. The next result shows that there is a lower bound $\delta^* \in (0, 1)$ such that, for any $\delta > \delta^*$, there is *some* β with which a random early-mover scheme is implementable. Contrast this with Proposition 3, which says that the first-best informational efficiency is not implementable with any δ (even if close to 1).

Theorem 2. There exists $\delta^* \in (0, 1)$ such that, for $\delta \geq \delta^*$, there exists $\beta(\delta)$ such that the random early-mover scheme with $\beta = \beta(\delta)$ is implementable.

Proof. Let $W_i(2, 2)$ ($W_i(2, 1)$) be i's expected continuation payoff from time t on in this scheme, in case the play of $a = (2, 2)$ ($a = (2, 1)$) is proposed at t. Note that the stationary structure of this scheme makes W_i independent of t. We have the following recursive characterization of W_i:

$$\begin{aligned} W_i(2,2) &= v_i(2,2) + \delta(\beta_2 W_i(2,2) + (1-\beta_2) W_i(2,1)) \\ W_i(2,1) &= v_i(2,1) + \delta(\beta_1 W_i(2,2) + (1-\beta_1) W_i(2,1)). \end{aligned}$$

Solving the system, we obtain:

$$\begin{aligned}
W_i(2,2) &= \frac{1}{(1-\delta)(1+\delta(\beta_1-\beta_2))}((1-\delta(1-\beta_1))v_i(2,2)+\delta(1-\beta_2)v_i(2,1)) \\
W_i(2,1) &= \frac{1}{(1-\delta)(1+\delta(\beta_1-\beta_2))}((1-\delta\beta_2)v_i(2,1)+\delta\beta_1 v_i(2,2)).
\end{aligned}$$

For this scheme to constitute an equilibrium of the repeated game, it is necessary and sufficient that firm 1 does not deviate when $a=(2,1)$ is proposed, and firm 2 does not deviate when $a=(2,2)$ is proposed (the other potential deviations are clearly non-profitable). That is:

$$\begin{aligned}
\frac{1}{(1-\delta)(1+\delta(\beta_1-\beta_2))}[(1-\delta\beta_2)v_1(2,1)+\delta\beta_1 v_1(2,2)] &\geq \frac{1}{1-\delta}v_1(1,1) \\
\frac{1}{(1-\delta)(1+\delta(\beta_1-\beta_2))}[(1-\delta(1-\beta_1))v_2(2,2)+\delta(1-\beta_2)v_2(2,1)] &\geq v_2(2,1)+\frac{\delta}{1-\delta}\underline{v}_2;
\end{aligned}$$

equivalently:

$$\begin{aligned}
(1-\delta\beta_2)v_1(2,1)+\delta\beta_1 v_1(2,2) &\geq (1+\delta(\beta_1-\beta_2))v_1(1,1) \\
(1-\delta(1-\beta_1))v_2(2,2)+\delta(1-\beta_2)v_2(2,1) &\geq (1+\delta(\beta_1-\beta_2))[(1-\delta)v_2(2,1)+\delta\underline{v}_2].
\end{aligned}$$

Call them the *incentive-compatibility* conditions.

Let $\tilde{\rho}=\frac{1-\delta(1-\beta_1)}{1+\delta(\beta_1-\beta_2)}$. Because the conditions are linear in β, we can focus our attention on the "boundary" β in the following sense.

Lemma 1. If there exists $\beta(\in[0,1]^2)$ that satisfies the two incentive-compatibility conditions, then they can be satisfied by β with either (i) $\beta_1=1$ and $\beta_2=\frac{1+\delta-\tilde{\rho}^{-1}}{\delta}$, or (ii) $\beta_2=0$ and $\beta_1=\frac{\tilde{\rho}-1+\delta}{\delta(1-\tilde{\rho})}$.

Proof. Because the constraints are both linear in β, if there exists β that satisfies the two incentive-compatibility conditions, then such β can be found at a boundary point: either (a) $\beta_1=1$ and β_2 satisfies firm 2's incentive constraint with equality, (b) $\beta_1=0$ and β_2 satisfies firm 2's incentive constraint with equality, (c) $\beta_2=1$ and β_1 satisfies firm 2's incentive constraint with equality, or (d) $\beta_2=0$ and β_1 satisfies firm 2's incentive constraint with equality.

Case (a) corresponds to Case (i) in the statement, and it makes sense if and only if $\frac{1+\delta-\tilde{\rho}^{-1}}{\delta}\in[0,1]$; equivalently:

$$\frac{\delta(v_2(2,1)-\underline{v}_2)}{v_2(2,1)-v_2(2,2)}\geq\frac{1}{1+\delta}.$$

Case (d) corresponds to Case (ii) in the statement, and it makes sense if and only if $\frac{\tilde{\rho}-1+\delta}{\delta(1-\tilde{\rho})}\in[0,1]$; equivalently:

$$\frac{\delta(v_2(2,1)-\underline{v}_2)}{v_2(2,1)-v_2(2,2)}\in[1-\delta,\frac{1}{1+\delta}].$$

.

Cases (b) and (c) are impossible, because $\beta_1 = 0$ violates firm 1's incentive condition (with any β_2), and $\beta_2 = 1$ violates firm 2's (with any β_1).[17] $\square$

As made clear in the proof, no β can satisfy the two incentive conditions if $\frac{\delta(v_2(2,1)-\underline{v}_2)}{v_2(2,1)-v_2(2,2)} < 1-\delta$, which proves:

$$\delta^* = \left(1 + \frac{v_2(2,1) - \underline{v}_2}{v_2(2,1) - v_2(2,2)}\right)^{-1} \quad (\in (0,1)).$$

$\square$

In Case (i), $\beta_2 < 1$ means that, right after the play of $a = (2,2)$, there is a strict positive probability that firm 2 hires early in the next year, though they may again play $a = (2,2)$ with the other probability. $\beta_1 = 1$ means that, right after the play of $a = (2,1)$, the firms surely play $a = (2,2)$ in the next year. This case can be interpreted as the case where firm 2's incentive problem is not so severe. Even though the first-best information efficiency is impossible, we only need to occasionally allow firm 2's early hiring.

Case (ii) corresponds to the case where firm 2's incentive problem is severe. $\beta_2 = 0$ means that, right after the play of $a = (2,2)$, they necessarily play $a = (2,1)$. β_1 could possibly be less than 1, meaning that firm 2 may need to be rewarded multiple periods of early hiring.

The condition for the case division in the proof clearly suggests that, with higher δ, the first case is more likely.

3.3.2 Optimality

The above theorem shows that some particular choice of β satisfies the incentive-compatibility conditions. We next show that this choice is in fact the *optimal* choice in a certain sense. To formally state this, let us introduce some additional definitions and notation.

A *scheme* proposes how the firms should play at each period t, potentially depending on their past plays. Let $\alpha_\tau \in \{1,2\}^2$ represent the proposed actions at τ, and let $a_\tau \in \{1,2\}^2$ represent the actual plays at τ. Their combination for $\tau = 1, \ldots, t-1$, denoted by $h_t = (\alpha_1, a_1; \ldots; \alpha_{t-1}, a_{t-1})$, is what the period-$t$ proposal can be contingent on.

Definition 1. A *scheme* is $\mu = (\mu_t(h_t))_{t,h_t}$, where for each t and h_t, $\mu_t(h_t) \in \Delta(\{1,2\})^2$ is interpreted as a probability distribution over the period-t proposal (i.e., α_t is randomly drawn according to $\mu_t(h_t)$).

[17] Recall $v_2(2,2) < v_2(2,1), \underline{v}_2$.

Each firm i's (pure) strategy is $\sigma_i = (\sigma_{it}(h_t, \alpha_t))$, where $\sigma_{it}(h_t, \alpha_t) \in \{1, 2\}$ describes i's actual play given the past proposals (up to t) and past plays (up to $t-1$): $(h_t, a_t) = (\alpha_1, a_1; \ldots; \alpha_{t-1}, a_{t-1}; \alpha_t)$. Say that σ_i is *obedient to* μ if $\sigma_{it}(h_t, a_t) = a_{it}$ for all t, h_t, and all $a_t = (a_{1t}, a_{2t})$ in the support of $\mu_t(h_t)$. σ is obedient to μ if both σ_1, σ_2 are obedient to μ.

We say that a scheme μ is *implementable* if it is a (subgame-perfect) equilibrium that the firms play obedient strategies to μ. To formally state this, for each t, let $\Pr_t(h_t, \alpha_t; \mu, \sigma)$ denote the distribution over (h_t, α_t) induced by (μ, σ).[18] Analogously, for $\tau \geq t$, the conditional distribution over (h_τ, α_τ) given (h_t, α_t) is denoted by $\Pr_\tau(h_\tau, \alpha_\tau | h_t, a_t; \mu, \sigma)$. Of course, it makes sense only if (h_t, α_t) is a *sub-history* of (h_τ, α_τ) (in the sense that, if $(h_t, \alpha_t) = (\alpha_1, a_1; \ldots; \alpha_{t-1}, a_{t-1}; \alpha_t)$, then $(h_\tau, \alpha_\tau) = (\alpha_1, a_1; \ldots; \alpha_{t-1}, a_{t-1}; \alpha_t, \ldots)$); let $\Pr_\tau(h_\tau, \alpha_\tau | h_t, \alpha_t; \mu, \sigma) = 0$ otherwise. Let:

$$V_{it}(h_t, \alpha_t; \mu, \sigma) = \sum_{\tau \geq t} \delta^{\tau - t} \sum_{(h_\tau, \alpha_\tau)} \Pr_\tau(h_\tau, \alpha_\tau | h_t, \alpha_t; \mu, \sigma) v_i(\sigma(h_\tau, \alpha_\tau))$$

be i's continuation payoff from t on conditional on (h_t, α_t), in the scheme μ associated with the strategy profile σ.

Definition 2. A scheme μ is *implementable* if, for an obedient σ^*, we have:

$$V_{it}(h_t, \alpha_t; \mu, \sigma^*) \geq V_{it}(h_t, \alpha_t; \mu, \sigma_i, \sigma^*_{-i})$$

for all t, h_t, all α_t in the support of $\mu_t(h_t)$, and all strategy σ_i.

Theorem 3. For any $\delta \geq \delta^*$, among the implementable random early-mover schemes, it is optimal to set β as in Case (i) or (ii) in the previous theorem.[19] Furthermore, the random early-mover scheme with the optimal β attains the best informational efficiency among all implementable schemes.

Proof. Let us first show the first claim about the optimal β. Because the initial period's play of a random early-mover scheme is $a = (2, 2)$ for sure, the (ex ante) informational efficiency is given by:

$$\begin{aligned} & W_1(2,2) + W_2(2,2) \\ = \;& \frac{(1 - \delta(1 - \beta_1))(v_1(2,2) + v_2(2,2)) + \delta(1 - \beta_2)(v_1(2,1) + v_2(2,1))}{(1 - \delta)(1 + \delta(\beta_1 - \beta_2))} \\ = \;& \frac{1}{1 - \delta} \left[\tilde{\rho}(v_1(2,2) + v_2(2,2)) + (1 - \tilde{\rho})(v_1(2,1) + v_2(2,1)) \right], \end{aligned}$$

where $\tilde{\rho} = \frac{1 - \delta(1 - \beta_1)}{1 + \delta(\beta_1 - \beta_2)}$. The optimal β is obtained by maximizing this subject to the two incentive conditions.

[18]Note that $\sum_{(h_t, \alpha_t) \in H_t \times \{1,2\}^2} \Pr_t(h_t, \alpha_t; \mu, \sigma) = 1$ for each t.

[19]Recall that Case (i) sets $\beta_1 = 1$ and $\beta_2 = \frac{1 + \delta - \tilde{\rho}^{-1}}{\delta}$, which applies when $\frac{\delta(v_2(2,1) - \underline{v}_2)}{v_2(2,1) - v_2(2,2)} \geq \frac{1}{1+\delta}$; and Case (ii) sets $\beta_2 = 0$ and $\beta_1 = \frac{\tilde{\rho} - 1 + \delta}{\delta(1 - \tilde{\rho})}$, which applies when $\frac{\delta(v_2(2,1) - \underline{v}_2)}{v_2(2,1) - v_2(2,2)} \in [1 - \delta, \frac{1}{1+\delta}]$.

Notice that, at the optimum, firm 2's incentive constraint must bind with equality; otherwise, $\beta_1 = \beta_2 = 1$ should be the solution (since the objective improves with higher β_1, β_2, and $\beta = (1,1)$ satisfies firm 1's incentive condition), but this contradicts that the first-best informational efficiency is impossible. Therefore:

$$\begin{aligned} & (1-\delta(1-\beta_1))v_2(2,2) + \delta(1-\beta_2)v_2(2,1) = (1+\delta(\beta_1-\beta_2))[(1-\delta)v_2(2,1) + \delta\underline{v}_2] \\ \Leftrightarrow \quad & \tilde{\rho} = \frac{\delta(v_2(2,1)-\underline{v}_2)}{v_2(2,1)-v_2(2,2)}, \end{aligned}$$

and thus, the maximized objective:

$$\begin{aligned} & W_1(2,2) + W_2(2,2) \\ = \quad & \frac{1}{1-\delta}\left[v_1(2,1) + v_2(2,1) + \frac{\delta(v_2(2,1)-\underline{v}_2)}{v_2(2,1)-v_2(2,2)}(v_1(2,2)+v_2(2,2)-v_1(2,1)-v_2(2,1))\right]. \end{aligned}$$

As in the previous lemma, $\delta \geq \delta^*$ implies that either Case (i) or (ii) applies. In either case, firm 2's incentive constraint is binding, and hence the above maximum objective is attained.

Next, we show the second claim. Fix any $\delta > \delta^*$, and suppose, for the sake of contradiction, that there exists an implementable scheme $\mu = (\mu_t(h_t))_{t,h_t}$ that attains strictly higher informational efficiency (or equivalently, strictly higher discounted sum of the total firm payoffs).

Let us introduce some additional notation. For each t and (h_t, α_t), let:

$$\rho_t(h_t,\alpha_t) = (1-\delta)\sum_{\tau\geq t}\delta^{\tau-t}\sum_{(h_\tau,\alpha_\tau)}\Pr_\tau(h_\tau,\alpha_\tau|h_t,\alpha_t;\mu,\sigma^*)1_{\{\alpha_\tau=(2,2)\}}.$$

Note that $\rho_t(h_t,\alpha_t) \leq 1$; in particular, $\rho_t(h_t,\alpha_t) \in [1-\delta, 1]$ if $\alpha_t = (2,2)$. Let

$$\overline{\rho} = \sup_{t,h_t,\alpha_t} \rho_t(h_t,\alpha_t).$$

Fix $\varepsilon > 0$, and fix any $t^*, h^*_{t^*}, \alpha^*_{t^*}$ such that $\rho_{t^*}(h^*_{t^*}, \alpha^*_{t^*}) > \overline{\rho} - \varepsilon$. Without loss, let $\alpha^*_{t^*} = (2,2)$.

The (ex ante) informational efficiency in this scheme is upper-bounded by:

$$\frac{1}{1-\delta}\left[\overline{\rho}(v_1(2,2)+v_2(2,2)) + (1-\overline{\rho})(v_1(2,1)+v_2(2,1))\right].$$

By firm 2's incentive constraint:

$$\begin{aligned} 0 \quad & \leq \quad V_{2t^*}(h^*_{t^*}, \alpha^*_{t^*}; \mu, \sigma^*) - V_{2t^*}(h^*_{t^*}, \alpha^*_{t^*}; \mu, \sigma^*_1, \sigma_2) \\ & \leq \quad [\frac{\rho_{t^*}(h^*_{t^*},\alpha^*_{t^*})}{1-\delta}v_2(2,2) + \frac{1-\rho_{t^*}(h^*_{t^*},\alpha^*_{t^*}))}{1-\delta}v_2(2,1)] - [v_2(2,1) + \frac{\delta}{1-\delta}\underline{v}_2] \\ & < \quad -\frac{\overline{\rho}-\varepsilon}{1-\delta}(v_2(2,1)-v_2(2,2)) + \frac{\delta}{1-\delta}(v_2(2,1)-\underline{v}_2), \end{aligned}$$

which implies

$$
\begin{aligned}
& 0 \leq -\overline{\rho}(v_2(2,1) - v_2(2,2)) + \delta(v_2(2,1) - \underline{v}_2) \\
\Leftrightarrow \quad & \overline{\rho} \leq \frac{\delta(v_2(2,1) - \underline{v}_2)}{v_2(2,1) - v_2(2,2)},
\end{aligned}
$$

because $\varepsilon > 0$ can be taken arbitrarily. Therefore, the (ex ante) informational efficiency in this scheme is upper-bounded by:

$$
\frac{1}{1-\delta}\left[v_1(2,1) + v_2(2,1) + \frac{\delta(v_2(2,1) - \underline{v}_2)}{v_2(2,1) - v_2(2,2)}(v_1(2,2) + v_2(2,2) - v_1(2,1) - v_2(2,1))\right],
$$

which contradicts that this scheme strictly improves over the optimal random early-mover scheme. □

4 Proposal

Summarizing the key observations in the previous sections, this section proposes a possible scheme that can potentially sustain informationally desirable outcomes in the judicial clerkship environment.

Randomization Each year, all Federal Circuit Court judges who open clerkship positions are randomly categorized to multiple groups (e.g., two). Those groups are randomly ordered, and assigned their deadlines (e.g., March 1st and May 1st). The randomization is correlated across years as suggested in the theoretical analysis in the last section. In particular, the group which is assigned the later deadline has a higher chance of being assigned the earlier deadline, and vice versa.

The idea of randomization has been common for judges, for example, in their case assignments. In this sense, the mechanism may easily be understood and accepted for them.

Definition of deadline A deadline for a judge is the earliest time that an offer from this judge can expire. A judge can interview at any time earlier than the deadline, and a student can accept an offer before the deadline. However, a judge cannot force a student to accept or reject an offer before the deadline.

Categorization of judges The categorization of judges in the initial year may either be uniformly random, or depend on some exogenous measures that are related to their eagerness of hiring better clerks. As suggested in the previous analysis, those more eager judges should have higher chances of being assigned earlier deadlines. For example, the total number of cases handled in the previous years may be a candidate for such a measure.

Principle The idea behind the mechanism is that an earlier deadline gives a judge some advantage. By randomly rotating, we incentivize judges to stick to the hiring plan. First, for those who are assigned the earlier deadline, they would have incentive not to move even earlier, because, as long as everyone else obeys the rule, there is no chance that other judges preempt them. For those who are assigned the later deadline, they have higher chances of hiring earlier next year, mitigating the incentive of deviating this year.

The goal of this new scheme is to sustain informationally desirable outcomes and to make all the judges better off in the Pareto dominance sense.

Other punishments While we think of this rotation mechanism as a critical aspect of a potentially successful reform, there could be other possible punishments to sanction any misbehavior, which would make the implementation even easier. For example, (a) the list of names of deviating judges may be published; (b) deviating judges may be assigned the latest deadline next years, or they are not even allowed to open a clerkship position; (c) non-deviating judges are subsidized in some way.

5 Conclusion

This paper investigates the Federal Judicial Clerkship market, which suffers from unravelling for a long time, both empirically and theoretically. The main takeaways are the following. First, even though a reform agreement of later hiring seems to significantly improve the production efficiency, its merit is asymmetric, and hence the reform might not be Pareto improving. Especially, the lower-productivity judges are to be hurt by such a reform, making its sustainable implementation difficult.

Based on this observation, we construct a theoretical model and its repeated version. The first-best efficiency would be difficult even from the repeated-game perspective with arbitrarily patient players, but we find that it could be a useful scheme to occasionally allow early hiring to some low-productivity employers.

A proposal is provided based on the insights gained in this paper's analysis. Needless to say, the practical implementation of such a random early-mover scheme would not be straightforward. For example, one might think it is unfair to treat some employers better than others by allowing for early hiring. There is also a related issue of how to identify the employers who need such better treatments. It is beyond the scope of the paper to discuss detailed ideas for fully practical implementation, but we believe that this paper could be useful as the first step toward a sustainable agreement. More generally, the fact that similar schemes are adopted in other industries such as in Major League Baseball Draft and National Basketball Association Draft suggests that its implementation is not entirely unrealistic (though their contexts are quite different). More detailed investigation of practical implementation is left for future research.

References

Aoyagi, M. (2003): "Bid Rotation and Collusion in Repeated Auctions," *Journal of Economic Theory*, 112, 79–105.

Avery, C., C. Jolls, R. A. Posner, and A. E. Roth (2001): "The market for Federal jucidial law clerks," *The University of Chicago Law Review*, 68, 793.

——— (2007): "The New Market for Federal Judicial Law Clerks," *The University of Chicago Law Review*, 447–486.

Chiu, Y. S. and W. Weng (2019): "Endogenous preferential treatment in centralized admissions," *RAND Journal of Economics*, 40, 258–282.

Cohen, A. (2023): "The Pervasive Influence of Ideology at the Federal Circuit Courts," *Working paper 31509, NBER Working Paper Series*.

Damiano, E., H. Li, and W. Suen (2005): "Unravelling of Dynamic Sorting," *Review of Economic Studies*, 72, 1057–76.

Fréchette, G. R., A. E. Roth, and M. U. Ünver (2007): "Unraveling yields inefficient matchings: evidence from post-season college football bowls," *The RAND Journal of Economics*, 38, 967–982.

Fudenberg, D., D. Levine, and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1039.

Halaburda, H. (2010): "Unraveling in Two-Sided Matching Markets And Similarities of Preferences," *Games and Economic Behavior*, 69, 365–93.

Li, H. and S. Rosen (1998): "Unraveling in Matching Markets," *American Economic Review*, 88, 371–87.

Li, H. and W. Suen (2000): "Unraveling in Matching Markets," *Journal of Political Economy*, 108, 1058–91.

Niederle, M. and A. E. Roth (2009): "Market Culture: How Rules Governing Exploding Offers Affect Market Performance," *American Economic Journal: Microeconomics*, 1, 199–219.

Roth, A. (2018): "Marketplaces, Markets, and Market Design," *American Economic Review*, 108, 1609–58.

Roth, A. and X. Xing (1994): "Jumping the Gun: Imperfections and Institutions Related to the Timing in Market Transactions," *American Economic Review*, 84, 992–1044.

Skrzypacz, A. and H. Hopenhayn (2004): "Tacit collusion in repeated auctions," *Journal of Economic Theory*, 114, 153–169.

Suen, W. (2000): "Unraveling in Matching Markets," *RAND Journal of Economics*, 31, 101–20.