When Matching Markets Unravel?
Theory and Evidence from Federal Judicial Clerkships*

Daniel L. Chen, Yinghua He, Takuro Yamashita

Preliminary and Incomplete.

Abstract

We study the judge-clerk match, a market plagued by unraveling. Evidence from a unique dataset on match and production shows that (1) agents on either side have similar preferences over those on the other side, (2) the matching game for judges is close to zero-sum, (3) this fierce competition among judges explains the unraveling in this market. We develop a theoretical model investigating how homogeneity of preferences (and competition) affects unraveling in matching markets. We show that a static mechanism, as proposed in many previous reforms, is impossible to solve the problem of unraveling in a market with a high degree of homogeneity. By contrast, a dynamic mechanism that takes advantage of judges' repeated participation in the market over time is proven promising. Based on our findings, we propose a new market design for the judge-clerk match.

1 Introduction

The federal judicial clerkship market presents a setting for market design unique for its perennial failure of reforms. We provide empirical evidence on the aspects of this market that differ from

*Daniel L. Chen, daniel.chen@iast.fr, Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France. Yinghua He, yinghua.he@rice.edu, Rice University. Takuro Yamashita, takuro.yamashita@tse-fr.eu, Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France.
First draft: June, 2019. Current draft: July, 2019.
Most recent version at: http://users.nber.org/∼dlchen/papers/When_Matching_Markets_Unravel.pdf.

other matching markets. In particular, we show that judges have similar preferences over clerks and that past reforms benefit average judges but hurt some. This implies that those reforms are not Pareto improving. We develop a theoretical model explaining why the past reforms have failed, taking into account the homogeneity of judge preferences. Our first result is that a static mechanism, similar to those proposed in the past reforms, cannot succeed in such a matching market. We note that judges participate in the matching market every year, which provides a new avenue for market design. Indeed, we design a dynamic mechanism that takes advantage of the repeated nature of the game can successfully avoid unraveling. Based on these results, we provide a detailed proposal for the judge-clerk match.

A law clerk assists judges on a range of tasks including researching issues, drafting opinions, and making legal determinations. Most law clerks are recent graduates who performed near the top of their class in law school. The positions are highly sought after as they can lead to professional opportunities. Some federal judges receive thousands of applications for a single position and even the least sought-after clerkship will receive over 150 applications. Each judge presently hires four clerks for a year, which leads us to a many-to-one matching problem: There are roughly 167 judges (similar to firms), each of whom is hiring 4 law students on a one-year contract from a much larger pool of candidates. The matching can be considered as a non-transferable utility problem because each clerk receives fixed salary.

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. Due to extreme competition, by judges to get the best candidates and by candidates to get the best judges, sometimes judges can require a candidate provide an answer to the question, “Will you accept an offer?” 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.

Several failed reforms have been attempted to regulate the earliest date at which law students could be hired. The market promptly unraveled in each of these prior reforms, in 1983, 1986, 1990, and 2005. While the reforms varied in their specific implementation, they generally had a deadline like “no job offers, tentative or final, shall be made to law clerk applicants before May

1st of the applicant's second year" or "judges should not consider applications before September 15 of the students' third year of law school." These failures have sparked an active theoretical and experimental literature (for example, Avery, Jolls, Posner, and Roth, 2001, 2007; Fréchette, Roth, and Ünver, 2007). This literature observes that some Circuits (Fifth, Seventh, and Eleventh) were noted to "cheat" in the reform years.

No prior research has examined the impact of unraveling on judicial production yet. Since it is difficult to attribute lower quality of judicial decision-making to mismatches (from the perspective of judges' and applicants' preferences), much less, calculate what is a mismatch, establishing prima facie whether successful matching would yield social improvement is one contribution of our empirical research. We hope it motivates further research on this important topic and to encourage legal decision-makers to conform to, rather than flout, their own rules.

2 Empirical Evidence

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 two types of regressions with yearly data: market-level time series regressions and judge-level panel regressions. A year is defined as September to August, taking into account that a clerk can arrive between June and October.

2.1 Data and Institutional Background (to be completed)

We use a dataset on all 380,000 published decisions (over a million judge votes) in U.S. Circuit Courts since 1880. We have the full citation network between the cases. We have detailed metadata for each case, from which we use in particular the court, publication date, and authoring judge.

2.2 Production, Total Output, and Variance (Inequality among Judges)

We test how important is match-specific productivity. The output of judge i together with the matched clerks, $\mu_t(i)$, in year t is:

$$q(i,\mu_t(i)) = \alpha_{i,t} + \sum_{j\in\mu_t(i)} \beta_j + \sigma \sum_{j\in\mu_t(i)} \varepsilon_{i,j},$$

where $q(i,\mu_t(i))$ can be measured by citations to i's cases published by the team $(i,\mu_t(i))$. To take the production function to data, we have:

$$q(i,\mu_t(i)) = \alpha_i + X_{i,t}\gamma + \sum_{j\in\mu_t(i)} Z_j\beta + \sum_{j\in\mu_t(i)} (X_{i,t}\times Z_j)\theta + f(t) + \kappa Reform_t + \xi_{i,t}, \tag{1}$$

where $X_{i,t}$ captures the time-variant characteristics of judge i such as age/experience; Z_j is a vector of characteristics of clerk j, e.g., law school quality; $f(t)$ captures time trend, which could be linear or quadratic; $Reform_t$ is a binary dummy which equals to one if year t is a reform year. θ captures the complementarity between judge and clerk's observed attributes. κ captures the average effect of a reform.

We are also interested in the social welfare which is better approximated by the total output:

$$\sum_{i\in N} q(i,\mu_t(i)) = \sum_{i\in N}\alpha_i + \sum_{i\in N} X_{i,t}\gamma + \sum_{i\in N}\sum_{j\in\mu_t(i)} Z_j\beta + \sum_{i\in N}\sum_{j\in\mu_t(i)} (X_{i,t}\times Z_j)\theta + |N|f(t) + |N|\kappa Reform_t + \sum_{i\in N}\xi_{i,t}, \tag{2}$$

This leads to a time-series regression with aggregated data. After estimating the coefficients, we measure the effect of unravelling on inequality (or variance in $q(i,\mu_t(i))$ in each year) and decompose it into observed part (i.e., X and Z) and unobserved part (the rest). Independently, we also investigate how unravelling affect inequality among judges by reduced-form regressions.

2.3 Time Series Analysis of Aggregate Production

We start with analyzing equation (2) with the aggregate data. Specifically, we ignore the observed heterogeneity (i.e., X and Z) and only include in the regression $Reform_t$ and a quadratic time trend. Moreover, we use ARIMA (1,1) to allow for serial correlation.

As measures of productivity, we focus on the annual total of the following: cases published, citations (within, i.e., from cases within the Circuit), citations (outside, i.e., from cases outside the Circuit), and cases reversed (by the Supreme Court). We focus on the coefficient on $Reform_t$, as it measures the effect of reform on overall production.

Table 1: Aggregate Effect of Reform (All Circuits)

Dependent variable (yearly)	Cases Published	Citations (within)	Citations (outside)	Citations (total)	Cases Reversed
	(1)	(2)	(3)	(4)	(5)
Reform	400.7	789.1	378.3	1296.7	-26.31
	(1673.6)	(2922.9)	(1659.0)	(4137.2)	(20.01)
N	44	44	44	44	44

Notes: ARIMA (1,1) with year and year-squared. Dependent variable calculated as total 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 are in parentheses.

The results are presented in Table 1. It shows that in terms of point estimates, the reforms increase the total productivity: (1) the total number case publications and (2) average citations per case increase, while the (3) probability of SCOTUS reversal decreases. However, these estimates are not statistically significant.

Furthermore, we find the same results among non-cheating circuits, but the results are reversed for cheating circuits, which were significantly adversely affected by reforms. The results are in Table 2, and some coefficients are significant at the 10% level.

Table 2: Aggregate Effect of Reform: Cheating Circuits:

Dependent variable (yearly)	Cases Published	Citations (within)	Citations (outside)	Citations (total)	Cases Reversed
	(1)	(2)	(3)	(4)	(5)
Reform	-145.8	-2445.7*	-528.9	-2920.5*	-6.678*
	(425.4)	(1346.2)	(674.4)	(1770.3)	(3.996)
N	44	44	44	44	44

Notes: ARIMA (1,1) with year and year-squared. Dependent variable calculated as total 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 are in parentheses. *: significant at the 10% level.

2.4 Panel regressions

With annual data on individual judges, we can estimate equation (1) with panel data techniques. Specifically, we control for judge fixed effects, experience, and experience squared.

Table 3: Aggregate Effect of Reform

Dependent variable (yearly)	Cases Published	Citations (within)	Citations (outside)	Citations (total)	Cases Reversed
	(1)	(2)	(3)	(4)	(5)
Reform	1.863***	0.688***	0.0433	0.732***	-0.00267***
	(0.241)	(0.132)	(0.0808)	(0.184)	(0.000683)
Judge Fixed Effects	X	X	X	X	X
N	13695	13220	13220	13220	13695
R-sq	0.683	0.250	0.312	0.265	0.185

Notes: Dependent variable calculated as the yearly mean per case for cases worked on during a market year (September to August) in Columns 2-5 and yearly total in Column 1. Reversed refers to cases worked on in a year eventually reversed by the Supreme Court. Standard errors are in parentheses. ***: significant at the 1% level.

The results are complement to those from the time-series regressions. As Table 3 shows, the effect of the reforms appears as an increase in (1) the total number case publications and (2) average citations per case, while (3) the probability of SCOTUS reversal decreases.

Table 4: Effect of Reform by Cheating Status

Dependent variable (yearly)	Cases Published	Citations (within)	Citations (outside)	Citations (total)	Cases Reversed
	(1)	(2)	(3)	(4)	(5)
Reform	1.869***	0.882***	0.126	1.008***	-0.00297***
	(0.262)	(0.155)	(0.0932)	(0.214)	(0.000810)
Reform x Cheating Circuit	-0.0261	-0.904***	-0.386**	-1.290***	0.00152
	(0.668)	(0.279)	(0.184)	(0.407)	(0.00143)
Judge Fixed Effects	X	X	X	X	X
N	13695	13220	13220	13220	13695
R-sq	0.683	0.250	0.312	0.265	0.185

Notes: Dependent variable calculated as the yearly mean per case for cases worked on during a market year (September to August) in Columns 2-5 and yearly total in Column 1. Reversed refers to cases worked on in a year eventually reversed by the Supreme Court. Standard errors are in parentheses. **: significant at the 5% level. ***: siginficant at the 1% level.

Furthermore, we find the same results among non-cheating circuits, but the results are reversed for cheating circuits (Table 4). These estimates are robust to sequentially adding controls for year fixed effects (Table 5) and judicial experience (Table 6).

Table 5: Effect of Reform: Controlling for Year Fixed Effects

Dependent variable (yearly)	Cases Published (1)	Citations (within) (2)	Citations (outside) (3)	Citations (total) (4)	Cases Reversed (5)
Reform	1.487***	0.763***	0.0115	0.775***	-0.00317***
	(0.259)	(0.156)	(0.0911)	(0.212)	(0.000814)
Reform x Cheating Circuit	0.0159	-0.901***	-0.383**	-1.284***	0.00154
	(0.607)	(0.269)	(0.169)	(0.380)	(0.00144)
Judge Fixed Effects	X	X	X	X	X
N	13695	13220	13220	13220	13695
R-sq	0.683	0.250	0.312	0.265	0.185

Notes: Dependent variable calculated as the yearly mean per case for cases worked on during a market year (September to August) in Columns 2-5 and yearly total in Column 1. Reversed refers to cases worked on in a year eventually reversed by the Supreme Court. Standard errors are in parentheses. **: significant at the 5% level. ***: siginficant at the 1% level.

Table 6: Effect of Reform: Controlling for Year Fixed Effects and Quadratic Experience

Dependent variable (yearly)	Cases Published (1)	Citations (within) (2)	Citations (outside) (3)	Citations (total) (4)	Cases Reversed (5)
Reform	2.222***	0.520***	-0.00564	0.514***	-0.00218**
	(0.305)	(0.123)	(0.0871)	(0.179)	(0.000906)
Reform × Cheating Circuit	-1.560**	-0.597**	-0.452***	-1.049***	-0.00149
	(0.699)	(0.236)	(0.146)	(0.326)	(0.00143)
Judge Fixed Effects	X	X	X	X	X
N	8798	8385	8385	8385	8798
R-sq	0.690	0.435	0.604	0.531	0.128

Notes: Dependent variable calculated as the yearly mean per case for cases worked on during a market year (September to August) in Columns 2-5 and yearly total in Column 1. Reversed refers to cases worked on in a year eventually reversed by the Supreme Court. Standard errors are in parentheses. **: significant at the 5% level. ***: siginficant at the 1% level.

2.5 Assortative Matching and Unraveling

Anecdotally, DC Circuit judges could hold out the longest because the top applicants wanted to go there, and they could fly there on the first day available for interviews and meet with multiple judges on that one day. Judges outside DC therefore felt compelled to hire early to have a chance at these top applicants.

We are able to investigate the effect of one reform year, where we have data on clerk characteristics for 1995-2010. This data is novel in that it doubles the time frame of data originally collated by Katz and Stafford (2010) and analyzed by Bonica et al. (2016).

The hypothesis that we test is that the matching is more assortative in the reform period. For example, we can assume that graduates from top 5 law schools are on average higher quality. We can then look at the fraction of top 5 school graduates of the clerk roster in this time period. The results are dramatic. Three results emerge. First, cheating Circuits appear to be of lower quality, as they tend to have fewer graduates from top law schools. Second, in reform years, the fraction of clerks who graduated from Harvard, Yale, Stanford, Columbia, and NYU are dramatically higher than in other periods. This is possibly due to the ability for judges all over the U.S. to interview these students. Third, after the unravelling, cheating Circuits have a somewhat higher share of clerks who graduated from Harvard, Yale, Stanford, Columbia, and Chicago.

Figure 1: Fraction of Clerks by Law School

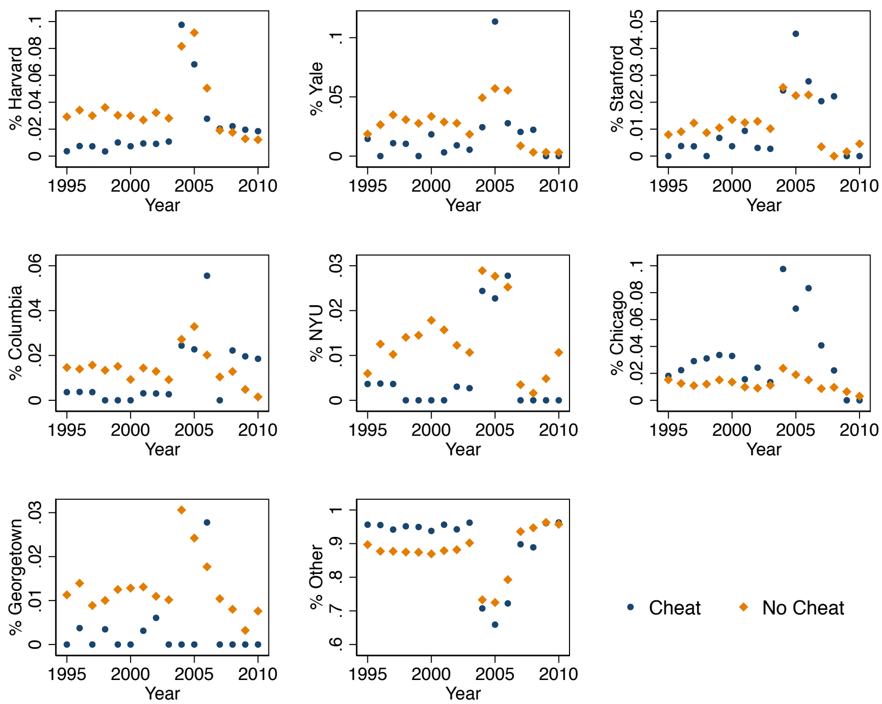

3 A Theoretical Investigation of Unraveling

Consider a dynamic matching problem between law-school students and judges. Every year, some students match with judges for internship. In practice, several times, the judges attempted to commit to the same timing of their job market hiring, but failed to do so, essentially because of some deviating judges who interview some students with early deadlines, i.e., "unraveling".

Our goal in this note is to argue that, under certain conditions, (i) voluntary agreement of a common timing is indeed difficult to be an equilibrium if no punishment can be put on deviators, and (ii) dynamic "stochastic rotation" mechanism can be a good (dynamic) mechanism. The sufficient conditions we obtain look reasonable in this current matching context, but maybe not in other contexts such as medical residents matching. This observation is hence useful in discussing why unraveling occurs in some markets but not (or less) in others.

To provide the main insight in a simple model, assume that there exist only two judges and N students. Every match is assumed to be one-to-one. Every student prefers judge 1 to judge 2, and judge 2 to "unmatched". For normalization, let a student's payoff be 1 if he is matched with judge 1, 0 if with judge 2, and $u_0 < 0$ if unmatched. Each student $i = 1, \ldots, N(> 2)$ is endowed with (unobservable) ability $\theta_i \in \{0, 1\}$, and the payoff of judge j is given by θ_i if he is matched with student i (and by 0 if unmatched).

There exist two periods of time, $t = 1, 2$. Without loss of generality, assume that, at the beginning of $t = 2$, each i's true ability type θ_i is revealed. At the beginning of $t = 1$, only a partially revealing signal, $s_i \in [0, 1]$, is observed, where $s = (s_i)_{i=1}^N$. Without loss, we assume that $\Pr(\theta_i = 1 | s_i) = s_i$. We also assume that $s|\theta$ has a full support, i.e., $\Pr(s \in A|\theta) > 0$ for any open interval $A \subseteq [0, 1]^N$.

Remark 1. These assumptions (on the information structure, the number of players, the number of periods...) are imposed in order to provide the main insight in a simple analysis. Our conclusion is probably robust to less extreme assumptions.

3.1 Single Deadline for Hiring: An Impossibility Result

In view of informational efficiency, the best scenario is where both judges wait until $t = 2$ to hire students. If that happens, given the homogeneous preferences on both sides, a stable matching is fully assortative given the information at t, i.e., judge 1 is matched with the best students, and judge 2 is matched with the second-best student.[1]

[1]More formally, letting I denote the set of all high-ability students,

1. if $|I| \leq 1$, then the high-ability student (if exists) is matched with $j = 1$ for sure; each student $i \notin I$ is matched with $j = 1$ with probability $\frac{1-|I|}{N-|I|}$, and with $j = 2$ with probability $\frac{1}{N-|I|}$; and
2. if $|I| \geq 2$, then each $i \in I$ is matched with $j = 1$ with probability $\frac{1}{|I|}$, and with $j = 2$ with probability $\frac{1}{|I|}$; each $i \notin I$ is unmatched for sure.

However, the next proposition shows that judge 2 would have an incentive to deviate from such an agreement, i.e., unraveling.

Proposition 1. There exists u_0^* such that, for any $u_0 < u_0^*$, it is not an equilibrium for both judges to always hire at $t = 2$. In particular, for some nonempty open set of realization of s (which occurs with a positive probability by the full-support assumption), judge 2 has a strict incentive to hire at $t = 1$.

The proof suggests some sources of this unraveling phenomenon. First, it is crucial that even a very promising student at $t = 1$ has an incentive to accept an offer from the worse judge, even though, by rejecting it, this student may have a good chance of ending up with the better judge at $t = 2$. The reason why a student accepts it is that, for him, the main comparison is not between judge 1 (at $t = 2$) and judge 2 (at $t = 1$), but between no matching (at $t = 2$) and judge 2 (at $t = 1$). Because even a worse judge is much better than no matching, the student has an incentive to accept an early offer from judge 2. This logic suggests a potential explanation as to why some markets with severe demand-supply imbalance, such as judge-clerk markets, suffer from unraveling more prominently than other markets, such as medical internship markets.

However, another question still remains: why a judge, even though he has a strong bargaining power over students due to demand-supply imbalance, has an incentive to make an early offer rather than waiting until $t = 2$. Indeed, imagine, hypothetically, that the judges know that there will always be at least two students with ability one. Then, it is always better for both to wait until $t = 2$ to know which students have ability one. This means that, for unraveling to happen, it must be that high-ability students are rare. Indeed, the key subset A^* in the proof (with which observed, judge 2 has an incentive to deviate) comprises realization s such that there is one student i who has high realization of s_i, and that the other students have low realizations of s_{-i}. This means that, evaluated at $t = 1$, judge 2 is likely to end up with an ability-zero student if he waits until $t = 2$; conversely, if he deviates and approaches student i, then he can steal this promising student at $t = 1$.

In conclusion, two specific properties of the market — demand-supply imbalance and rarity of high-ability students — are the important properties for our unraveling logic.

Remark 2. Although we focus on two-judges environments to simplify the argument, it is straightforward to extend it to cases with more than two judges. In such a case, the lowest-ranked judge would have an incentive to deviate from an agreement of waiting until $t = 2$.

3.2 Dynamic Incentives

Given that the judges are in a long-term relationship, it seems natural to investigate if their dynamic incentives can mitigate the problem, and if so, what kind of reward-punishment schemes would be the most effective.

To provide some idea, imagine that a judge can move earlier than the other judges in year $y+1$ in case he complies with the pre-determined timing of hiring in year y, while otherwise he would be punished (e.g., by the other judges moving very early). If this judge is patient enough, this dynamic incentive scheme could potentially deter his deviation. At the same time, such a (non-monetary) reward scheme introduces some inefficiency, because this early-moving judge is to be matched with a student without much information.

To analyze the dynamic incentive more formally, assume that each judge j maximizes his discounted payoff sum, denoted by $\sum_{y=1}^{\infty} v_{jy}(t_y)$, where $t_y = (t_{1y}, t_{2y})$ denotes the time of hiring at year y. Assume, as in the previous section, that u_0 is low enough so that any offer is to be accepted. Then, essentially the game is just between the two judges of choosing their timing of hiring at each year y.

In what follows, we further assume that the instantaneous payoffs are given by the following table:

	$t_{2y} = 2$	$t_{2y} = 1$
$t_{1y} = 2$	(a_1, a_2)	(c, b)
$t_{1y} = 1$	(b, c)	(b, d)

where $a_1 > b > c > a_2 > d$ and $a_1 + a_2 > b + c > b + d$. The last inequality means that it is Pareto improving for both judges to hire later because of the informational reason.

The idea behind this table and the first set of inequalities can be explained as follows. The top left cell is first-best, where both hire late, but as the most preferred judge, judge 1 gets the best student. If judge 1 moves early (and judge 2 hires at $t = 2$), then judge 1 may hire a bad student in the ex post sense due to uncertainty, which in turn benefits judge 2. On the other hand, if judge 2 moves early (and judge 1 hires at $t = 2$), then judge 2 may hire a good student in the ex post sense, which in turn hurts judge 1. Indeed, if judge 2 moves early, judge 1 also prefers to move early. However, if both move early, the outcome is Pareto inferior to the one where both wait.

Remark 3. As a more concrete situation, imagine that, each year y, there are three students where only one student, say $i=1$, has $s_1 = H$, and the other two students have $s_2 = s_3 = L$, with $1 > H > L > 0$.[2] Then, the judges' year-y payoffs (i.e., the expected student qualities in the corresponding matching) are given as follows:

	$t_{2y}=2$	$t_{2y}=1$
$t_{1y}=2$	$(1-(1-H)(1-L)^2, L^2+2L(1-L)H),$	$(2L-L^2, H))$
$t_{1y}=1$	$(H, 2L-L^2)$	(H, L)

The above inequality is satisfied, for example, where $H=0.6$ and $L=0.3$.

Our goal is to identify the best way to coordinate their hiring timing. Note first that, if the game is played only once, then there is a mixed-strategy equilibrium where judge 1 plays $t_1 = 2$ with probability $\frac{c-d}{c-d+b-a_2}$, and judge 2 plays $t_2 = 2$ with probability $\frac{b-c}{a_1-c}$, yielding judge 1's payoff b and 2's payoff $x \equiv \frac{bc-a_2 d}{c-d+b-a_2}$. Note also that, importantly, "$t_y = (2,2)$ forever" cannot be an equilibrium *even if* δ *is arbitrarily close to one*. This is because judge 2 can use a (deviating) strategy of "hiring at $t=1$ only this year, and then at $t=2$ in all the subsequent years".

In what follows, we use the static mixed equilibrium as the "threat" after any deviation. This is without loss of generality, because in this game, each player's payoff in this mixed-strategy equilibrium coincides with his min-max payoff.

Consider the following form of cooperation: every year (unless any deviation occurs), the judges flip an unfair coin to play $(t_1, t_2) = (2,2)$ with probability q, and play $(2,1)$ with probability $1-q$; any deviation results in the "static mixed equilibrium forever". We obtain the maximum sustainable q as δ tends to one.

First, consider judge 1. His incentive of deviation is highest when $(2,1)$ is supposed to be played. Thus, for him not to deviate, we must have:

$$(1-\delta)c + \delta(a_1 q + c(1-q)) \geq (1-\delta)b + \delta b. \qquad \text{(IC(2,1))}$$

[2] This assumption does not satisfy the full-support condition in Theorem 1, but it is easy to show that the same impossibility result applies (recall that the full-support condition is sufficient for Theorem 1, but not necessary).

As $\delta \to 1$, the condition becomes equivalent to $q \geq \frac{b-c}{a_1-c}$. For judge 2, his incentive of deviation is highest when $(2,2)$ is supposed to be played. Thus, for him not to deviate, we must have:

$$(1-\delta)a_2 + \delta(a_2 q + b(1-q)) \geq (1-\delta)c + \delta\frac{bc - a_2 d}{c - d + b - a_2}. \qquad \text{(IC(2,2))}$$

As $\delta \to 1$, the condition becomes equivalent to $q \leq \frac{b-d}{b-d+c-a_2}$. Therefore, the maximum sustainable q is $\frac{b-d}{b-d+c-a_2}$ when δ is close to one.[3] Note that $\frac{b-d}{b-d+c-a_2} < 1$, consistent with our previous observation that "$(2,2)$ forever" is not sustainable even with δ almost one.

3.2.1 "Informationally Optimal" Equilibrium for Fixed δ

A natural question is whether there exist "better" equilibria, especially when δ is fixed to be less than one.

First, as a way of ranking equilibria, here we consider the *informational* optimality, that is, we characterize the equilibrium that maximizes the total expected ability indexes of the hired students (which equals the total expected payoffs of the judges).

The optimal "Markovian" equilibrium To provide some intuition for better equilibria in this sense than the one in the previous section (i.e., the one with q), imagine that the probability of playing $(2,2)$ in each year y, previously denoted by q, is now made dependent on the play in year $y-1$ (in case of no deviation) in the following way: the judges play $(2,2)$ with probability q_1 in year y if they played $(2,1)$ in $y-1$, while they play $(2,2)$ with probability q_2 if they played $(2,2)$ in $y-1$, where $q_1 > q > q_2$.

Recall that, in the previous construction with q, the maximum possible q is solely determined by the incentive compatibility of judge 2 not to deviate when $(2,2)$ is supposed to be played. On the other hand, judge 1's incentive compatibility is more easily satisfied if q is larger. Thus, in the alternative construction with q_1 and q_2, we can have a lower q_2 to improve judge 2's incentive, and at the same time, to compensate this decrease in q_2 (and hence informational efficiency), we have higher q_1.

It is useful to characterize the optimal (q_1, q_2) (and hence the optimal "Markovian" equilib-

[3] Note that we have $\frac{b-c}{a_1-c} \leq \frac{b-d}{b-d+c-a_2}$.

rium). Let W^0 denote the set of all feasible average payoffs (not necessarily supportable in equilibria), i.e.,

$$
\begin{aligned}
W^0 = \{w \in [0,1]^2 \quad | \quad & \exists \gamma \in [0,1]^4;\ \sum_{k=1}^{4} \gamma_k = 1, \\
& w = \gamma_1(a_1,a_2) + \gamma_2(c,b) + \gamma_3(b,c) + \gamma_4(b,d) \\
& w \geq (b, (1-\delta)c + \delta x)\}
\end{aligned}
$$

Let $W(2,2) = (W_1(2,2), W_2(2,2))$ denote the players' average continuation payoffs if $(2,2)$ is supposed to be played in a given year, and similarly for $W(2,1)$:

$$
\begin{aligned}
W_1(2,2) &= (1-\delta)a_1 + \delta(q_2 W_1(2,2) + (1-q_2)W_1(2,1)) \\
W_1(2,1) &= (1-\delta)c + \delta(q_1 W_1(2,2) + (1-q_1)W_1(2,1)) \\
W_2(2,2) &= (1-\delta)a_2 + \delta(q_2 W_2(2,2) + (1-q_2)W_2(2,1)) \\
W_2(2,1) &= (1-\delta)b + \delta(q_1 W_2(2,2) + (1-q_1)W_2(2,1)),
\end{aligned}
$$

or equivalently,

$$
\begin{aligned}
W_1(2,2) &= \frac{(1-\delta+\delta q_1)a_1 + \delta(1-q_2)c}{1+\delta(q_1-q_2)} \\
W_1(2,1) &= \frac{\delta q_1 a_1 + (1-\delta q_2)c}{1+\delta(q_1-q_2)} \\
W_2(2,2) &= \frac{(1-\delta+\delta q_1)a_2 + \delta(1-q_2)b}{1+\delta(q_1-q_2)} \\
W_2(2,1) &= \frac{\delta q_1 a_2 + (1-\delta q_2)b}{1+\delta(q_1-q_2)}.
\end{aligned}
$$

Assume that the first year begins by playing $(2,2)$. Then, if no one deviates, their average payoffs are $W(2,2)$. As before, the relevant incentive constraints are the following ones:

$$
\begin{aligned}
W_1(2,1) &\geq b \\
W_2(2,2) &\geq (1-\delta)c + \delta\frac{bc - a_2 d}{c - d + b - a_2}.
\end{aligned}
$$

In what follows, we assume that $W(2,2), W(2,1) \in W^0$, which essentially assumes that δ is not too small. Later, we consider the case without Assumption 1.

Assumption 1. $W(2,2), W(2,1) \in W^0$

The optimal Markovian equilibrium is obtained by

$$
\begin{aligned}
w^M = \quad & \max_{q_1, q_2} \quad \frac{(1-\delta+\delta q_1)(a_1+a_2)+\delta(1-q_2)(c+b)}{1+\delta(q_1-q_2)} \\
& \text{sub. to} \quad \frac{\delta q_1 a_1+(1-\delta q_2)c}{1+\delta(q_1-q_2)} \geq b && \text{(IC(2,1))} \\
& \qquad \frac{(1-\delta+\delta q_1)a_2+\delta(1-q_2)b}{1+\delta(q_1-q_2)} \geq (1-\delta)c+\delta\frac{bc-a_2 d}{c-d+b-a_2}. && \text{(IC(2,2))}
\end{aligned}
$$

As suggested by the above discussion, the optimal (q_1, q_2) is such that the two constraints are binding:

$$
\begin{aligned}
\delta q1 &= \frac{(1-\delta)a2+\delta b}{(a1-c)((1-\delta)c+\delta x)-a2(b-c)}(b-c) \\
1-\delta q2 &= \frac{(1-\delta)a2+\delta b}{(a1-c)((1-\delta)c+\delta x)-a2(b-c)}(a1-b).
\end{aligned}
$$

Two natural questions remain given this exercise. First, the policy considered here is "Markovian" in the sense that the distribution over action profiles at $t+1$ depends only on which action profile is played at t, but not on what happened at $t-1$ or before. One may wonder if such a longer-memory strategy profile can generate a more Pareto-superior equilibrium payoffs. The answer is negative, and the argument is standard: for example, in order to satisfy IC(2,2) at period t, we must "promise" judge 2 that his continuation payoff from period $t+1$ on is at least greater than certain amount, but this necessary amount does not depend on what happened at $t-1$ or before. Put differently, making the promised continuation vary with the events at $t-1$ or before does not provide any strict improvement.

Second, even within Markovian policies, the one considered here is special in that it only involves two action profiles, $t=(2,2)$ and $(2,1)$. One may wonder if more complicated strategy profiles which sometimes involve $t=(1,2)$ or $(1,1)$ can generate a better equilibrium payoffs. The answer is again negative, which is formally shown in the next subsection.

The optimal equilibrium Here, we show that the optimal Markovian equilibrium in the previous subsection, i.e., the one with average payoffs w^M, is indeed optimal across all subgame-perfect equilibria (with δ not too low so that Assumption 1 is satisfied).

In general, the set of all (subgame-perfect) equilibrium payoffs with a fixed discounting δ is given by Abreu, Pearce, and Stacchetti (1990). Let $W \subseteq W^0$, and define $B(W) \subseteq W^0$ as follows:

$$
\begin{aligned}
B(W) = \{\beta \in W^0 \quad | \quad & \exists t \in A \equiv \{1,2\}^2,\ \exists \omega : A \to W; \\
& \beta = (1-\delta)v(t) + \delta\omega(t); \\
& \beta_i \geq \max_{t_i' \in \{1,2\}} (1-\delta)v_i(t_i', t_{-i}) + \delta\omega(t_i', t_{-i}),\ \forall i\}
\end{aligned}
$$

By Abreu, Pearce, and Stacchetti (1990), the set of all subgame-perfect equilibrium payoff vectors (with public randomization), denoted by $W^* \subseteq W^0$, is given by the largest $W \subseteq W^0$ that satisfies $W = B(W)$.

In what follows, for each β, the corresponding $t \in A$ and $\omega : A \to W$ in the definition of $B(W^*)$ is denoted by t_β, ω_β.

Theorem 1. With Assumption 1, w^M corresponds to the optimal equilibrium.

Remark 4. As in the previous section, the assumption of two judges is unrealistic but greatly simplifies the analysis to highlight the main point. Although the qualitative results are robust, with more than two judges, which equilibrium is the informationally efficient equilibrium can depend not only δ but also the other payoff parameters, and in this sense, it is probably too complicated to thoroughly analyze all the cases.

To provide some idea, imagine a case with three judges. Assume that judge 2 and 3 have an incentive to unilaterally deviate from $t = (2,2,2)$ ("everyone waits"), and that if one judge moves early, then every other judge also has an incentive to move early. One candidate for the informationally efficient equilibrium is to rotate among $t = (2,2,2), (2,1,2), (2,2,1)$ (i.e., either "everyone waits", "only judge 2 moves early", and "only judge 3 moves early"). Another candidate may be to rotate among $t = (2,2,2), (2,1,1), (2,2,1)$ (i.e., "both judges 2 and 3 move early", instead of "only judge 2 moves early"). If the rotation probabilities are the same, then the first candidate is obvious better, but the second candidate may involve less severe incentive constraint and hence may be able to achieve higher probability of $t = (2,2,2)$. Overall, which one is more efficient can depend on the payoff parameters: for example, if the total payoffs of the judges (which is the measure of information efficiency) are almost the same between $(2,1,2)$ and $(2,1,1)$, while the incentive compatibility constraint is much severer given $(2,1,2)$ than given $(2,1,1)$, then the second candidate can be the most informationally efficient equilibrium (and vice versa).

In this sense, the full analysis with more than two judges can be complicated, but the fundamental feature of the optimal scheme, stochastic rotation of early hiring, seems quite robust.

With low δ? In the previous two subsections, Assumption 1 guarantees that w^M is feasible, which essentially says that δ is not too low. In case δ is low, we argue that any agreement is not really sustainable.

To obtain this conclusion, first, consider an equilibrium where they are supposed to play either (1,2) or (1,1) on the equilibrium path (if it exists).[4] Technically, in our current model, this could be incentive compatible: for example, if (1,2) is supposed to be played, judge 1 has an incentive to deviate to (2,2), but if δ is sufficiently high, then a possibility of future punishment can prevent him from the deviation. However, such a punishment looks strange given our goal of delaying hiring as much as possible. Besides, in practice, detecting such a deviation and punishing it is perhaps difficult. For example, judge 1 can always claim that he tried to hire someone at $t = 1$ but did not succeed. Similarly, from (1,1), judge 2 has an incentive to deviate to (1,2). Indeed, if we change the model so that only earlier hiring activities than the agreement are detectable, then those action profiles cannot be played in any equilibrium. In this spirit, in what follows, we exclude any such equilibrium.

We first note the following lemma.

Lemma 1. If Assumption 1 is violated, then at least one of the following two cases applies:

- for any $w \in W^0$ that satisfies IC(2,2), w' does not satisfy IC(2,1), where $w = (1-\delta)(a_1, a_2) + \delta w'$; or
- for any $w \in W^0$ that satisfies IC(2,1), w' does not satisfy IC(2,2), where $w = (1-\delta)(c, b) + \delta w'$.

[4] More precisely, we consider an equilibrium where, in an event which happens with a positive probability on the equilibrium path, either (1,2) or (1,1) are played surely conditional on that event.

In case Assumption 1 is violated, first, assume that the first case in the statement of the lemma applies. In this case, for any $w \in W^*$, its associated t^w cannot be (2,2), which immediately implies that any $w \neq (b, (1-\delta)c + \delta x)$ cannot be in W^*.

Similarly, the second case in the statement of the lemma also implies that W^* can only contain $(b, (1-\delta)c + \delta x)$.

In sum, if δ is so low that Assumption 1 is violated, it is impossible to stochastically rotate between $(2,2)$ and $(2,1)$, because one of the judges would deviate given any such rotation scheme. As a consequence, only "repeating static Nash equilibrium" is the only possibility.

4 A Proposal to Solve the Market Unraveling

Contrasting with prior reforms, which relied on a static mechanism, the core of our proposal is a dynamic one.

1. **Randomization**: We first specify a set of deadlines, e.g., March 1, April 1, and May 1. Then, each judge who is willing to participate is assigned a deadline, according to an order that is determined by randomization. To insure equality, we make the randomization correlated across years (e.g., "the earliest mover in one year gets a lower chance of being an early mover in the following year"), but everyone, including the earliest mover has a positive probability of being the earliest mover in the following year.

2. **Definition of deadline**: A deadline means that an offer from the judge will expire on that date. A judge can interview at any time earlier than the deadline, and a student can accept an offer before the deadline. However, forcing a student to accept or reject an offer before the deadline is prohibited.

3. **Principle**: The idea behind the mechanism is that an earlier deadline gives a judge some advantage. By the random rotation, we incentivize judges to stick to the hiring plan. Given the impossibility result for static mechanisms (Proposition 1), this kind of "dynamic" mechanism may be our only hope. As the judges have been using randomization in case assignment, the mechanism may be familiar to them.

4. **Unraveling in terms of interviews**: It might be the case that a judge with a very late deadline has incentives to interview and to make offers early. This may undermine the advantage of earlier movers. However, this incentive is mitigated by several factors: (a) It is costly to interview students. Interviewing early will require the judge to interview those who later will be matched with other judges. (b) The extent of competition from later movers is limited by the number of offers they can make. There can be some overbooking, but the total number of offers is limited by the total number of positions. Since a judge will care about the composition of his clerk team, this means the second offer depends on the decision of the student with the first offer. (c) A judge may get better signals of the qualities of a student as time passes. (d) There can also be some incentives/punishments as described below.

5. **Possible punishment/incentive**: (a) Publish the list of names of deviating judges. (b) Deviating judges will have a lower, but non-zero, probability of being the first mover for some years. (c) Interviews conducted close to the judge's deadline will be subsidized. (d) No Supreme Court clerkship position for clerks with deviating judges (this is already partially in place as of 2019).

5 Concluding Remarks

We use a unique dataset on match and production to study the judge-clerk match. We find that (1) agents on either side have similar preferences over those on the other side, (2) the matching game for judges is close to zero-sum, (3) this fierce competition among judges explains the unraveling in this market.

In a theoretical model, we show how homogeneity of preferences (and competition) affects unraveling in matching markets. A static mechanism, as proposed in many previous reforms, is impossible to solve the problem of unraveling in a market with a high degree of homogeneity. By contrast, a dynamic mechanism that takes advantage of judges' repeated participation in the market over time is proven promising. Based on our findings, we propose a new market design for the judge-clerk match.

References

ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58, 1041–1063.

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.

BONICA, A., A. S. CHILTON, J. GOLDIN, K. ROZEMA, AND M. SEN (2016): "The political ideologies of law clerks," *American Law and Economics Review*, 19, 96–128.

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.

KATZ, D. M. AND D. K. STAFFORD (2010): "Hustle and Flow: A Social Network Analysis of the American Federal Judiciary," *Ohio State Law Journal*, 71, 08–004.

Appendix

Proofs

Proof of Proposition 1. Suppose that $j = 2$ deviates at $t = 1$ and approaches the best student at $t = 1$, say i (i.e., $s_i \geq s_{i'}$ for all i'). If this student accepts it, this judge's expected payoff (evaluated at $t = 1$) is s_i. Indeed, later we show that, if u_0 is low enough, this student surely accepts this offer of judge 2 at $t = 1$.

To see if this deviation is profitable, observe that, if $j = 2$ waits until $t = 2$, then one of the following three cases happens:

(Case 1) $\theta_i = 1$ and there is no other student with ability 1.

(Case 2) $\theta_i = 1$ and there is another student whose ability is 1.

(Case 3) $\theta_i = 0$.

Case 1 occurs with probability $s_i \prod_{i' \neq i}(1 - s_{i'})$, and in this case, judge $j = 2$ is better off by deviation by payoff 1. In Case 2, this judge's payoff is the same with or without deviation. Case 3 occurs with probability $1 - s_i$, and in this case, judge $j = 2$ is worse off by deviation by payoff at most 1. Therefore, this judge is better off by deviation in expectation at $t = 1$ if:

$$s_i \prod_{i' \neq i}(1 - s_{i'}) \geq 1 - s_i,$$

which is satisfied by a nonempty open interval of s, say $A^* \subseteq [0,1]^N$.

We complete the proof by showing that this student i surely accepts the offer from $j = 2$ at $t = 1$. If this student accepts the offer, then his payoff is 0. If he rejects it, then his payoff is u_0 in case $\theta_i = 0$ and two or more students have ability 1. Because this case occurs with a positive probability

$$(1 - s_i)\left(1 - \prod_{i' \neq i}(1 - s_{i'}) - \sum_{i''=1}^{N} s_{i''} \prod_{i' \neq i, i''}(1 - s_{i'})\right),$$

his expected payoff is negative if $-u_0$ is sufficiently large. More precisely, there exists u_0^* such that, for any $u_0 < u_0^*$ and any $s \in A^*$, judge $j = 2$ is better off by deviating at $t = 1$. $\square$

Proof of Theorem 1. Let $w^* \in W^*$ denote the optimal equilibrium payoffs. We first show that $t^{w^*} = (2,2)$.

Lemma 2. $t^{w^*} = (2,2)$.

Proof of Lemma 2: Suppose not. Then, t^{w^*} is either $(2,1),(1,2),(1,1)$. Let $t^{w^*} = (2,1)$, for example. Then, we have

$$w^* = (1-\delta)(c,b) + \delta\omega^{w^*}(2,1).$$

Because we have $w_1^* + w_2^* \geq w_1^M + w_2^M > c + b$, we must have $\omega_1^{w^*}(2,1) + \omega_2^{w^*}(2,1) > w_1^* + w_2^*$.

However, because $\omega_1^{w^*}(2,1)+\omega_2^{w^*}(2,1) \in W^*$, this contradicts that w^* is the optimal equilibrium.

The same holds even if $t^{w^*}=(1,2)$ or $(1,1)$. This proves the lemma. $\diamondsuit$

Recall that, by construction, IC(2,2) is binding in w^M, i.e., $w_2^M=(1-\delta)c+\delta\frac{bc-a_2d}{c-d+b-a_2}$. Because w^* satisfies IC(2,2) too, we have $w_2^* \geq w_2^M$.

Note also that w^M is in the Pareto frontier of W^0, because the players only play (2,1) and (2,2) on this equilibrium path. Because (2,2) achieves a higher instantaneous total payoff than (2,1), this implies that any w with $w_1+w_2 > w_1^M+w_2^M$ is not in W^0.

Therefore, we have $w^*=w^M$. $\square$

Proof of Lemma 1. Suppose the first condition does not hold. Then necessarily, if we take w so that it satisfies IC(2,2) with equality and that it lies in the Pareto frontier of W^0, then $w'=\frac{1}{\delta}((1-\delta)(a_1,a_2)-w)$ satisfies IC(2,1).

Similarly, suppose that the second condition does not hold either. Then necessarily, if we take w so that it satisfies IC(2,1) with equality and that it lies in the Pareto frontier of W^0, then $w'=\frac{1}{\delta}((1-\delta)(c,b)-w)$ satisfies IC(2,2).

Then, we can easily find q_1,q_2 so that Assumption 1 is satisfied. $\square$