Recognition and Performance: Evidence from a Regression Discontinuity in College Football

Lester Lusher[a]*, Matthew Naven[b]

[a]University of Hawaii at Manoa

[b]University of California, Davis

Abstract

A growing literature has investigated the impacts of informing agents about their relative performance and rank on effort and performance. We study the impact of national recognition in NCAA football on a team's future performance. Identification comes from utilizing a regression discontinuity design where within weeks we compare teams that received recognition against those which just barely did not receive recognition. In addition, the threshold for national recognition adjusts weekly which allows us to use across-week comparisons of similar teams. Across a variety of specifications, we find that team recognition significantly boosts a team's likelihood of winning the following week. Placebo tests on alternate rankings show no effect, further suggesting a true impact of team recognition.

Keywords: Feedback; Relative Performance; Rank; College Football

JEL codes: L83, D01

*Corresponding author: Lester Lusher, University of Hawaii at Manoa, Saunders Hall 542, 2424 Maile Way, Honolulu, HI 96822 United States. Email: lrlusher@hawaii.edu; Matthew Naven, Email: msnaven@ucdavis.edu

1 Introduction

Recent years have seen a large growth in studies investigating the relationship between recognition and performance. The primary motivation for understanding this relationship is straightforward: beyond extrinsic or monetary incentives, can simply providing rankings or feedback be used to further motivate agents? While the literature has only started to investigate this relationship in the past decade, many institutions in society have long implemented feedback systems under the assumptions that feedback boosts effort. Most notably, employers and educators often provide feedback on employees' and students' relative performance and achievement. Correspondingly, the literature has largely focused on the workplace or classroom environments and have generally found that recognizing employees/students by revealing relative performances and/or ordinal ranks leads to boosts in performance.[1]

In order to identify the effects of providing information on relative performance, the entirety of the previous literature has, to our knowledge, made a simple static juxtaposition between two settings: one where agents were recognized against one where agents were not recognized.[2] Though recognition came in different forms in these studies (e.g. recognizing the top x students/employees), their identification effectively compares the average performance of *all* agents (regardless of whether they were *ex post* recognized or not) in one population with recognition against another population of agents in a setting without recognition. Consequently, these studies answer the question "What happens to effort in response to the introduction of relative performance recognition?" A slightly different, yet unanswered, question of interest is "*Within* a setting of recognition, how does (a lack of) recognition affect future performance?" In other words, how does the subsequent effort of agents who performed just well enough to receive positive recognition compare to those who were close to, but did not receive, recognition? Understanding the answer to this question has

[1]These studies include Hoogveld and Zubanov (2017), Bradler et al. (2016), Jalava et al. (2015), Bandiera et al. (2015), Ashraf et al. (2014), Kuziemko et al. (2014), Charness et al. (2013), Bhattacharya and Dugar (2013), Tran and Zeckhauser (2012), Kuhnen and Tymula (2012), Kosfeld and Neckermann (2011), Azmat and Iriberri (2010), and Eriksson et al. (2009).

[2]A notable exception to this is a working paper from Moreira (2017) who, like us, utilizes a regression discontinuity design within a setting where agents are recognized. Moreira (2017) utilizes an educational setting in Brazil.

obvious implications for the potential long-run general equilibrium consequences of agent recognition. For instance, if *ex post* recognition increases subsequent effort, and increased effort begets future success and recognition, then the dynamic effort of agents who initially receive recognition could continually increase across periods.

This paper investigates the effects of recognition on subsequent effort and performance by utilizing nearly 10,000 games played by National College Athletics Association (NCAA) college football teams from 1990 to 2015. Each week, the Associated Press (AP) recognizes the top 25 teams based on their past and projected performances for the season. These teams are selected from a vote by sportswriters across the country, where each sportswriter's votes generate a number of vote-points for each team, and the 25 teams with the most vote-points are recognized nationally.

We implement a regression discontinuity design which simultaneously utilizes two dimensions of variation. First, *within* each week, we compare the subsequent performance of teams which just barely ranked in the top 25 against teams which just barely ranked outside the top 25. Second, *across* weeks, we contrast the subsequent performance of teams which ranked in the top 25 against teams which received the exact same amount of vote-points but which did not place in the top 25 that week. Our identification effectively utilizes a regression discontinuity design where the threshold for recognition adjusts every week based on the AP's votes.

Results suggest that being recognized increases a team's future performance. In our preferred estimates teams that are barely ranked in the AP poll outperform their barely unranked counterparts by about ten percentage points in the likelihood of winning their next game. This estimate isolates the effect of being recognized that is independent of a team's ability which, if unaccounted for, would bias traditional OLS estimates; our estimates are about 20 percentage points lower than one would find by running the naive OLS regression of winning your next game on being ranked. While our estimates are somewhat sensitive to bandwidth choice, we find that they are consistently positive and grow in magnitude as the bandwidth increases. Furthermore, placebo tests that estimate the impact of being ranked 5th, 10th, 15th, or 20th show no effect on future performance which suggests that we are indeed capturing the true impact of team recognition.

The remainder of this paper proceeds as follows. The next section describes our data and some institutional background about the AP NCAA college football rankings. Section 3 describes our econometric specifications and identification strategy. Section 4 presents out results, and section 5 concludes.

2 Data Source and Institutional Background

Data come from the Associated Press (AP) National College Athletics Association (NCAA) college football rankings. Each week, a panel of predetermined sportswriters provide their ranking of the top 25 college football teams in the nation. Team points are calculated from these ballots with a first place vote awarding 25 points, a second place vote awarding 24 points, and so on down to 1 point for a 25th place vote. After totaling the number of points for each team, the teams are ranked from 1 to 25 based on their total number of points, and this comprises the AP rankings. Every week, some teams receive votes, but do not receive enough total points to be in the top 25. These teams are unranked but listed as *receiving votes*. We have weekly data on 433 weeks of poll rankings dating back to the 1990 preseason poll. The data only include teams who received at least one vote during that week so some marginal teams may enter and exit the dataset from week to week.

The poll includes data on a team's record during the week of the poll. We use this information to calculate whether or not a team won their *next* football game that takes place the week after the poll is released. We calculate a team as winning if their current week's number of wins is greater than their prior week's number of wins. We make the assumption that if a team is in the dataset one week but disappears from the dataset the next week it is because the team lost their next game and for that reason received no votes. Importantly, each weekly rank from the AP poll gives no tangible benefit to teams beyond the recognition that they receive, so teams should not have an increased chance of performing better in the future beyond any mental benefits that accrue.

We use the number of total points a team receives as the running variable in our regression discontinuity (RD) analysis instead of team ranking so that there is variation in the relative position of teams who are all listed as receiving votes. Because the number of points necessary in order to be ranked 25th varies from week to week, we rescale the points variable such that it measures the points relative to the cutoff number of points required to be ranked 25th each week. We calculate this weekly cutoff by taking the average number of points for the 25th ranked team and the team with the highest number of points that is listed as receiving votes (i.e. the unranked team with the most points). As a robustness check, we include the number of actual points as a control variable because the rescaling prevents collinearity issues with our treatment indicator.

The initial dataset consists of 14,362 team-by-poll week observations. For our regression analyses we drop all observations in the final poll of each season (845 observations), as teams play no games after these polls. We also drop all observations who have a bye during the next week (2,141 observations), as these teams have no win, loss, or tie outcome in their next game. We drop any observation where a team played multiple games in the next week if the team's outcomes differed between the two games such as a win and a loss or a win and a tie (12 observations). In a few rare cases teams have fewer wins, losses, or ties than in the prior week, so we drop those observations as well as the team's observation from one week prior (10 observations), as this could contaminate the correct measurement of whether or not a team won their next game. Finally, we drop any observations that are missing our primary outcome of interest, winning your next game, as well as our running variable[3] (1,058 and 374 observations respectively). Our baseline sample before imposing bandwidth restrictions includes 9,922 observations from 385 poll weeks and 26 unique football seasons.

[3]For 21 of the 433 weeks of data, we cannot calculate the running variable because there are no unranked teams that received votes.

Table 1: Summary Statistics

	Rank					
	Unranked	21-25	16-20	11-15	6-10	1-5
Winning Percentage	.735	.779	.816	.853	.9	.971
	[.194]	[.177]	[.154]	[.138]	[.114]	[.0577]
Points	12.1	211	493	808	1,130	1,455
	[22.9]	[84.5]	[106]	[117]	[125]	[139]
Points from the Ranked Cutoff	-78.4	135	419	734	1,056	1,381
	[35.9]	[88.2]	[116]	[125]	[132]	[148]
Won Next Game	.393	.811	.733	.726	.751	.808
	[.489]	[.392]	[.442]	[.446]	[.433]	[.394]
Lost Next Game	.607	.186	.262	.27	.243	.186
	[.489]	[.389]	[.44]	[.444]	[.429]	[.389]
Tied Next Game	.000387	.00328	.00494	.00388	.00579	.00562
	[.0197]	[.0572]	[.0701]	[.0622]	[.0759]	[.0748]
Observations	2583	1221	1417	1547	1554	1600

2.1 Summary Statistics

Table 1 reports summary statistics for the baseline sample. By definition, unranked teams have fewer points that are below the cutoff for being ranked, so differences in points by ranking are not surprising. There is, however, a substantial difference in the likelihood of winning for ranked teams. While part of this is due to the fact that ranked teams are inherently better, the most noticeable difference is the fact that barely ranked teams (i.e. those ranked 21-25) have the highest winning percentage, even higher than the top five ranked teams. Furthermore, these barely ranked teams are about 40 percentage points more likely to win their next game than their unranked counterparts.

Figure 1 shows a binned scatter plot of the proportion of teams that won their next game by relative number of points. There are seven quantiles for the unranked teams in order to match the mean number of unranked teams in each poll, and the ranked quantiles were created to have approximately equal numbers of observations as the unranked quantiles. The scatter plot implies that there is a sharp discontinuity in the likelihood of winning your next game when you appear as a ranked team in the AP poll. While teams who

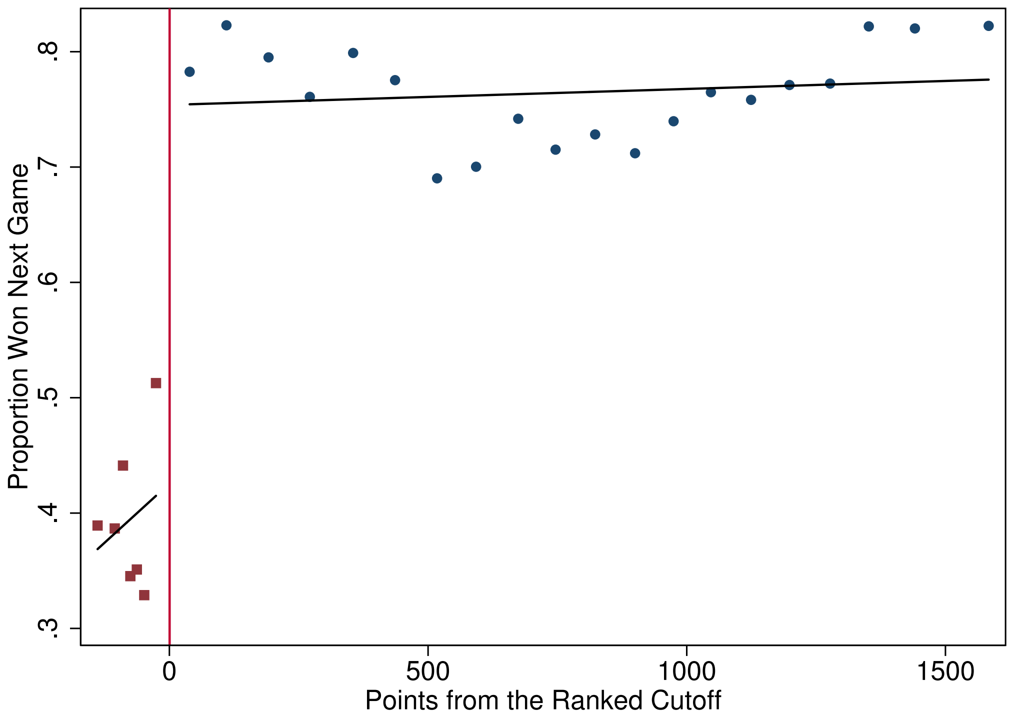

Figure 1: Proportion Won Next Game by Ranking

are just shy of being ranked win their next game on average about 55% of the time, the worst teams to get ranked win their next game on average about 75% of the time despite the fact that these teams received very similar point totals. It is important to note the the visual estimate of the discontinuity may differ from regression estimates due to the fact that the fitted line in the figure is estimated on the binned averages while our regression analysis is estimated on the underlying team microdata.

3 Econometric Specification and Identification

Due to the mechanical relationship between a team's performance and weekly ranking, a naive linear regression of winning your next game on being ranked will lead to a biased estimate of the impact of being ranked, as being ranked is endogenous to the omitted variable of team ability which will also impact the likelihood of winning. For this reason, we use a regression discontinuity (RD) design in order to isolate the pure effect of being ranked that is independent of a team's ability. An RD is based upon the assumption that potential outcomes are smooth around the cutoff where teams go from being unranked to ranked. This assumption essentially implies that treatment is randomly assigned at the limit of the cutoff between ranked and unranked teams. We use a local polynomial specification (Hahn et al., 2001; Porter, 2003) that takes the form

$$Y_{iwtc} = \alpha + f(p_{iw}) + \beta Ranked_{iw} + g(p_{iw}) Ranked_{iw} + \boldsymbol{X}_{iwtc} + \varepsilon_{iwtc} \tag{1}$$

where

$$Ranked_{iw} = \mathbb{1}[p_{iw} > 0]$$

and Y_{iwtc} denotes whether a team i in week w of season s and conference c won their next game, $f(p_{iw})$ and $g(p_{iw})$ are flexible polynomials in the number of points from the ranked cutoff, $\boldsymbol{X}_{iwtc}$ is a vector a covariates that includes the raw number of points received, season (i.e. year) fixed effects, and conference fixed effects[4], and ε_{iwtc} is a normally distributed idiosyncratic error term.

We do not constrain the derivatives of the polynomials to be equal on both sides of the cutoff, hence the separate functional forms for $f(p_{iw})$ and $g(p_{iw})$. Observations are weighted using a triangular kernel that puts more weight on observations closer to the cutoff, as this provides the optimal boundary correction (Cheng et al., 1997). We cluster standard errors by the running variable following Lee and Card (2008).

Evidence of manipulation to the number of points a team receives would violate the assumption that teams are essentially randomly assigned treatment status at the ranked cutoff. Because teams have no incentive to purposefully lose games[5] and there are a fixed number of ranked teams, there is little concern of

[4] Conference fixed effects are defined as "Power 5 Conference" fixed effects and include the following conferences: Atlantic Coast Conference (ACC), American Athletic Conference (AAC), Big 12, Big Ten, Pacific-12 Conference (Pac-12), Southeastern Conference (SEC), and other. Due to conference mergers and realignment, we assign teams in the Big 8 to the Big 12, teams in the Big East to the AAC, and teams in the Pac-10 to the Pac-12. The other category consists of the Big Sky, the Big West, Conference USA, the Mid-American Conference, the Mountain West Conference, the Patriot League, the Southwest Conference, the Sun Belt Conference, the Western Athletic Conference, and independents.

[5] In professional sports, teams may have an incentive to purposefully lose games when the draft order is partially determined by achievement, in which poorly performing teams are more likely to have an earlier position to draft new players. In college sports new players are recruited by the current coaching staff, and the best teams are generally rewarded with better recruits. Thus college teams do not face the same adverse incentive as professional teams.

manipulation of the running variable. Even so, statistically testing for the difference in densities at the cutoff (McCrary, 2008) will mechanically estimate a difference in densities due to the fact that only one team can be ranked 25th (the right limit of the discontinuity) but multiple teams can be unranked (the left limit of the cutoff). For this reason, we do not report results from the McCrary test.

4 Results

4.1 Ordinary Least Squares Results

Table 2 reports the naive OLS estimates from a regression of winning your next game on being ranked. The coefficient estimates suggest that being ranked has a large and statistically significant impact on your likelihood of winning your next game of about 30 percentage points. The effect is relatively stable regardless of the addition of additional covariates that may be correlated with the likelihood of winning your next game. This coefficient estimate combines two effects, however. The first is our effect of interest - the psychological effect of receiving recognition for performance. The second effect is due to team ability and confounds traditional OLS estimates - teams that are more talented will be more likely to win games and for this reason will also be more likely to be ranked. If our dataset consisted of all Division I football teams instead of just the subset of teams who received at least one vote in the AP poll, then the coefficient estimate would doubtlessly be even larger in magnitude due to the addition of many unranked, low-ability teams.

4.2 Regression Discontinuity Results

In order to isolate the effect of being ranked that is purely due to recognition and uncorrelated with team ability, we now report our main results that use an RD design in order to estimate the effect of being ranked. Table 3 reports coefficient estimates from a kernel weighted local linear regression. In this table bandwidths

Table 2: OLS Regression

	Won Next Game							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ranked	0.312***	0.278***	0.351***	0.312***	0.327***	0.277***	0.326***	0.326***
	(0.033)	(0.032)	(0.031)	(0.033)	(0.032)	(0.032)	(0.032)	(0.032)
Points Control	-	Y	-	-	Y	Y	-	Y
Year FE	-	-	Y	-	Y	-	Y	Y
Conference FE	-	-	-	Y	-	Y	Y	Y
Observations	9,922	9,922	9,922	9,922	9,922	9,922	9,922	9,922

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

are calculated separately below and above the cutoff[6] in order to minimize mean squared error (Calonico et al., 2014, 2016, 2017). Our baseline estimate in column one suggests that there is no significant impact of being ranked on the probability of winning your next football game. Our preferred estimate in column two, which controls for the raw number of points and therefore uses variation both within and across weeks, suggests that teams are a statistically significant 11 percentage points more likely to win their next game after being ranked. In fact, with the exception of only conference fixed effects, the coefficient estimates are all positive and statistically significant after including additional controls, with the effect reaching as large as 24 percentage points. While the effect size is rather large, especially given that the winning percentage for unranked teams is about 40%, it is noticeably smaller than both the OLS results in table 2 as well as the discontinuity suggested by table 1 and figure 1.

Table 4 illustrates why this is the case. The coefficient estimates come from the same specification as in table 3, but instead of calculating optimal bandwidths conditional on the included regressors using the underlying data we use the average unranked and ranked bandwidths across the eight specifications in table 3 for all regressions in table 4. By keeping the bandwidth constant we can isolate the change in the coefficient estimates that is purely due to the inclusion of covariates as opposed to the combined effect

[6] We estimate separate bandwidths on each side of the cutoff due to the fact that there are vastly more ranked teams than unranked teams.

Table 3: Regression Discontinuity Local Linear Regression with Optimal Bandwidth

	Won Next Game							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ranked	0.038	0.113^{**}	0.167^{***}	0.033	0.242^{***}	0.113^{**}	0.168^{***}	0.240^{***}
	(0.068)	(0.049)	(0.053)	(0.068)	(0.046)	(0.049)	(0.053)	(0.046)
Polynomial Degree	1	1	1	1	1	1	1	1
Points Control	-	Y	-	-	Y	Y	-	Y
Year FE	-	-	Y	-	Y	-	Y	Y
Conference FE	-	-	-	Y	-	Y	Y	Y
Unranked Bandwidth	48.6	82.8	64.8	48.8	86.1	84.6	65.6	86.8
Ranked Bandwidth	294	300	384	303	381	307	394	392
Effective Unranked Observations	549	1,495	988	549	1,600	1,555	1,011	1,607
Effective Ranked Observations	1,383	1,404	1,786	1,416	1,770	1,430	1,850	1,838

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

of including covariates and changing the size of the bandwidth. When using this constant bandwidth the baseline effect size in column one is larger and now statistically significant, and the inclusion of controls between specifications makes a much smaller difference in effect size. Thus the increase in the baseline effect size is likely due to the increased bandwidth, as the largest coefficient estimates in table 3 are the regressions that also have the largest optimal bandwidth.

Somewhat unintuitively, year fixed effects still impact the coefficient estimate even when using a constant bandwidth. This is due to the fact that there are more unranked teams receiving votes in later years and teams in our data are less likely to win their next game in later years, both perhaps due to increased parity amongst the top football programs. The omitted variable of year fixed effects biases the estimates downward due to the correlation with the dependent variable and the independent variable of interest.

As a robustness check, we also report results from a local quadratic regression. Table 5 reports results from this specification. For these regressions the bandwidth was again calculated optimally on both sides of the cutoff in order to minimize mean squared error. The local quadratic results are slightly less conclusive than the local linear estimates, as a few estimates are negative (although insignificantly so). Nevertheless, the specifications that include year fixed effects produce broadly similar results to our preferred specification in

Table 4: Regression Discontinuity Local Linear Regression with Constant Bandwidth

	Won Next Game							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ranked	0.102*	0.091*	0.190***	0.100*	0.183***	0.087	0.186***	0.178***
	(0.054)	(0.053)	(0.052)	(0.055)	(0.051)	(0.054)	(0.052)	(0.051)
Polynomial Degree	1	1	1	1	1	1	1	1
Points Control	-	Y	-	-	Y	Y	-	Y
Year FE	-	-	Y	-	Y	-	Y	Y
Conference FE	-	-	-	Y	-	Y	Y	Y
Unranked Bandwidth	71	71	71	71	71	71	71	71
Ranked Bandwidth	344	344	344	344	344	344	344	344
Effective Unranked Observations	1,193	1,193	1,193	1,193	1,193	1,193	1,193	1,193
Effective Ranked Observations	1,601	1,601	1,601	1,601	1,601	1,601	1,601	1,601

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

table 3 and indicate that being ranked increases your probability of winning by about 13 percentage points.

As a final specification test, we run local linear regressions using placebo cutoffs for the following ranks: 5th, 10th, 15th, and 20th. These estimates can be found in table 6. In these regressions, we rescale the running variable such that the cutoff occurs between the 5th and 6th ranked teams (columns 1 and 2), the 10th and 11th ranked teams (columns 3 and 4), the 15th and 16th ranked teams (columns 5 and 6), and the 20th and 21st ranked teams (columns 7 and 8). If there is evidence of an effect at these placebo cutoffs, this would cast significant doubt upon the effects that we found in our prior tables. Encouragingly, we find no evidence of a discontinuity at the placebo cutoffs, as the estimates are all statistically insignificant and vary in sign. This gives confidence that our estimates reflect the true impact of being recognized for achievement and do not reflect some unobserved omitted variable.

Lastly, figure 2 and table 7 present evidence that team characteristics vary smoothly across the threshold from unranked to ranked. Figure 2 is analogous to figure 1 but instead of plotting the likelihood of a team winning their next game instead plots the teams winning percentage that week. Although the linear fit of the binned data suggests that there is a small discontinuity in winning percentage between ranked and unranked teams, the line of best fit appears to be heavily influenced by the bin of teams with the fewest points relative

Table 5: Regression Discontinuity Local Quadratic Regression with Optimal Bandwidth

	Won Next Game							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ranked	-0.013	0.040	0.134**	-0.013	0.142**	0.040	0.130*	0.139**
	(0.070)	(0.063)	(0.067)	(0.071)	(0.065)	(0.064)	(0.067)	(0.066)
Polynomial Degree	2	2	2	2	2	2	2	2
Points Control	-	Y	-	-	Y	Y	-	Y
Year FE	-	-	Y	-	Y	-	Y	Y
Conference FE	-	-	-	Y	-	Y	Y	Y
Unranked Bandwidth	93.6	113	94.9	94	98.2	114	95.4	98.9
Ranked Bandwidth	504	511	641	512	652	515	622	630
Effective Unranked Observations	1,731	2,217	1,749	1,742	1,833	2,234	1,759	1,838
Effective Ranked Observations	2,322	2,349	2,965	2,357	3,014	2,370	2,889	2,924

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Placebo Regression Discontinuity Local Linear Regression with Optimal Bandwidth

	Won Next Game							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ranked	-0.002	-0.002	-0.025	-0.020	0.041	0.049	-0.060	-0.059
	(0.043)	(0.043)	(0.034)	(0.034)	(0.035)	(0.035)	(0.043)	(0.042)
Placebo Cutoff Rank	5	5	10	10	15	15	20	20
Polynomial Degree	1	1	1	1	1	1	1	1
Points Control	-	Y	-	Y	-	Y	-	Y
Year FE	-	-	-	-	-	-	-	-
Conference FE	-	-	-	-	-	-	-	-
Unranked Bandwidth	291	291	251	265	213	202	88.3	92.8
Ranked Bandwidth	107	107	297	287	431	436	451	460
Effective Unranked Observations	1408	1408	1209	1270	1037	1001	459	485
Effective Ranked Observations	551	551	1449	1404	2042	2065	2099	2133

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

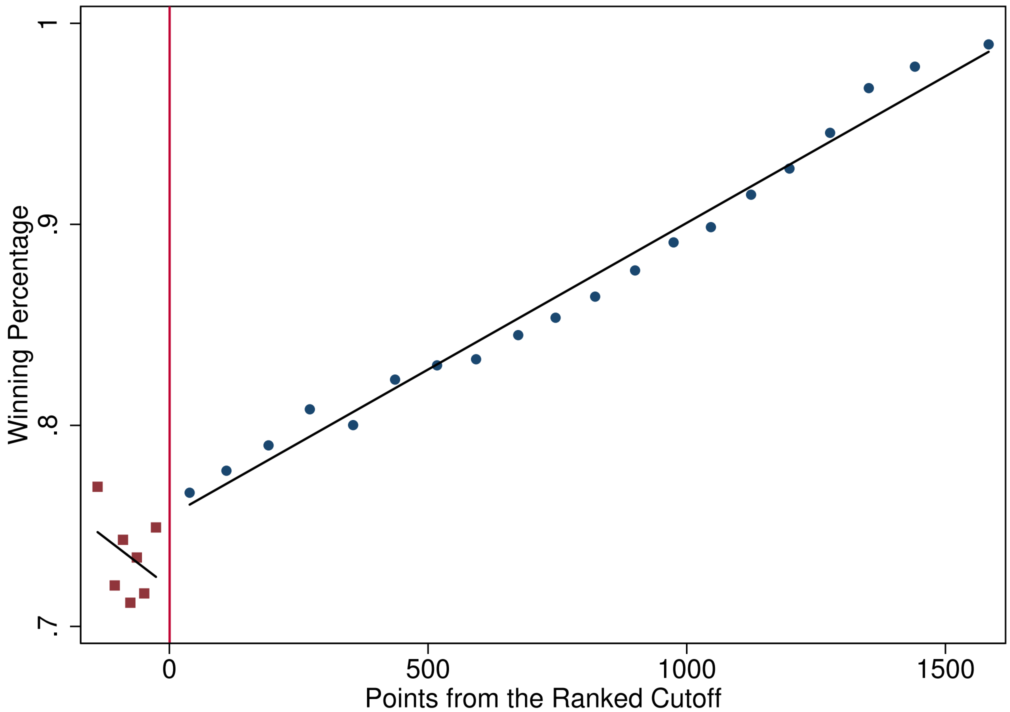

Figure 2: Winning Percentage by Ranking

to the cutoff. A visual inspection of the data shows no discontinuous jump in winning percentage at the cutoff. Table 7 confirms that this is the case using the underlying microdata, as the effect of being ranked on winning percentage is small, insignificant, and negative, which would bias our results toward zero if anything. We do not report results that include year fixed effects due to the fact that we are unable to compute the optimal bandwidth below the threshold when including year fixed effects.

5 Conclusions

A growing strand of literature has empirically identified the impacts of ranking and recognition on agent performance. To our knowledge, the entirety of these studies have made a static juxtaposition between two settings: one where all agents were (potentially) recognized versus a setting where all agents were not recognized. This study is the first to causally identify the impacts of being ranked (vs. unranked) *within* a

Table 7: Regression Discontinuity Local Linear Regression with Optimal Bandwidth

	Won Next Game			
	(1)	(2)	(3)	(4)
Ranked	-0.007	-0.012	-0.010	-0.018
	(0.020)	(0.019)	(0.020)	(0.020)
Polynomial Degree	1	1	1	1
Points Control	-	Y	-	Y
Year FE	-	-	-	-
Conference FE	-	-	Y	Y
Unranked Bandwidth	83.8	89.7	83	82.1
Ranked Bandwidth	492	481	499	484
Effective Unranked Observations	999	1,075	962	950
Effective Ranked Observations	2,070	2,023	2,107	2,034

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

setting with recognition. That is, compared to a counterfactual of just barely not receiving recognition, we identify how subsequent effort and performance responds to recognition.

To do this, we utilize nearly 10,000 games played by National College Athletics Association (NCAA) college football teams from 1990 to 2015. Every week, the top 25 teams are recognized by the Associated Press (AP) based on their past and projected performances for the season. Underlining these rankings is a running variable of vote-points, which are generated by a series of national sportswriters; every week the top 25 teams with the most vote-points are ranked and recognized one through 25. We implement a regression discontinuity design which simultaneously utilizes two dimensions of variation: within each week, we compare the subsequent performance of teams which just barely ranked in the top 25 against teams which just barely ranked outside the top 25, and across weeks we compare the performance of teams which ranked in the top 25 against teams which received the exact same amount of vote-points but did not place in the top 25.

We find that being recognized increases a team's future performance. In our preferred estimates teams that are ranked in the AP poll experience a ten percentage point increase in the likelihood of winning their next game when compared to their barely unranked counterparts. By utilizing a regression discontinuity design, this estimate isolates the effect of being recognized that is independent of a team's ability which, if unaccounted for, would bias traditional OLS estimates. Our results are consistent across a series of robustness checks, bandwidth choices, and placebo tests.

These findings suggest an additional benefit to ranking and recognition above and beyond those made from previous studies. While prior studies found increases in effort and performance in the classroom/workplace in response to the *possibility* of receiving recognition (i.e. statically juxtaposing all agents in a setting where some agents may receive recognition vs. a setting with no recognition), our findings suggest that, within a setting of recognition, agent performance is further impacted by receiving recognition *ex post*. This has several implications for the long-run general equilibrium impacts of recognizing agents. Namely, in all settings with repeated interactions and evaluations, receiving recognition will lead to further increased effort and performance in subsequent periods. For instance, our findings plausibly suggest that in an educational setting students who *ex post* do well on examinations and receive recognition will subsequently increase their effort on future exams.

References

ASHRAF, N., O. BANDIERA, AND S. S. LEE (2014): “Awards unbundled: Evidence from a natural field experiment,” *Journal of Economic Behavior & Organization*, 100, 44–63.

AZMAT, G. AND N. IRIBERRI (2010): “The importance of relative performance feedback information: Evidence from a natural experiment using high school students,” *Journal of Public Economics*, 94, 435–452.

BANDIERA, O., V. LARCINESE, AND I. RASUL (2015): “Blissful ignorance? A natural experiment on the effect of feedback on students' performance,” *Labour Economics*, 34, 13–25.

BHATTACHARYA, H. AND S. DUGAR (2013): “Contests for ranks: Experimental evidence,” *Southern Economic Journal*, 79, 621–638.

BRADLER, C., R. DUR, S. NECKERMANN, AND A. NON (2016): “Employee recognition and performance: A field experiment,” *Management Science*, 62, 3085–3099.

CALONICO, S., M. D. CATTANEO, AND M. H. FARRELL (2017): “On the effect of bias estimation on coverage accuracy in nonparametric inference,” *Journal of the American Statistical Association*.

CALONICO, S., M. D. CATTANEO, M. H. FARRELL, AND R. TITIUNIK (2016): “Regression discontinuity designs using covariates,” *URL http://www-personal. umich. edu/~ cattaneo/papers/Calonico-Cattaneo-/Farrell-Titiunik_2016_wp. pdf*.

CALONICO, S., M. D. CATTANEO, AND R. TITIUNIK (2014): “Robust nonparametric confidence intervals for regression-discontinuity designs,” *Econometrica*, 82, 2295–2326.

CHARNESS, G., D. MASCLET, AND M. C. VILLEVAL (2013): “The dark side of competition for status,” *Management Science*, 60, 38–55.

CHENG, M.-Y., J. FAN, J. S. MARRON, ET AL. (1997): “On automatic boundary corrections,” *The Annals of Statistics*, 25, 1691–1708.

ERIKSSON, T., A. POULSEN, AND M. C. VILLEVAL (2009): “Feedback and incentives: Experimental evidence,” *Labour Economics*, 16, 679–688.

HAHN, J., P. TODD, AND W. VAN DER KLAAUW (2001): “Identification and estimation of treatment effects with a regression-discontinuity design,” *Econometrica*, 69, 201–209.

HOOGVELD, N. AND N. ZUBANOV (2017): “The power of (no) recognition: Experimental evidence from the university classroom,” *Journal of Behavioral and Experimental Economics*, 67, 75–84.

JALAVA, N., J. S. JOENSEN, AND E. PELLAS (2015): “Grades and rank: Impacts of non-financial incentives on test performance,” *Journal of Economic Behavior & Organization*, 115, 161–196.

KOSFELD, M. AND S. NECKERMANN (2011): “Getting more work for nothing? Symbolic awards and worker performance,” *American Economic Journal: Microeconomics*, 3, 86–99.

KUHNEN, C. M. AND A. TYMULA (2012): “Feedback, self-esteem, and performance in organizations,” *Management Science*, 58, 94–113.

KUZIEMKO, I., R. W. BUELL, T. REICH, AND M. I. NORTON (2014): ““Last-place aversion”: Evidence and redistributive implications,” *The Quarterly Journal of Economics*, 129, 105–149.

Lee, D. S. and D. Card (2008): "Regression discontinuity inference with specification error," *Journal of Econometrics*, 142, 655–674.

McCrary, J. (2008): "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of econometrics*, 142, 698–714.

Porter, J. (2003): "Estimation in the regression discontinuity model," *Unpublished Manuscript, Department of Economics, University of Wisconsin at Madison*, 5–19.

Tran, A. and R. Zeckhauser (2012): "Rank as an inherent incentive: Evidence from a field experiment," *Journal of Public Economics*, 96, 645–650.