

## **Assessing Racial Differences in the Tenures of College Football Coaches Using Survival Analysis**

Mac Wetherbee \*

### **Abstract**

It is widely suspected that ethnic/racial minority college football coaches are treated differently in the job market than their White counterparts. In this paper, I test whether minority head football coaches have shorter mean tenure lengths than White head coaches. I first build a new dataset consisting of all Division I Football Bowl Subdivision (FBS) college football head coaches that coached at least one season between 1979 and 2019, expanding the sample from previous studies in this area. I present descriptive statistics for this new dataset. I then test the relationship between ethnicity and tenure length using the Cox proportional-hazards model. Even after controlling for relevant covariates, results indicate that there is a difference in mean tenure lengths between minority coaches and White coaches. Minority head football coaches are significantly more likely to have their coaching tenures end at any given length than White head football coaches are, all else equal. Due to limited sample sizes for other ethnic groups, this difference is driven primarily by the difference in tenure lengths between White and Black college football coaches.

**Key Words:** College Football, NCAA, Survival Analysis, Racial Differences, Cox Model, Unfair Firing Practices

### **1. Introduction & Literature Review**

In this paper, I use survival analysis to address the question of whether college football coaches from minority ethnic/racial groups—hereafter sometimes referred to simply as minority coaches—have shorter tenures (are fired more quickly) than their White counterparts. I first present a newly constructed dataset that includes both a coach's ethnicity and his tenure length (how long he lasted in the position) as well as various relevant covariates. I then conduct the survival analysis, describe the results, and discuss the implications.

Previous research on the hiring and firing of athletic coaches, especially in those studies focused on potential ethnic discrimination, has found mixed and conflicting results. Several authors (e.g. Day, 2018; Cook and Glass, 2013; Bozeman and Fay, 2013; Kopkin, 2014) find evidence of some sort of discrimination or difference between the careers of White coaches and the careers of minority coaches. Others find no evidence of race being an important factor (e.g. Solow et al., 2011; Holmes, 2011). Yet a third category finds that minority coaches might actually receive benefits from their minority status (e.g. Mixon and Trevino, 2004).

The evidence of minority status being a hindrance in the job market broadly falls into two categories. First, there is evidence that the career trajectories of minority coaches and White coaches differ well before they reach the higher levels of the coaching profession. According to this logic, these early career differences explain much of the difference in the coaches' later outcomes. Day (2018, p. 1) uses sociological sequence analysis to show that “white coaches are more likely to follow upward career trajectories while black coaches are more likely to get stuck in careers characterized by low-level positions.” His results show that, even given the same pre-coaching experience, minority coaches are more likely to remain at the lower levels of the coaching hierarchy than their White counterparts, who are more likely to get promoted. Thus, the differing propensities for Black coaches and White

---

\*University of Denver

coaches to reach the head coaching ranks have more to do with differences in the positions and promotions they receive early in their careers, rather than (only) discrimination at the upper ranks of the coaching profession.

Bozeman and Fay (2013) provide a potential mechanism to explain this. They argue that the lack of minority head coaches is due to differences in their early career positions. According to this logic, Black coaches and White coaches tend to have different lower-level positions when they are first hired as assistant coaches. White coaches tend to coach certain positions (e.g. quarterbacks in football), while minority coaches tend to coach different positions (e.g. running backs in football). Quarterback is a more prominent position, so the coach who coaches that position is more likely to get noticed and therefore promoted. Thus, there is a path dependent career trajectory for both White and minority coaches: their future careers depend heavily on which positions they coach early in their careers, and the positions that they coach are affected by their ethnicities. This, of course, leads to inequities throughout the careers of the coaches.

The other category consists of the research showing discrimination specifically at the head coach level. Cook and Glass (2013) analyze the types of college basketball teams that coaches of different ethnicities are hired to coach. They find that minority coaches are most likely to receive head coaching positions for two types of teams: historically Black colleges and universities (HBCUs) and teams with losing records in the previous season(s). The HBCUs are tangential to the point here, but the teams with losing records provide important evidence that minority coaches may not be given equal opportunity. If minority coaches primarily receive job opportunities with teams that have historically been unsuccessful, the coaches may face an uphill battle to show their skills. In turn, minority coaches could appear to be less successful than they otherwise would/could have been, leading to a decreased probability of receiving a second opportunity elsewhere. Furthermore, the authors provide evidence that, once a minority coach is fired, the school is more likely to replace the coach with a White coach (the 'savior effect').

Similarly, Kopkin (2014) finds that a Black college football coach is 5.28% more likely to be fired in any given year than a White coach with the same experience. Interestingly, he finds that this difference is statistically significant in the first three years of a coach's tenure and in the seventh and eighth years of a coach's tenure, but it is not statistically significant in years 4-6. It is possible that this is an artifact of the individual dataset, but it would be interesting to test this relationship to see if it can be replicated. Regardless, his study provides another example of potential discrimination against minority coaches at the head coaching level.

On the other hand, at least one study has found the opposite effect: that minority coaches actually receive benefits from their minority status. Mixon and Trevino (2004) found that Black coaches were 9.6% less likely to be fired than non-Black coaches in a given year, though the authors specifically point out that this does not rule out the possibility of discrimination in the hiring process. This result directly contradicts that of Kopkin (2014). Hopefully further research (including this paper) can help to adjudicate, reconcile, and explain these contradictory findings.

Other authors have found no evidence of ethnicity having an effect on coaching decisions in either direction. For example, Solow et al. (2011) find that race is not a statistically significant predictor of whether a coordinator (the level below head coach) will be promoted to head coach in professional football. Thus, "conditional on reaching coordinator status, there is no evidence that race influences head coach hiring decisions" (Solow et al., 2011, p. 332). They also found that the 'Rooney Rule,' which was implemented to require teams to interview minority candidates for head coaching positions, did not actually increase the number of minority head coaches. Holmes (2011) similarly does not find evidence of a

relationship between ethnicity and the probability of a coach being fired, though he mostly treats ethnicity as a control variable.

Other research has found similarly interesting patterns—not directly related to ethnicity—with respect to the hiring and firing of coaches. Avery et al. (2016) find that better academic performance of student-athletes decreases the probability that a coach will be fired, but does not affect the probability of receiving a job offer from a different school. Halgin et al. (2020) find that ties to an elite coach help a younger coach get promotions, but only if the media recognizes such ties. Holmes (2011) finds that recent success decreases the probability of a coach being fired, but older success (more than three years in the past) actually increases the probability of a coach being fired. This is presumably because the coach created higher expectations than had previously existed at the school. Finally, Fee et al. (2006) find that the probability of a head coach being fired is based mostly on team performance, while an assistant coach's probability of promotion to head coach is based primarily on individual-level characteristics and individual success.

## 2. Dataset

### 2.1 Dataset Construction

I construct the dataset of coaches from scratch. I include in the sample every head coach at an NCAA Division I FBS (formerly known as Division I-A) football program who coached at least one year after 1978. I select 1979 as the starting year because this was the year that the first Black coach was hired at the Division I FBS level. Because the dataset includes all coaches who coached in 1979 or later, it includes some coaches who were hired before 1979, but whose tenures included 1979. I include variables for the name of the school where a coach coached, the coach's name, the year that the coach was hired, the coach's race/ethnicity, and how long the coach remained at that school. Although I do not utilize it in this paper, I also include a variable in the dataset for the reason that the coach left his position. I also supplement these data with variables that I ultimately use as control variables: the coach's age at the time of his first season with the team, how many years the coach spent as an assistant coach prior to being hired, how many years the coach spent as a head coach prior to being hired, the coach's winning percentage during his tenure, the team's historical winning percentage prior to the coach's arrival, the average number of points scored per game by the team during the coach's tenure, the average number of points allowed per game by the team during the coach's tenure, how many bowl/playoff games the team won during the coach's tenure, and whether the coach is an alumnus of the school. The data are current as of the end of the 2019-2020 football season.

To collect these data, I used various sources. These included newspaper articles, college athletic department websites, wikipedia, sports-reference.com, and the ESPN College Football Encyclopedia (MacCambridge, 2005). Individuals for whom only a birth year (not a specific date) could be found are assumed to have a birthday prior to September 1st.

I include coaches who were hired to start with the 2020-2021 season. However, even though they are listed in the dataset, they are not included in the analysis. This occurs for two reasons: first, data collection started before the start of the 2020-2021 college football season started, so there was not yet data on the coaches' performance in the 2020 season. Second, the 2020 season was affected by the Covid-19 pandemic. Some teams did not play at all, some teams played limited schedules, and some teams played full or nearly full schedules. Some players opted out, and some coaches were more comfortable with players opting out than others. Thus, the 2020 season may not have been equivalent to other seasons in terms of coaches' probabilities of retaining their position. It is therefore

preferable to stop the analysis prior to this pandemic-affected season.

The data is inherently structured using a coach-tenure unit of analysis, with each coach-tenure representing one observation. Thus, this dataset is time-invariant: each variable is assumed to be constant throughout the tenure of a coach. If a coach later receives another position, then the values of the variables might be different for that observation than they were for the first tenure.

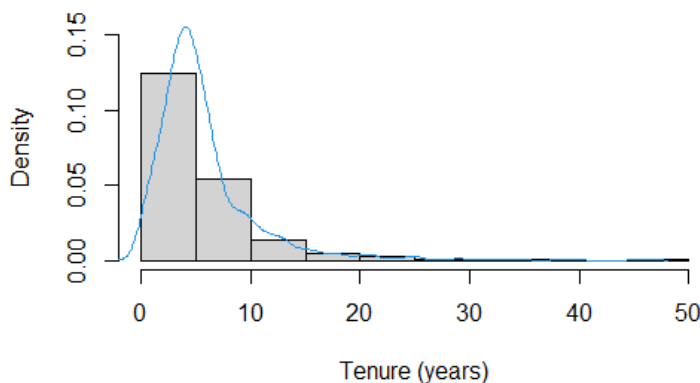
## 2.2 Descriptive Statistics

Descriptive statistics for the overall sample and by group are provided in Tables 1 and 2 below. Most notable, perhaps, are the sample sizes for each group in Table 1. College football coaches have overwhelmingly been White, with over 90% of the sample being comprised of White coaches. Easily the second largest group is Black head coaches at about 7.5% of the sample. The other three ethnicities—Hispanic, Pacific islander, and multiple races—each have fewer than ten observations.

Other than Pacific Islanders, who have a sample size of  $n=2$ , the White coaches have the longest mean tenure length at 5.89 years. This is at least two years greater than the mean tenure lengths for each of the Black, Hispanic, and multiple race categories. The standard deviation for the Black coaches is also notable: the spread of their tenure lengths is much lower than the spread for White and Hispanic coaches. This may be an artifact of the dataset, or it may show that they have fewer coaches with very long tenures than the White, Hispanic, and Pacific Islander categories.

**Table 1:** Tenure Lengths by Ethnicity

	Mean (Tenure Length)	St. Dev. (Tenure Length)	n
Full Sample	5.70	4.68	948
White	5.89	4.80	864
Black	3.55	1.98	71
Hispanic	3.89	5.13	9
Pacific Islander	8.00	5.66	2
Multiple Race	3.25	0.35	2



**Figure 1:** Histogram of Tenure Length

**Table 2:** Descriptive Statistics

	Min	Q1	Median	Q3	Max	Mean	St. Dev.
Year Hired	1958	1985	1998.5	2011	2020	1997.65	14.41
Tenure	0	3	5	7	46	5.70	4.68
Coach Age	29	40	44	49	72	44.72	7.16
Asst. Coach Years	0	9	13	19	39	13.73	6.55
Head Coach Years	0	0	2	7	30	4.35	5.77
Win %	0	0.36	0.49	0.61	0.93	0.48	0.18
Historical win %	0	0.50	0.55	0.60	0.79	0.55	0.09
Poins/game	8.73	20.52	24.54	29	46.86	24.96	6.17
Allowed/game	10.18	21.53	25.65	29.79	52.67	25.83	5.95
Bowl Wins	0	0	0	1	24	1.18	2.15
Alum	0	0	0	0	1	0.14	0.34

### 3. Research Questions

1. Does the tenure length of college football coaches differ by ethnicity?
2. If tenure length does differ by ethnicity, what is the difference in expected coaching duration between different ethnicities?

### 4. Methods

To answer these questions, I use the Cox proportional-hazards model. The coach's ethnicity acts as the primary independent variable of interest and the tenure length acts as the right-censored dependent variable. Because of the limited sample size in each ethnicity category other than White and Black, I collapse all of the minority groups together into a single 'Minority' dummy variable. Other independent variables are included as control variables to limit possible confounding effects.

The first set of control variables accounts for a coach's age and experience. I include a variable for the coach's age as of September 1st of his first season. A second control variable in this category is the number of seasons that the coach spent as an assistant coach prior to being hired for that position. The third control variable in this category represents the number of seasons that the coach spent as a head coach prior to being hired for that position. For these coaching experience measures, I count any seasons spent coaching football at the collegiate or professional levels. I do not include experience spent coaching a different sport or experience coaching at the high school level.

A second set of control variables accounts for the coach's success during his tenure. The first control variable in this category measures the team's winning percentage during the coach's tenure. I calculate this by dividing the number of wins by the sum of the wins and losses ( $win\% = \frac{wins}{wins+losses}$ ), effectively ignoring ties. Similarly, I also include variables for the average number of points scored and the average number of points allowed during the coach's tenure. Finally, I also include a variable for the percentage of seasons during the coach's tenure that ended with a bowl or playoff win.

Finally, I include two contextual variables that may also affect a coach's probability of being fired or retained. The first is a dummy variable for whether the coach is an alumnus of the school. The second is a measure of the team's historical winning percentage in all seasons before the coach was hired. This second variable should provide a measure of ex-

pectations: teams with higher historical winning percentages likely have loftier expectations than those with lower historical winning percentages.

After conducting the survival analysis itself, I interpret the results in terms of statistical significance and in terms of substantive effect sizes. I also compute expected durations for both White and minority coach categories. This may provide a more substantively meaningful depiction of the differences between the categories.

## 5. Results

### 5.1 Survival Analysis Results

Results from the Cox proportional-hazards model indicate support for the hypothesis that the tenure lengths of minority coaches are shorter than the tenure lengths of White coaches. This result is statistically significant at a conventional  $\alpha = 0.05$  level. The model controls for relevant covariates that measure context, coaching experience, and team success, meaning that the statistically significant difference by ethnicity is not due to these other factors. Thus, we have evidence that, all else equal, minority college football coaches do not last as long in their positions as White coaches do. We can see these results in Table 3.

We can also assess the substantive magnitude of this effect using the  $e^{coef}$  column. This odds ratio tells us the increase (or decrease) in probability of having the tenure end at any given point that would occur with a one-unit change in the predictor variable. Thus, the odds ratio of 1.388 for minority status tells us that minority coaches are 38.8% more likely to have their tenures end at a given point. This is a noticeable and substantively meaningful difference between the groups, providing additional compelling evidence that minority coaches have shorter tenure lengths even when controlling for relevant covariates.

Several of the control variables also achieve statistical significance. Most notably, the school's historical winning percentage has a very small p-value and a very large odds ratio. The positive coefficient and odds ratio larger than one tell us that schools/teams with higher historical winning percentages are much more likely to have their coach's tenure end. This is likely due to high expectations: teams that have been historically successful are more willing to remove an unsuccessful coach (or, similarly, that the threshold for a coach to be considered successful at that program is higher).

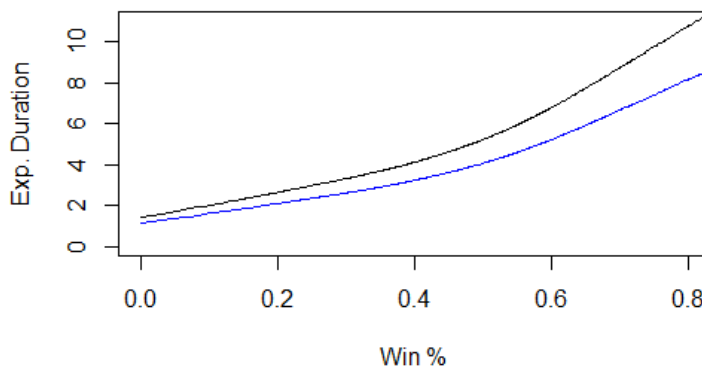
Furthermore, a coach's winning percentage during his tenure and the percentage of seasons that end with a bowl/playoff win both have statistically significant negative coefficients (and corresponding odds ratios below one). Thus, unsurprisingly, a coach's success during his tenure is associated with longer tenures: those who are more successful also last longer. None of the other control variables achieve statistical significance.

### 5.2 Expected Durations

We can also generate expected durations from the Cox model. For many people, this is more intuitive than thinking about odds ratios and/or p-values. To compute these expected durations, I use the `coxed` package in R (Kropko and Harden, 2020, 2019; Harden and Kropko, 2019). In Figure 2, we can see that the expected durations for the minority coaches are lower than the expected durations for the White coaches at all winning percentages. It is also true that the difference in expected duration between the two groups is larger, in absolute terms, at higher values of winning percentage. In substantive terms, this would mean that the effect of being a minority on tenure length is most noticeable for successful coaches: coaches with lower winning percentages are likely to have their tenures end quickly regardless of ethnicity.

**Table 3: Cox Model Results**

Predictor	Coef	$e^{coef}$	SE	P-value
Minority	0.328	1.388	0.137	0.017
Age	0.009	1.009	0.008	0.269
Ast. Coach Years	-0.000	1.000	0.008	0.987
Head Coach Years	-0.002	0.998	0.010	0.828
Win %	-2.916	0.054	0.846	0.001
Hist. Win %	2.580	13.196	0.451	< 0.000
PPG	0.024	1.024	0.017	0.157
Opp. PPG	0.021	1.022	0.017	0.197
Alum	-0.117	0.889	0.105	0.265
Bowl Win / Season	-0.820	0.441	0.277	0.003

**Figure 2: Expected Durations**

For example, take a coach with a winning percentage of 65%. If this coach is White, then the model predicts that the coach would average 7.74 years in his position. However, if he is a minority, then the model would expect a duration of 5.92 years. This difference of almost two years is, once again, substantively noticeable.

## 6. Discussion

### 6.1 Implications

The results suggest that there is a systematic difference in the mean tenure lengths of minority coaches and White coaches. Specifically, White coaches last longer in their positions than minority coaches, all else equal. This difference is substantively and statistically significant.

However, this result does not *necessarily* show that discrimination is present. In other words, even though we have convincing evidence that the mean tenure length is shorter for minorities, we do not have specific evidence showing the reason for this difference. More specifically, we do not have evidence that there is bias or discrimination leading to this difference.

Nonetheless, the result is suggestive. By controlling for coach quality/success, con-

textual factors, and the team's historical expectations, the model has controlled for some of the potential alternative explanations. We have evidence that the differences in tenure lengths are not due to different levels of success or different levels of expectations. Thus, in the absence of alternative explanations, one might conclude that discrimination is the most likely explanation for the observed relationship.

It is also worth noting that this analysis already ignores any bias and discrimination in the hiring process—I only compare coaches who have actually been hired into head coaching positions. Thus, the schools that hire minority coaches might actually be the ones who are less discriminatory, since they actively made the decision to hire a minority coach. If minority coaches still have shorter tenure lengths at these schools, this points to the likelihood that the situation could be even more extreme at other schools that have so far not made the decision to hire a minority coach.

## 6.2 Potential Limitations

The main limitation of this study is that it treats coaches leaving their positions as monolithic. In reality, there are many reasons that a coach might end his tenure at a specific college: he might be fired, he might accept a position with a different college/team, he might retire, etc. The analysis in this paper ignores this and implicitly treats each coach leaving his position as if it were the school's decision (i.e. the coach was fired). This assumption would be particularly problematic if the reasons that coaches leave their positions are associated with ethnicity. Unfortunately, this could be possible. This is more reason to consider the results in this paper to be suggestive of, rather than proof of, racial discrimination.

A much smaller limitation is the need to group all minority groups together into a single 'minority' category. In an ideal world, we would be able to look at each ethnic group as its own category. In this study, however, there were simply not enough college football coaches in many minority categories (e.g. two Pacific Islander coaches) to consider them separately. Grouping them together as a single category makes the assumption that the effect of being a minority is uniform across the different ethnic groups. This assumption is likely somewhat unrealistic, but it is nonetheless necessary due to the unequal group sizes. This is the consequence of relying on observational rather than experimental data.

## 6.3 Directions for Future Research

The best way to address the main limitation above would be to use competing risks survival analysis instead of traditional survival analysis. Like traditional survival analysis, competing risks models are more typically used in biomedical research in which there are multiple possible different causes of death (see Austin et al., 2016). However, such models can be applied to any situation in which there are multiple different possible outcomes that can occur at given time lengths. In the context of coaching tenure, the model would estimate the probability of each different possible reason for a coach leaving his position at a given tenure length based on the individual's covariate profile. For example, based on the coach's ethnicity, winning percentage, etc., what is the probability that the coach is fired, takes a new job, or continues in his current position?

This type of model would address the limitation of not considering the reasons that a coach leaves his position. Thus, if findings continued to show an increased estimated probability of minority coaches being fired (not just having their tenures end), this would provide further support for the idea that there may be some discrimination in the firing processes of college football coaches. Conversely, if a competing risks model did not show such a difference with respect to being fired, then we may conclude that the statistically



significant differences in mean tenure length shown in this paper are likely due to factors other than bias and discrimination.

Additionally, further research can also build upon this study by branching out from football and basketball, on which most studies of athletic coaching tenure have focused. Future studies could assess whether the results change if we focus on baseball, gymnastics, hockey, tennis, etc. This would provide useful knowledge on whether potential discrimination is homogenous across sports or varies by sport.

Qualitative and survey research could also be used to better understand the causal mechanisms behind the differences that I find in this study. Although we may try to infer bias as the causal mechanism, we only actually have evidence that there is a difference in mean tenure length. Qualitative research and survey research could both dive more deeply into the causal mechanism behind this to understand *why* minority coaches have shorter tenure lengths.

Finally, future research could address whether the effect of being a minority has changed over time. It seems plausible that potential discrimination against minority coaches may have systematically changed over time—likely (though not necessarily) that it has decreased over time. It would therefore be useful to understand this change. Then, if there is a change across time periods, the magnitude of this change would be relevant, as would the extent to which the effect is still non-zero. In other words, even if discrimination has decreased, it might still be present. Measuring it across time would provide a useful perspective on whether it has improved, and how much further there is still left to improve.

### References

- Austin, P. C., Lee, D. S., and Fine, J. P. (2016). Introduction to the Analysis of Survival Data in the Presence of Competing Risks. *Circulation*, 133(6):601—609.
- Avery, C., Cadman, B., and Cassar, G. (2016). Academics vs Athletics: Career Concerns for NCAA Division I Coaches.
- Bozeman, B. and Fay, D. (2013). Minority football coaches' diminished careers: Why is the "pipeline" clogged? *Social Science Quarterly*, 94(1):29–58.
- Cook, A. and Glass, C. (2013). Glass Cliffs and Organizational Saviors: Barriers to Minority Leadership in Work Organizations? *Social Problems*, 60(2):168–187.
- Day, J. C. (2018). Climbing the ladder or getting stuck: An optimal matching analysis of racial differences in college football coaches' job-level career patterns. *Research in Social Stratification and Mobility*, 53:1–15.
- Fee, C. E., Hadlock, C. J., and Pierce, J. R. (2006). Promotions in the Internal and External Labor Market: Evidence from Professional Football Coaching Careers. *The Journal of Business*, 79(2):821–850.
- Halgin, D. S., Borgatti, S. P., Mehra, A., and Soltis, S. (2020). Audience Perceptions of High-Status Ties and Network Advantage: The Market for Coaching Jobs in the NCAA (2000-2011). *Journal of Organizational Behavior*, 41(4):332–347.
- Harden, J. J. and Kropko, J. (2019). Simulating Duration Data for the Cox Model. *Political Science Research and Methods*, 7(4):921—928.
- Holmes, P. (2011). Win or go home: Why college football coaches get fired. *Journal of Sports Economics*, 12(2):157–178.

- Kopkin, N. A. (2014). You're Fired! The Impact of Race on the Firing of Black Head Coaches in Major College Football. *The Review of Black Political Economy*, 41(4):373–392.
- Kropko, J. and Harden, J. J. (2019). coxed: An R Package for Computing Duration-Based Quantities from the Cox Proportional Hazards Model. *The R Journal*, 11(2).
- Kropko, J. and Harden, J. J. (2020). Beyond the Hazard Ratio: Generating Expected Durations from the Cox Proportional Hazards Model. *British Journal of Political Science*, 50(1):303—320.
- MacCambridge, M., editor (2005). *ESPN College Football Encyclopedia*. ESPN Books, New York.
- Mixon, F. G. and Trevino, L. J. (2004). How Race Affects Dismissals of College Football Coaches. *Journal of Labor Research*, 25(4):645–656.
- Solow, B. L., Solow, J. L., and Walker, T. B. (2011). Moving on up: The Rooney Rule and Minority Hiring in the NFL. *Labour Economics*, 18(3):332–337.