

Career Implications of Having a Female-Friendly Supervisor: Evidence from the NCAA*

Steven Bednar[†] and Dora Gicheva[‡]

December 2013

Abstract

In this paper we study differences in observed performance and turnover patterns based on workers' gender and supervisors' propensity to hire and retain females. Using longitudinal data on athletic directors and head coaches of NCAA Division I women's teams we estimate athletic director fixed effects in the share of female head coaches, holding constant any time-invariant school-specific factors. While directors matter for the fraction of female head coaches, there is no statistical difference in the distribution of the estimated fixed effects by gender of the athletic director. Further, our findings indicate that the careers of male and female workers may progress differently depending on the propensity of their supervisors to work with and mentor females, and more focus should be placed on this propensity rather than on supervisor gender.

JEL Classification Codes: J44, J70, M50

*Draft. Please do not cite without permission.

[†]Department of Economics, Elon University, 2075 Campus Box, Elon, NC 27244, USA. E-mail: sbednar@elon.edu. Phone: +1 (336) 278-5935. Fax: +1 (336) 278-5952.

[‡]Department of Economics, University of North Carolina at Greensboro, PO Box 26170, Greensboro, NC 27402, USA. E-mail: d_gichev@uncg.edu. Phone: +1 (336) 334-4865. Fax: +1 (336) 334-5580.

1 Introduction

Greater representation of females at top levels of a firm has been found to have positive effects on female workers' career outcomes (e.g. Matsa and Miller 2011, Bell, Smith, Smith and Verner 2008). A similar relationship has been established in other settings, such as higher education (e.g. Neumark and Gardecki 1998, Bettinger and Long 2005, Hoffmann and Oreopoulos 2009). When workers and supervisors are similar along an easily observable dimension, such as gender, mentoring may be more effective (Athey, Avery and Zemsky 2000), or supervisors may be better at inferring employees' unobserved qualities (Cornell and Welch 1996).

Conversely, Blau and DeVaro (2007) find that promotions and promotion expectations for male and female workers in their data are unaffected by supervisor gender, and others, such as Rothstein (1997), have documented that females often earn higher wages when working under a male supervisor, but in the absence of detailed employer-level information it is not possible to rule out the explanation that supervisor gender is acting as a proxy for job type. Furthermore, the term "Queen bee syndrome" has been used to describe women who have reached the top ranks of their company, particularly in a male-dominated occupation, and intentionally hinder the career progress of females in lower levels of the job ladder (Staines, Tavriss and Jayaratne 1974).¹

Few papers look beyond simple observables such as gender or race, incorporating several easily identifiable characteristics into a composite measure of social proximity. Bandiera, Barankay and Rasul (2009) exploit manager and worker similarities across nationality, residential location and employment start date to show that social connections between lower-level workers and their manager matter less when the manager's goal is to maximize firm performance. Behncke, Frölich and Lechner (2010) find that

¹The applied psychology and human resource management literature has further addressed the relationship between gender and the amount and benefits of mentoring, but there is disagreement over the importance of the gender composition of the mentorship relationship (e.g. Dreher and Ash 1990, Dreher and Cox Jr. 1996, Ragins and Cotton 1999).

in Switzerland, the unemployed are more likely to find a job if they share gender, age, education and nationality with their caseworker.

In this paper we introduce the idea that visible characteristics may not be sufficient to define “type” in the mentoring relationship. We focus on one workplace setting whose features make it very suitable to study empirically the relationship between supervisor characteristics and the career progression of lower-level managers. We use a unique data set that identifies the athletic director, as well as the identity of the head coach and performance record for six women’s sports at all NCAA Division I programs during a twelve-year period. The structure of the data make it possible to track athletic directors and coaches across institutions and to compare outcomes at a set of workplaces where the internal hierarchy is homogeneous, mobility is fairly high, and a consistent quantifiable measure of performance exists.

We utilize the longitudinal nature of our data to estimate how much athletic directors influence the gender composition of head coaches above and beyond the university culture, an approach similar to Bertrand and Schoar (2003). It is important for our identification that athletic directors are responsible for personnel decisions and have some level of discretion into how much focus the school places on each sport’s success, for example through decisions about how much effort to put into fundraising to improve training facilities. We then use the inferred “female-friendliness” of the top-level administrators in our data as an explanatory variable in regressions of performance and retention of head coaches. We show that female coaches who are hired by a female-friendly athletic director see an increase in performance, as measured by team winning percentage, over the course of the coach-director match. The relationship is reversed for male coaches. Interestingly, these trends are not observed when the coach’s tenure exceeds that of the athletic director, consistent with a model in which supervisors make their hiring and mentoring decisions jointly. We also find that female-friendly athletic directors are more likely to retain a female coach conditional on observing poor performance than directors

whose type falls at the other end of the spectrum.

In our data, female athletic directors are no more likely to be female-friendly than their male counterparts. This indicates that studies focusing only on gender may fail to account for important components of the mentoring relationship. Overall, our findings indicate that the careers of male and female workers may progress differently depending on the inherent propensity, assumed to be exogenous, of their supervisors to work with and mentor females. While our study focuses on a very specific labor market setting, our findings should motivate researchers to place more focus on this propensity in addition to directly observable supervisor and worker characteristics such as gender.

Our paper adds to the broader strand of the literature that evaluates the links between gender and the performance of leaders in the context of the labor market. The findings of previous studies, which focus on the gender of top-level employees and not on the more complexly defined type that we propose, have been mixed. For example, Wolfers (2006) does not find evidence that large female-lead firms have different stock market performance than firms headed by males. Matsa and Miller (2013), on the other hand, show that gender quotas on corporate boards can affect firms' strategies and performance.

The rest of the paper is organized as follows: Section 2 lays out a theoretical framework that outlines how type-based mentoring can affect career outcomes in the athletic administrator and head coach setting, and section 3 provides more details about the data. The main empirical analysis follows in section 4, where we first establish the importance of athletic director fixed effects, followed by tests of the model's implications for coaches' career progression. Section 5 concludes.

2 Theoretical Framework

We consider a firm with two job levels, one of which is managerial; in the context of the empirical application of this paper, the two positions are athletic director and head

coach. Each coach c is characterized by an unobserved ability parameter a_c that determines the speed of learning-by-doing, and observed type s , where $s = \{\text{male, female}\}$. For simplicity we assume that the distribution of a does not differ by type s , but this assumption is not crucial for the predictions of the model. Athletic directors are characterized by their type $m = \{0, 1\}$, which determines their ability or willingness to mentor female coaches. Mentors of type $m = 0$ are more effective at improving the performance of male coaches over time, and female coaches acquire human capital more quickly under type $m = 1$ supervisors, conditional on the ability parameter a .

Observed performance (team winning percentage) every season is a noisy signal of the coach’s productivity, denoted by η :

$$w_t = \eta_{ct} + \nu_{ct}.^2 \tag{1}$$

Productivity increases over time with mentoring:

$$\eta_{ct} = \theta f(x_{ict}; a_c),$$

where x_{ict} is the length of time the coach has worked under the current athletic director i . Initial ability a and mentoring are additively separable in the model proposed by Athey et al. (2000). The human capital production function here is also similar to the Gibbons and Waldman (1999) setup, where the speed of learning-by-doing differs by worker type. When the parameter θ depends on s and m , our model corresponds to Athey et al.’s (2000) idea of “type-based” mentoring, and we can write (1) as

$$w_t = \theta(s, m) f(x_{ict}; a_c) + \nu_{ct}.$$

²The error term ν can be modified to include a time-invariant school-specific component.

Then the rate of learning depends on s and m :

$$\begin{aligned}\theta(s = \text{male}; m = 0) &> \theta(s = \text{male}; m = 1); \\ \theta(s = \text{female}; m = 0) &< \theta(s = \text{female}; m = 1); \\ \theta(s = \text{male}; m = 0) &> \theta(s = \text{female}; m = 0); \\ \theta(s = \text{male}; m = 1) &< \theta(s = \text{female}; m = 1).\end{aligned}$$

We do not assume that all values of θ are positive; performance may actually deteriorate under some mentor-mentee type combinations. In addition, we do not make any assumptions about the comparative magnitudes of $\theta(s = \text{male}; m_1)$ and $\theta(s = \text{female}; m_2)$, where $m_1 \neq m_2$.

Let a_{ct}^e denote the expected inherent ability of coach c after observing season t performance:

$$a_{ct}^e = E(a_c | a_{c,t-1}^e, w_t).$$

Then the expected value of retaining coach c for one more season is $V(a_{ct}^e, \theta(s, m))$, while the expected value of hiring a new coach is constant at V_0 . We observe a separation if $V(a_{ct}^e, \theta(s, m)) < V_0$.

The first testable prediction of the model is that coach performance improves more quickly over the course of the worker-supervisor match when female coaches are paired with athletic directors of type $m = 1$ and male coaches work under administrators with $m = 0$. If we assume additive separability between a and $\theta f(x_t)$, similar to Athey et al. (2000), we have that

$$\frac{\partial^2 w_t}{\partial x_t \partial \theta} = f'(x_t) > 0.$$

Second, the model predicts that the probability of separation is decreasing in w_t , since

$$\frac{\partial V(a_t^e, \theta)}{\partial w_t} = V_1 \cdot \frac{\partial a_t^e}{\partial w_t} > 0.$$

Third, the model predicts that the probability of separation increases less after a low realization of w_t when θ is larger:

$$\frac{\partial^2 V(a_t^e, \theta)}{\partial w_t \partial \theta} = V_{12} \cdot \frac{\partial a_t^e}{\partial w_t} < 0.$$

This follows from $V_{12} < 0$: the value of retaining a coach decreases less quickly when the expectation a_t^e is adjusted downward if θ is larger. In section 4.2 we first test the latter two predictions (section 4.2.2), and next we test the first prediction in section 4.2.1. It should be noted that this model does not imply taste-based discrimination. The underlying assumption is that certain types of supervisors are more effective at improving the performance of a given type of worker through mentoring, but we do not discuss the reasons for the variation in mentor types.

3 Data

3.1 Sample Construction

It is likely that the role that institutional culture plays for the share and success of female workers is non-negligible, whether we are looking at firms or athletic programs and head coaches. In order to identify the influence of athletic directors separately from school-level factors, we use a panel data set that tracks athletic directors across programs. We construct an athletic director - coach - university matched panel data set for Division I universities that spans the period from the 1993-1994 academic year to the 2010-2011 academic year. Other existing matched worker-firm data such as the Longitudinal EmployerHousehold Dynamics (LEHD) panel contain establishment-level firm characteristics but no information about a worker's supervisor and do not follow supervisors across establishments. Previous studies that link supervisor gender to worker outcomes, such as Giuliano, Levine and Leonard (2009) and Giuliano and Ransom (2013), are

focused on one employer and any firm-specific factors cannot be differenced out.

The gender of each athletic director and the tenure at their current school are identified through web searches. For six women's sports - basketball, field hockey, lacrosse, soccer, softball and volleyball - we record the gender of the head coach, current tenure, total number of years of experience as a head coach, and the team's winning percentage for each season. This information is provided on the NCAA website³. We restrict our analysis to women's sports because for the men's Division I sports for which the NCAA provides head coach and season-by-season performance data, the fraction of female coaches is either equal to or is very close to zero; more information is available in the *NCAA Member Institutions' Personnel Report* (2011).

These data allow us to track athletic directors and head coaches as they change schools. We record the starting month, when available,⁴ and year for each athletic director-university pair, including directors who were observed in the position in the 1993-94 year, which allows us to construct an accurate tenure measure of tenure. To improve the precision of the tenure variable, we take into account the exact months when administrators assumed and vacated their positions. Field hockey, soccer and volleyball are played during the fall season; basketball is a winter sport, and lacrosse and softball are spring sports. An athletic director is assigned the fall season if her start date is before September 1; basketball if the start date is before November 1 and the spring season of a given academic year if she started in January of that year or earlier. In the rare case where a coach separates from the school mid-season (less than 1.5 percent of all observations), we record the winning percentage and coach information for both subsets of the season.

Our identification approach uses only schools for which at least one of the athletic directors is observed at multiple programs, so we impose this restriction on the data.

³<http://web1.ncaa.org/stats/StatsSrv/careersearch>

⁴We were able to find information on the exact start date for three quarters of the director-school pairs in the sample. For all others, we assumed that the employment spell started during the summer.

For this subset of schools we use data for all years in which the school participated in Division I athletics with at least one team among the six sports we consider. The main sample consists of 13,086 team-level observations and 548 unique athletic directors, of which 486 occupy the same position for more than two years. Athletic directors who are observed in a program for fewer than two years are likely to be interim and as such may have comparatively less decision power. There are 137 administrators who are observed in the top position at multiple schools⁵. Using the subsample of directors observed at multiple schools allows us to separately identify athletic director fixed effects from time-varying institutional trends in regressions related to hiring practices and coach success, as fixed effects for athletic directors who do not switch schools are indistinguishable from period-specific school fixed effects.

3.2 Sample Description

While female athletic directors are underrepresented in Division I athletics, there has been a trend toward greater diversity among athletic administrators holding this top position. On the other hand, the fraction of female coaches in our data has decreased somewhat. Figure 1 shows how the percent of female coaches and athletic directors in the sample has changed from the 1993-1994 academic year to the 2010-2011 academic year. We exclude lacrosse and field hockey from the head coach gender ratio because as Table 1 illustrates, these sports are nearly always coached by women. Aside from temporary variations, the fraction of female athletic directors stayed fairly constant at about 7.5 percent for the schools competing in the Football Bowl Series and the Football Championship Series, and increased dramatically from 6.5 percent at the beginning of the sample period and 0 between 1995-96 and 1998-99 to nineteen percent in the 2010-11 school year for programs without a football team. The proportion of female coaches started out at 65 percent and decreased by about 10 percentage points for the top two

⁵Of these 137 movers, 116 are observed at two schools, 18 are observed at three schools and 3 are observed at four schools.

subdivisions and remained constant at around 46 percent for the lowest-ranked group of NCAA Division I programs. The trends observed in our sample follow those published by the NCAA in the *NCAA Member Institutions' Personnel Report (2011)*. Figure 1 emphasizes that our empirical strategy of identifying athletic director fixed effects only off of movers is preferable to estimating a fixed effect for all administrators in the data, given that programs that differ on observable and most likely on unobservable characteristics experienced different trends during the sample period. The gender distribution in our sample is similar to the labor market more generally defined if a parallel is drawn between the athletic director position and that of top corporate executives on the one hand, and head coaches and lower-level managers on the other. Matsa and Miller (2011) report that in their sample of publicly traded US companies, the share of females among the firm's top five executives increases from 3.2 percent to 6 percent between 1997 and 2009, and Bertrand and Hallock (2001) point out that 41.4 percent of firm managers (occupation codes between 3 and 22) in the Current Population Survey were female in the early to mid-1990s.

Table 1 gives the female fraction of coaches by sport and the number of teams in the sample that compete in each sport. The fraction of female coaches ranges from 0.35 for soccer to 0.72 for softball to over 0.9 for lacrosse and field hockey. Since the latter two sports are coached by females almost by default, it is possible that for those coaches the athletic director's inherent propensity to work with females has different implications compared to other sports in the data. In section 4.2 we perform the empirical analysis both with and without lacrosse and field hockey and find that disparities in the treatment of female coaches are, in fact, likely.

Table 2 displays summary statistics by head coach gender. The sample is further divided based on whether the current athletic director is one of the 137 individuals observed at multiple schools or not. Only athletic directors who spent more than two full years at a given school are included in the calculations. Athletic directors who

serve at multiple institutions tend to work with slightly more experienced coaches (8.5 compared to 7.8 years of experience) and achieve higher winning percentages, regardless of coach gender, but it is unclear how much of the difference is due to program-specific factors. Turnover rates and average coach tenure are fairly similar among the two groups of administrators, especially among male coaches. Male coaches have a slightly higher average winning percentage and are slightly less likely to leave their current position: 89 percent of male coaches and 87 percent of female coaches are observed at the same school during the following year. The average coach tenure in all subsamples is a little over 5.5 years. The average observed coach-athletic director spell is 2.6 years for both male and female coaches. Most of the analysis focusses on the subsamples in columns 1 and 2 of Table 2: the 510 male and 850 female coaches working under one of the 137 athletic administrators observed at multiple schools. This sample includes 188 different institutions.

4 Empirical Analysis

4.1 Importance of Athletic Director Heterogeneity

4.1.1 Fixed Effect Estimation

Our measure of an athletic director’s “female friendliness” is based on directors’ heterogeneity in the propensity to hire and retain female coaches. The gender distribution of coaches at a given program is likely to be influenced by school-specific factors that may vary over time. Our identification strategy is based on athletic directors who are observed at multiple schools. We estimate a model of the fraction of females coaching sports in our data at school j in academic year t :

$$pct_female_{ijt} = \gamma_j + \eta_t + \delta_i + \varepsilon_{ijt}. \quad (2)$$

We include school and year fixed effects to account for differences across institutions and over time in the gender composition of coaching staffs. The variable δ_i is an athletic director fixed effect estimated for administrators who are observed holding this position for more than two full seasons at multiple Division I programs during the sample period. This approach is similar to Bertrand and Schoar (2003) and ensures that these athletic directors have enough time to make personnel and other decisions that will affect the athletic department. Our approach is more conservative than an alternative specification in which a fixed effect is assigned to everyone who serves in the position of athletic director during part of the sample period. These fixed effects are empirically identified, but we are unable to distinguish the influence of non-movers from unrelated institution-specific time trends concurrent with their tenure. It must, however, be pointed out that athletic directors are clearly not randomly assigned to schools, so we are not necessarily identifying a causal effect. Our approach is to pick up systematic differences in the share of female head coaches across programs and over time that we can link to individual athletic directors.

The dependent variable in equation (2) measures the fraction of female head coaches for the sports in our sample excluding lacrosse and field hockey, which are coached by females over 90 percent of the time, as shown in Table 1. The resulting gender ratio aggregates the share of female coaches over the academic year, and athletic directors who left program j before the end of year t are assigned the same value they would have received if they had left at the end of the year. This may lead to measurement error in the dependent variable if the incoming athletic director replaced any coaches. We choose this approach because the alternative, adjusting the gender ratio by the director's departure date, leads to greater inaccuracy. Of the sports used to construct pct_female_{ijt} , only softball is played during the spring season, and as Table 1 shows, more females coach softball than basketball, soccer or volleyball. If softball is excluded from the calculation of the share of female coaches for administrators who left in January

for example, we will systematically underestimate *pct_female* for the year of departure.

The fixed effects in equation (2) are constructed for athletic directors who spent more than two full years at each school where they are observed, which excludes interim directors. Observations for which athletic director tenure is equal to zero are also excluded from the estimation. Since it is possible that coaching staff adjustments require time, this restriction allows us to use data that are less likely to reflect the previous director's choices. The predictive power of the model increases with the above sample restrictions.

We also estimate a version of equation (2) in which the dependent variable is replaced with a school-specific average of the winning percentages for sports in our sample, including lacrosse and field hockey, in a given academic year.⁶ This model is intended to test whether individual athletic directors have an impact on performance that can be separated from any institution-specific trends. The estimated winning fixed effects are a measure of one dimension of athletic director quality.

The estimation results for the two outcomes of interest are summarized in Table 3. Columns 1 and 2 show results for the female fraction of coaches, while columns 3 and 4 reflect the use of average winning percentage as a dependent variable. In each case the first set of results (columns 1 and 3) does not include any director fixed effects. We report the adjusted R^2 from each model, as well as the F-statistic for the joint significance test of all director indicator variables.

The results in Table 3 show evidence of non-negligible athletic director effects. The estimates also suggest that athletic director heterogeneity has a stronger impact on the gender composition of head coaches than for team performance for the sports we use in the estimation. In the first set of regressions, the adjusted R^2 increases from 0.49 to 0.56 when the athletic director fixed effects are included, and the F-statistic for the set of 137 binary variables is 4.52. In the average winning percentage regressions, the increase in the adjusted R^2 is less pronounced (0.43 to 0.47), and the F-statistic for the set of

⁶When two athletic directors were present during an academic year, two separate winning percentage averages are constructed, including only the sports assigned to each administrator (see Section 3).

137 fixed effects is lower (2.64). Both F-statistics allow us to reject the null hypothesis that all fixed effects equal zero at very high levels of significance.

4.1.2 Distribution of the Estimated Fixed Effects

Given the evidence in Table 3 that athletic director heterogeneity is important for the gender composition and the performance of head coaches, we proceed to examine in more detail the distribution of the estimated fixed effects. The distribution of the $\hat{\delta}$ estimates is shown in Figure 2. The most and the least female-friendly administrators tend to be male, and there are no discernible gender differences in the distributions. The null hypothesis that the $\hat{\delta}$ s for males and females are drawn from the same distribution cannot be rejected by a Kolmogorov-Smirnov test (p-value of 0.5).

Since in Section 4.2 the estimated coefficients are used as explanatory variables, they are adjusted to account for the additional estimation variance. The adjustment is based on Bayesian shrinkage (Morris 1983) and is similar to the approach taken by most studies of teacher or principal effectiveness (e.g. Jacob and Lefgren 2008, Leigh 2010, Branch, Hanushek and Rivkin 2012). Let $\hat{\delta}_i^{OLS}$ be the estimated fixed effect from the regression in (2), constructed to have a mean of zero. The true fixed effect is δ_i , and e_i is measurement error because OLS does not estimate the coefficients precisely. Under the classical measurement error assumptions,

$$\hat{\delta}_i^{OLS} = \delta_i + e_i$$

and

$$\text{Var}(\delta) = \text{Var}(\hat{\delta}^{OLS}) - \text{Var}(e).$$

When the fixed effects are included in second-stage regressions, the measurement error

attenuation bias is given by

$$\text{plim } \hat{\beta} = \beta \left(\frac{\text{Var}(\delta)}{\text{Var}(\hat{\delta}^{OLS})} \right).$$

To adjust the estimated fixed effects, we estimate $\text{Var}(\hat{\delta}^{OLS})$ by finding the sample variance of the estimated fixed effects (denoted $\hat{\sigma}_{OLS}^2$). $\text{Var}(e)$ is estimated as the sample average of the squared standard errors of the $\hat{\delta}^{OLS}$ parameters (denoted $\hat{\sigma}_e^2$). Then

$$\hat{\sigma}_\delta^2 = \hat{\sigma}_{OLS}^2 - \hat{\sigma}_e^2$$

and the adjusted fixed effects are given by

$$\hat{\delta}_i^A = \hat{\delta}_i^{OLS} \left(\frac{\hat{\sigma}_\delta^2}{\hat{\sigma}_\delta^2 + \hat{s}e_i^2} \right),$$

where $\hat{s}e_i$ is the standard error of the first-stage parameter estimate for the i^{th} athletic director. The resulting $\hat{\delta}^A$ parameters are standardized to have mean 0 and standard deviation of 1 in the sample of athletic directors observed at multiple programs. The purpose of this normalization is make the interpretation of the coefficients from second-stage models more straightforward when the estimated fixed effects are used as regressors. The standardized adjusted fixed effects are denoted by $\hat{\delta}_{st}^A$.

To further test for statistically significant gender differences in the fixed effect distribution, we regress the adjusted and standardized parameters $\hat{\delta}_{st}^A$ on a gender indicator and a constant term. Table 4 shows results from an OLS specification (column 1) and quantile regressions at the 25th, 50th and 75th percentiles (columns 2-4). We weight each observation by the inverse of the OLS variance of the fixed effect estimate ($\hat{s}e_i^2$) so that more weight is placed on more precisely estimated observations. The gender coefficients are not statistically significant in any of the specifications. The R^2 of the OLS regression in column 1 is low (0.0010). While the sample of female athletic direc-

tors observed at multiple schools for multiple years is very limited, these results suggest that supervisor gender and “female-friendliness,” defined as having a high value of δ , are unrelated in our data.

We also estimate a model of the form

$$(\hat{\delta}_{winning\%}^A)_i = \beta_0 + \beta_1(\hat{\delta}_{st}^A)_i + u_i, \quad (3)$$

whose purpose is to test for a relationship between the estimated adjusted winning percentage fixed effects and the fixed effects from the female share of coaches regression. Observations are again weighted by the inverse of $\hat{s}e_i^2$. In panel B of Table 4 we show the OLS results from the model in (3) (column 1), as well as quantile regressions at the 25th, 50th and 75th percentiles (columns 2-4) in order to test for differences at various points of the athletic director quality distribution. We do not see statistical difference by $\hat{\delta}_{st}^A$ in the quality of athletic directors at the 25th percentile of the distribution of $\hat{\delta}_{winning\%}^A$, but there is evidence that more female friendly administrators tend to perform worse at the mean, median and 75th percentile.

Finally, to provide some external validity of the measure of female-friendliness, we use data on team revenues and expenditures reported by Division I institutions in accordance with the Equity in Athletics Disclosure Act.⁷ We use data on total expenses and revenues by sport at each school for the period 2006-07 to 2010-11. As stated on the EADA survey form, the expenses category “includes appearance guarantees and options, athletically related student aid, contract services, equipment, fundraising activities, operating expenses, promotional activities, recruiting expenses, salaries and benefits, supplies, travel, and any other expenses attributable to intercollegiate athletic activities.” For sports that are played by both men and women at a given institution, we construct the fraction of revenues or expenses attributed to the women’s team. We regress these measures on the standardized adjusted athletic director fixed effects, director gender,

⁷The data are publicly available at <http://ope.ed.gov/athletics/>.

as well as year and team (school-by-sport) indicators. We again use the inverses of the fixed effect estimation variances as weights.

The results in Table 5 show that spending on women’s sports is higher when the athletic director is more “female-friendly”, while revenues are unaffected by the athletic director type that we define. We interpret these results as evidence that the supervisor type we introduce is not narrowly defined to matter only for the fraction of female coaches at an institution; it appears to be correlated with other female-friendly practices.

4.2 Implications for the Career Progression of Female Head Coaches

4.2.1 Performance and the Length of the Coach-Athletic Director Match

The first theoretic prediction outlined in section 2 is that coach performance improves more quickly over the course of the worker-supervisor match when female coaches are paired with female friendly athletic directors and male coaches work under administrators of the opposite type, who have a comparative advantage in mentoring male workers. We investigate whether an athletic director’s attitude towards working with female coaches is related to the variation in performance of coaches over time. Under a type-based mentoring hypothesis in the context of our study, performance would improve faster for female mentees (head coaches) when the mentor (athletic director) is of a female-friendly type. For male coaches, performance would improve slower when the athletic director is female-friendly. A mentoring relationship is expected to be stronger when a coach was hired by the current athletic director rather than by a previous administrator. We use team winning percentage to measure performance; this is an appropriate measure based on our finding in Table 3 that athletic directors have an influence on winning percentage when all time-invariant school-specific factors are differenced out.

We estimate the following regression for coach c in sport s at school j in year t

working under athletic director i :

$$\begin{aligned} \text{WinningPct}_{cijt} = & \gamma_1 \text{WinningPct}_{cij,t-1} + \gamma_2 (\hat{\delta}_{st}^A)_i + \gamma_3 (T_{cijt} \times (\hat{\delta}_{st}^A)_i) \\ & + \gamma_4 \text{Female AD}_i + \gamma_5 (T_{cijt} \times \text{Female AD}_i) + \mathbf{Z}_{ci} \mathbf{\Gamma} + \eta_t + \eta_{sj} + \nu_{cijt}. \end{aligned} \quad (4)$$

The variable T measures the number of years the coach has worked under the current athletic director at school j . When the director's tenure exceeds the tenure of the head coach, T equals the coach's tenure; otherwise T is set to equal the tenure of the athletic director. The variables contained in the vector \mathbf{Z} include quadratics in the athletic director's tenure and the coach's tenure and experience. We estimate (4) separately for coaches who were hired by the current athletic director (coach tenure is less than or equal to athletic director tenure) and coaches who were hired by a previous director (coach tenure exceeds athletic director tenure). The match length T does not enter the model uninteracted because it equals one of the tenure variables, depending on the sample restriction. This model includes team (school by sport) indicators; the results are similar when we include separate school and sport-specific binary variables and coach fixed effects. We also control for athletic director gender and its interaction with the length of the coach - director match. Since we compare a coach's current performance to her performance during the previous year, we exclude from the sample observations for which coach tenure equals zero. We again show results with and without lacrosse and field hockey for female coaches. Since the same restriction has negligible impact on the sample of male coaches, for them we only show results with lacrosse and field hockey. The inverse of the OLS variance of the fixed effects estimated from (2) is used as weight.

The estimation results for the regression in (4) are shown in Table 6. The lagged performance variable is positive and highly significant for female coaches and positive and significant but smaller for males. For female coaches who were hired under the current athletic director (columns 1 and 2), performance is initially worse but improves faster when the athletic director is female-friendly (negative coefficient on $\hat{\delta}_{st}^A$ that is significant

at the 10 percent level and positive and significant coefficient on the interaction of the estimated fixed effect with the length of the coach-director match). The relationship is stronger when lacrosse and field hockey coaches are excluded from the sample. On average, a female coach hired by an athletic director whose assigned fixed effect is one standard deviation above the mean sees a decline in winning percentage equal to 0.024 after one full year with the team, but performance for this coach-director pair improves faster than the baseline (director fixed effect equal zero) after 5 years of working together. The relationship is reversed for male coaches, whose performance improves faster when the athletic director they work for is assigned a fixed effect below the mean. Such relationship is not observed based on the gender of the athletic director.

Columns 4-6 of Table 6 suggest that any improvement in coach performance over time is unrelated to the gender or female friendliness of the athletic director when the coach was hired by a previous administrator. The coefficients on the interaction $T_{cijt} \times (\hat{\delta}_{st}^A)_i$ are not statistically or economically different from zero in all three specifications. Male coaches hired by a previous director tend to perform better when the current athletic director is female or less female friendly. Overall the results are consistent with the type-based mentoring framework in which female-friendly administrators hire and mentor coaches of their preferred type.

4.2.2 Performance and Mobility

The theoretical model in section 2 also implies that the value of retaining a coach increases with the observed winning percentage w_t , and therefore we expect to see an inverse relationship between w_t and the probability of separation. When athletic directors, in their role of mentors in the context of our study, differ in the value they add to mentees (coaches) over time, we will see a difference in the rate at which coaches are dismissed following a bad season. In particular, female friendly athletic directors will be less likely to dismiss a female coach based on poor current season performance than

directors with low values of $\hat{\delta}_{st}^A$. To explore how a supervisor’s “female-friendliness” impacts coach turnover in the data, we estimate a survival model in which failure is defined as a separation. We assume a Weibull proportional-hazards survival distribution: given current coach tenure t , the hazard function takes the form

$$\lambda(t; \mathbf{X}) = \exp(\mathbf{XB})\alpha t^{\alpha-1}.$$

The parameter α is estimated along with the vector of coefficients \mathbf{B} . One advantage of the Weibull specification is that it allows for duration dependence. The explanatory variables contained in \mathbf{X} include controls for athletic director gender and $\hat{\delta}_{st}^A$, the difference in team winning percentage between the current and previous season, and interactions between the change in winning percentage and supervisor gender and female friendliness. We also include athletic director tenure and tenure squared, a quadratic in coach experience, and year, school and sport dummy variables.⁸ Including the difference in winning percentage compared to the previous season instead of current performance adds a dynamic component to the model, allowing us to control better for team-specific trends, but does not change the results substantially. However, the estimation sample has to exclude observations from 1994 because we do not observe team performance in 1993. The model is estimated separately for male and female coaches. Since lacrosse and field hockey are almost exclusively coached by females, one can argue that gender will play a different role in these sports compared to sports in which administrators are more likely to have an impact on the coach’s gender. For this reason we show two sets of results, with and without lacrosse and field hockey. Since we use $\hat{\delta}_{st}^A$ as a regressor, the sample is limited to observations for which this parameter is estimated, namely teams currently lead by athletic directors who are observed at another school at some point

⁸We estimated alternative models in which the school and sport indicators were replaced with team-specific binary variables. The results for female coaches were similar, but the sample of male coaches is not sufficiently large for this model to be fully identified. Sample limitations also made problematic the estimation of a more flexible relationship between tenure and separation compared to the monotonic hazard implied by the Weibull distribution.

during the sample period.

Table 7 shows estimation results for all female coaches (column 1), female coaches excluding lacrosse and field hockey (column 2), all male coaches (column 3) and male coaches excluding lacrosse and field hockey (column 4). Due to the small number of male coaches in lacrosse and field hockey, columns 3 and 4 are almost identical. The table reports the estimates for the coefficients contained in the vector \mathbf{B} ; positive coefficients are associated with variables that increase the hazard, or decrease the probability of survival (remaining with the team). None of the coefficients on the female director indicator are statistically significant, but they are negative in columns 1 and 2 and positive in the male coach specifications, suggesting that female athletic directors may be less likely to dismiss female coaches and more likely to dismiss male coaches, regardless of performance. There is a similar trend associated with the female-friendly measure, and it is significant at the 5 percent level in columns 3 and 4. This observation is consistent with the fact that the variable was constructed to reflect the female share of coaches that a given administrator works with. The difference between winning percentage in seasons τ and $\tau - 1$ is negative and statistically significant in all four specifications, indicating that relatively poor performance in year τ makes it more likely that a new coach is observed in year $\tau + 1$, and this applies to both genders.

The main coefficients of interest in Table 7 are on the interaction between the winning percentage change from the previous season and the female-friendly measure. This coefficient is positive and significant at the five percent level in column 1 and larger and significant at the one percent level in column 2, when lacrosse and field hockey coaches are excluded from the sample of female coaches. This suggests that for female coaches performance matters less for turnover when the athletic director is female-friendly. Interestingly, this result is stronger in the sports in which we observe nontrivial variation in coach gender. The corresponding coefficients in the male coach regressions are negative but close to zero and not statistically significant. The interaction between current

season performance, relative to the previous season, and director gender is statistically indistinguishable from zero, but the point estimates are positive and large for male coaches. The estimate for the parameter α is statistically different from 1 and indicates positive duration dependence for all subsamples: the probability of coach turnover tends to increase with tenure.

The implications from the results in Table 6 in the previous section should be considered in combination with the findings from the duration model. Female friendly administrators are less likely to dismiss female coaches based on poor performance so holding everything else constant, under the implied selection mechanism we should observe a relative decline in the observed average performance of female coaches working under female-friendly supervisors. In fact, we see the opposite trend in Table 6, which leads us to believe that more than a pure taste-based discrimination sorting mechanism is needed to explain the observed trends. In this study we have proposed type-based mentoring as the missing link between Tables 6 and 7.

5 Conclusion

Previous research has established that individuals at the top levels of the firm have strong potential to affect overall performance. There is a growing body of literature that explores the mechanisms through which these leaders matter. In many cases high-level managers make personnel decisions or interact with mid-level managers, possibly in a mentorship role. Managers may be better able to detect the true ability of a subordinate if they share a characteristic, such as gender. Worker-manager type matches may make the role of mentorship more effective or alternatively there may be discrimination against members of the firm who do not share the same characteristic as the top manager.

We extend this literature by allowing unobservable characteristics to drive supervisors' level of support for female mid-level managers, which places additional burdens on the data in order to determine who is more female-friendly. We find evidence in our

sample of NCAA Division I athletic directors and head coaches of heterogeneity in the propensity to hire and retain females, which is identified separately from institution-specific constant and time-varying factors by focusing on administrators observed at multiple programs. We show that a supervisor’s female-friendliness is a much stronger predictor of coaches’ success and turnover, compared to the role that athletic director gender plays. The relationships we document in this specific labor market may not be immediately generalizable to broader labor market settings, but the conclusions we draw about the productivity of mentoring relationships are novel and should motivate researchers to explore further the role of unobservables.

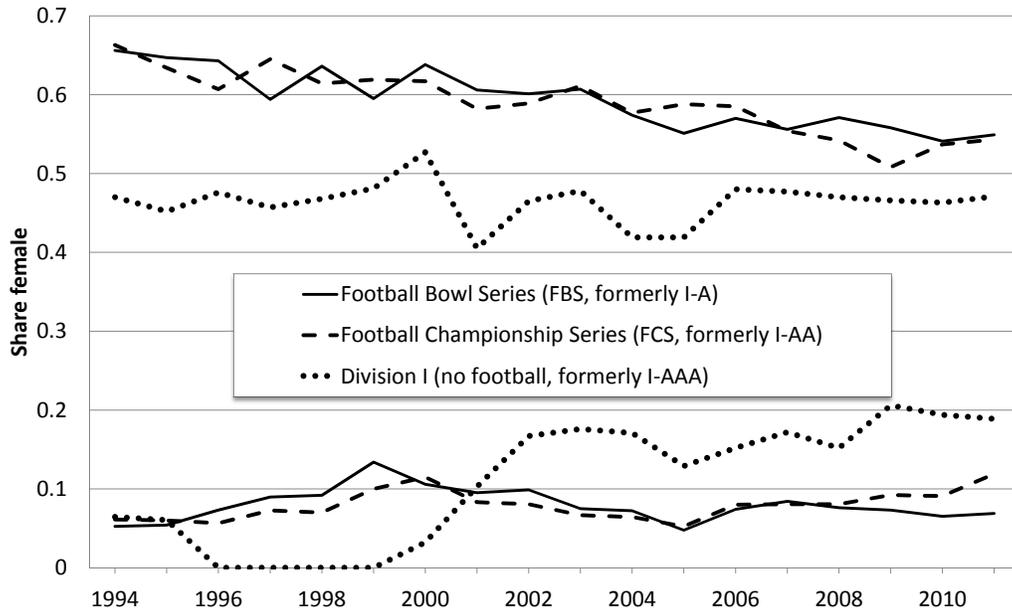
References

- Athey, Susan, Christopher Avery, and Peter Zemsky**, “Mentoring and Diversity,” *American Economic Review*, September 2000, *90* (4), 765–786.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul**, “Social Connections and Incentives in the Workplace: Evidence from Personnel Data,” *Econometrica*, 2009, *77* (4), 1047–1094.
- Behncke, Stefanie, Markus Frölich, and Michael Lechner**, “A Caseworker Like Me - Does The Similarity Between The Unemployed and Their Caseworkers Increase Job Placements?,” *Economic Journal*, December 2010, *120* (549), 1430–1459.
- Bell, Linda, Nina Smith, Valdemar Smith, and Mette Verner**, “Gender differences in promotion into top-management jobs,” Working Papers 08-21, University of Aarhus, Aarhus School of Business, Department of Economics January 2008.
- Bertrand, Marianne and Antoinette Schoar**, “Managing With Style: The Effect Of Managers On Firm Policies,” *The Quarterly Journal of Economics*, November 2003, *118* (4), 1169–1208.
- and **Kevin F. Hallock**, “The Gender gap in top corporate jobs,” *Industrial and Labor Relations Review*, October 2001, *55* (1), 3–21.
- Bettinger, Eric P. and Bridget Terry Long**, “Do Faculty Serve as Role Models? The Impact of Instructor Gender on Female Students,” *American Economic Review*, September 2005, *95* (2), 152–157.

- Blau, Francine D. and Jed DeVaro**, “New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires,” *Industrial Relations*, 2007, 46 (3), pp. 511–550.
- Branch, Gregory F., Eric A. Hanushek, and Steven G. Rivkin**, “Estimating the Effect of Leaders on Public Sector Productivity: The Case of School Principals,” Working Paper 17803, National Bureau of Economic Research February 2012.
- Cornell, Bradford and Ivo Welch**, “Culture, Information, and Screening Discrimination,” *Journal of Political Economy*, 1996, 104 (3), pp. 542–571.
- Dreher, George F. and Ronald A. Ash**, “A comparative study of mentoring among men and women in managerial, professional, and technical positions,” *Journal of Applied Psychology*, 1990, 75 (5), 539–546.
- and **Taylor H Cox Jr.**, “Race, gender, and opportunity: A study of compensation attainment and the establishment of mentoring relationships,” *Journal of Applied Psychology*, 1996, 81 (3), 297–308.
- Gibbons, Robert and Michael Waldman**, “A Theory Of Wage And Promotion Dynamics Inside Firms,” *The Quarterly Journal of Economics*, November 1999, 114 (4), 1321–1358.
- Giuliano, Laura and Michael R. Ransom**, “Manager Ethnicity and Employment Segregation,” *Industrial and Labor Relations Review*, 2013, 66 (2), 346–379.
- , **David I. Levine, and Jonathan Leonard**, “Manager Race and the Race of New Hires,” *Journal of Labor Economics*, October 2009, 27 (4), 589–631.
- Hoffmann, Florian and Philip Oreopoulos**, “A Professor Like Me: The Influence of Instructor Gender on College Achievement,” *Journal of Human Resources*, 2009, 44 (2).
- Jacob, Brian A. and Lars Lefgren**, “Can Principals Identify Effective Teachers? Evidence on Subjective Performance Evaluation in Education,” *Journal of Labor Economics*, 2008, 26, 101–136.
- Leigh, Andrew**, “Estimating teacher effectiveness from two-year changes in students’ test scores,” *Economics of Education Review*, June 2010, 29 (3), 480–488.
- Matsa, David A. and Amalia R. Miller**, “Chipping Away at the Glass Ceiling: Gender Spillovers in Corporate Leadership,” *American Economic Review*, May 2011, 101 (3), 635–39.
- and — , “A Female Style in Corporate Leadership? Evidence from Quotas,” *American Economic Journal: Applied Economics*, 2013, 5 (3), 136–169.

- Morris, Carl N.**, “Parametric Empirical Bayes Inference: Theory and Applications,” *Journal of the American Statistical Association*, 1983, 78 (381), pp. 47–55.
- NCAA Member Institutions’ Personnel Report**, Technical Report, National Collegiate Athletic Association, Indianapolis, Indiana February 2011.
- Neumark, David and Rosella Gardecki**, “Women Helping Women? Role Model and Mentoring Effects on Female Ph.D. Students in Economics,” *Journal of Human Resources*, 1998, 33 (1), 220–246.
- Ragins, Belle Rose and John L. Cotton**, “Mentor functions and outcomes: A comparison of men and women in formal and informal mentoring relationships,” *Journal of Applied Psychology*, August 1999, 84 (4), 529–550.
- Rothstein, Donna**, “Early career supervisor gender and the labor market outcomes of young workers,” in Francine Blau and Ronald Ehrenberg, eds., *Gender and family issues in the workplace*, New York, NY: Russell Sage Foundation, 1997, pp. 210–255.
- Staines, Graham, Carol Tavris, and Toby Jayaratne**, “The Queen Bee Syndrome,” *Psychology Today*, 1974, 7 (8), 63–66.
- Wolfers, Justin**, “Diagnosing discrimination: Stock returns and CEO gender,” *Journal of the European Economic Association*, 2006, 4 (2-3), 531–541.

Figure 1: Female Athletic Directors and Share of Female Coaches by Subdivision



Mean share of female head coaches (top three lines) and proportion of female athletic directors (bottom three lines) in the full sample of 3,135 school-year observations. The share of female coaches excludes lacrosse and field hockey.

Figure 2: Distribution of Athletic Director Fixed Effects by Gender

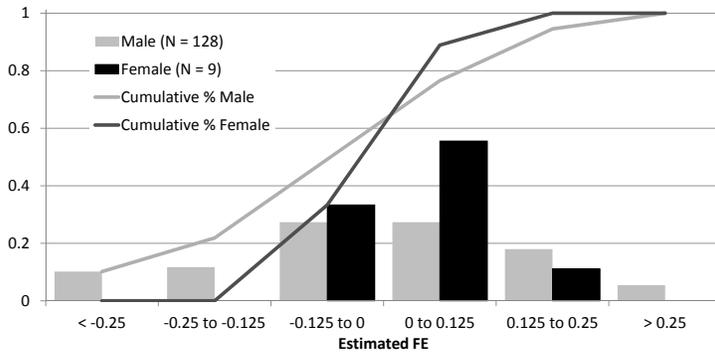


Table 1: Female Fraction of Coaches by Sport

Sport	N	Fraction female coaches
Basketball	3,386	0.667
Field Hockey	698	0.95
Lacrosse	702	0.912
Soccer	2,743	0.35
Softball	2,586	0.717
Volleyball	2,971	0.521
Total	13,086	0.606

Table 2: Descriptive Statistics by Head Coach Gender

	AD observed at multiple schools		AD observed at single school	
	Male Coach	Female Coach	Male Coach	Female Coach
Winning %	0.540 (0.199)	0.502 (0.192)	0.525 (0.195)	0.474 (0.194)
Last season with team	0.107	0.132	0.104	0.139
Tenure	5.710 (5.830)	5.800 (5.990)	5.720 (5.960)	5.560 (6.030)
Experience	8.570 (6.900)	8.460 (7.280)	7.930 (6.600)	7.780 (6.910)
Years working with AD	2.620 (2.750)	2.630 (2.740)	2.650 (2.920)	2.550 (2.730)
Number of obs.	2507	4098	2386	3426
Number of coaches	510	850	507	785
Number of schools		188		171
Number of ADs		137		349

Standard deviations in parentheses. Athletic directors who spent 2 years or less at one school are excluded.

Table 3: Importance of AD Fixed Effects, Share of Female Coaches and Average Winning Percentage

	% Female Coaches ^a		Average Winning %	
N	3135	3135	3135	3135
Adjusted R-sq.	0.49	0.56	0.43	0.47
F stat		4.52		2.64
Number of FEs		137		137
Year dummies	Yes	Yes	Yes	Yes
School dummies	Yes	Yes	Yes	Yes
AD fixed effects ^b	No	Yes	No	Yes

^aDoes not include lacrosse and field hockey

^bFixed effects only for athletic directors observed at multiple schools.

Table 4: Magnitude and Distribution of Estimated Fixed Effects

A. Mean and Quantile Regressions of $\hat{\delta}_{st}^A$				
	Mean	q25	q50	q75
Female AD	0.1423 (0.3826)	0.4205 (0.6409)	0.3136 (0.4986)	-0.1616 (0.5999)
Constant	-0.0226 (0.0886)	-0.5849*** (0.1484)	-0.0278 (0.1155)	0.6722*** (0.1389)
R^2	0.0010			
B. Mean and Quantile Regressions of Adjusted Winning % FEs				
	Mean	q25	q50	q75
$\hat{\delta}_{st}^A$	-0.0055* (0.0033)	-0.0052 (0.0048)	-0.0092** (0.0038)	-0.0111*** (0.0039)
Female AD	0.0198 (0.0147)	0.0181 (0.0212)	0.0009 (0.0171)	0.0183 (0.0173)
Constant	-0.0013 (0.0034)	-0.0168*** (0.0049)	0.0028 (0.0039)	0.0236*** (0.0040)
R^2	0.0321			

* p<0.10, ** p<0.05, *** p<0.01. $\hat{\delta}_{st}^A$ has mean 0 and standard deviation of 1 in the sample of athletic directors with estimated fixed effects. $N = 137$

Table 5: Female-Friendly Athletic Directors' Support for Women's Sports: EADA Data, 2006-07 to 2010-11

Dependent variable	Expenses	Revenues
$\widehat{\delta}_{st}^A$	0.0091** (0.0044)	-0.0178 (0.0131)
Female AD	0.0145 (0.0215)	0.0073 (0.0643)
N	2235	2233
R-squared	0.9327	0.8312

** Significant at the 5 percent level. The dependent variable in each regression is the fraction of expenses/revenues for a given sport that are attributed to the women's team. Only sports with a men's and women's team at an institution are included. All specifications include year and team fixed effects. The inverses of the fixed effect estimation variances are used as weights. There are 101 ADs and 107 schools included in the regressions.

Table 6: Performance and Length of Coach-Athletic Director Match by Timing of Coach Hiring

	Dep. variable: Winning % at t					
	Hired by current AD			Hired by previous AD		
	(1)	(2)	(3)	(4)	(5)	(6)
Winning % at $t - 1$	0.187*** (0.028)	0.224*** (0.032)	0.074* (0.040)	0.148*** (0.026)	0.149*** (0.028)	0.077** (0.032)
Female-friendly (stand.)	-0.028* (0.014)	-0.031* (0.016)	0.037 (0.027)	-0.015 (0.010)	-0.011 (0.011)	-0.035** (0.015)
Fem.-friendly x Yrs. match	0.005** (0.002)	0.007*** (0.003)	-0.007** (0.003)	0.000 (0.002)	0.001 (0.002)	-0.001 (0.002)
Female AD	0.017 (0.059)	0.085 (0.072)	-0.115 (0.130)	0.016 (0.038)	0.009 (0.040)	0.108** (0.047)
Female AD x Yrs. match	-0.013 (0.010)	-0.017 (0.011)	0.036 (0.029)	0.005 (0.006)	0.000 (0.008)	0.000 (0.009)
N	1538	1199	903	1897	1578	1222
R-squared	0.624	0.642	0.680	0.609	0.614	0.655
Coach gender	Female		Male	Female		Male
Lacrosse & field hockey	Yes	No	Yes	Yes	No	Yes

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The samples exclude observations with 0 years of coach tenure. The specifications also include controls for athletic director tenure and tenure squared, coach tenure, tenure squared, experience, and experience squared, as well as year and school by sport indicators.

Table 7: Performance and Mobility, Survival Analysis

	Female coaches		Male coaches	
	(1)	(2)	(3)	(4)
Female AD	-0.553 (0.438)	-0.882 (0.539)	0.816 (0.706)	0.768 (0.701)
Female-friendly (standardized)	-0.071 (0.124)	-0.165 (0.137)	0.505** (0.204)	0.510** (0.207)
Winning Diff.	-0.937*** (0.304)	-1.140*** (0.345)	-0.910** (0.432)	-0.965** (0.437)
Winning Diff. x Fem.-friendly (stand.)	0.617** (0.302)	0.945*** (0.355)	-0.114 (0.405)	-0.073 (0.408)
Winning Diff. x Female AD	0.259 (1.203)	-0.525 (1.409)	2.530 (1.738)	2.622 (1.739)
N	3432	2776	2125	2091
α	2.179	2.171	1.945	1.911
P-value ($\alpha = 1$)	0.000	0.000	0.000	0.000
log likelihood	-471.46	-355.15	-246.66	-242.28
Lacrosse & field hockey	Yes	No	Yes	No

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Results for Weibull survival model. The reported estimates are the coefficients on the explanatory variables. The specifications also include controls for athletic director tenure and tenure squared and a quadratic in coach experience, as well as indicators for year, school and sport.