



A Fair Game? Racial Bias and Repeated Interaction between NBA Coaches and Players

Administrative Science Quarterly
2017, Vol. 62(4)603–625
© The Author(s) 2017
Reprints and permissions:
sagepub.com/
journalsPermissions.nav
DOI: 10.1177/0001839217705375
journals.sagepub.com/home/asq


Letian Zhang¹

Abstract

There is strong evidence of racial bias in organizations but little understanding of how it changes with repeated interaction. This study proposes that repeated interaction has the potential to reduce racial bias, but its moderating effects may be limited to the treatment of individuals rather than of entire racial groups. Using data from 2,360 National Basketball Association (NBA) players and 163 coaches from 1955 to 2000, I find that players receive more playing time under coaches of the same race, even though there is no difference in their performance. This racial bias is greatly reduced, however, as the player and the coach spend more time on the same team, suggesting that repeated interaction minimizes coaches' biases toward their players. But it does not reduce coaches' racial biases in general. Even after years of coaching other-race players, coaches still exhibit the same levels of racial bias as they did upon first entering the league. These results suggest that repeated workplace interaction is effective in reducing racial bias toward individuals but not toward groups, making an important contribution to the literature on organizational inequality.

Keywords: inequality, racial bias, discrimination, intergroup contact, contact hypothesis, sports

Racial bias, the unequal treatment of individuals or groups on the basis of race, is an important driver of organizational inequality (Pager and Shepherd, 2008). Audit studies consistently find that white managers favor white candidates in hiring (Bertrand and Mullainathan, 2004; Pager, Western, and Bonikowski, 2009), and laboratory experiments show that subjects evaluate more favorably people of the same race (Bielby, 2000; Pager and Shepherd, 2008). Observational studies have found racial bias among sports referees (Price and Wolfers, 2010; Parsons et al., 2011; Pope, Price, and Wolfers, 2013; Pope and Pope, 2015), police officers (Donohue and Levitt, 2001; Antonovics and Knight, 2009), mortgage lenders (Yinger, 1996; Ladd, 1998), and physicians (Rubineau and Kang, 2012).

¹ Harvard University

Most of this evidence comes from single-interaction contexts, but in most organizations, managers and workers have frequent contact (Kalev, 2009). There are reasons to suspect that racial dynamics are different in initial and subsequent encounters (Copus, Ugelow, and Sohn, 2005; Landy, 2008). The contact hypothesis suggests that under the right conditions, repeated intergroup contact may improve racial attitudes (Allport, 1954; Amir, 1969). Many observers have therefore wondered if workplace interaction could reduce or even eliminate racial bias (Kalev, 2009), but few studies have explored bias in repeated interactions (Pager and Shepherd, 2008).

Repeated interaction could moderate racial bias in two ways. First, by establishing familiarity, it allows people to acquire individuating information, which should reduce their racial bias toward each other (Harrison, Price, and Bell, 1998; Harrison et al., 2002). Second, if people generalize their contact experiences into future encounters with other individuals, repeated contact could also change a decision maker's racial bias toward groups as a whole (Allport, 1954; Pager and Karafin, 2009). It is important to differentiate these processes because repeated interaction may affect one but not the other. The theory of subtyping, for example, suggests that people typically perceive familiar individuals as exceptions and do not associate them with their group identities (Weber and Crocker, 1983; Pager and Karafin, 2009), which would mean that repeated interaction would reduce bias toward only the familiar individuals, not their groups.

I consider these hypothesized processes by examining coaches' allocation of playing time in the National Basketball Association (NBA), which offers a unique opportunity to isolate the effect of race from potential individual confounders. Given that players and coaches frequently change teams, I can use player fixed-effects models to examine a player's treatment by coaches of different races, effectively controlling for any time-invariant individual traits (Halaby, 2004). In addition, a player's annual performance can be measured objectively, which affords robust controls of time-varying individual productivity (Kahn, 1991). These unique features help overcome the common critique that observational studies of racial bias cannot effectively control for individual quality (Farkas and Vicknair, 1996; Tomaskovic-Devey, Thomas, and Johnson, 2005; Pager and Shepherd, 2008; Charles and Guryan, 2011).

RACIAL BIAS AND REPEATED INTERACTION

There are two types of racial bias. First, decision makers may give more-favorable treatment to some racial groups over others because of their own racial preferences. Such bias is typically referred to as taste-based bias (Charles and Guryan, 2011), and a common example is people's tendency to like ingroups but dislike outgroups (Reskin, 2000, 2008). Second, decision makers may show bias toward individuals because they associate some racial groups with better qualities than other groups (Rissing and Castilla, 2014), a type of bias that can be further categorized into performance stereotyping and true statistical discrimination (Rubineau and Kang, 2012).¹ In performance stereotyping, sometimes referred to as error bias (England, 1992), decision makers use

¹ Some studies use the term "statistical discrimination" to broadly refer to both performance stereotyping and true statistical discrimination (e.g., Rissing and Castilla, 2014). For clarity, I use different terms for these two processes.

biased perceptions of a group to infer an individual's quality (Correll and Benard, 2006). In true statistical discrimination, decision makers draw from unbiased and objective group information to infer an individual's quality (Phelps, 1972; Aigner and Cain, 1977). But as Correll and Benard (2006) discussed, there has been limited empirical evidence of true statistical discrimination (Altonji and Pierret, 2001).

Scholars have documented racial bias across a range of contexts (e.g., Pager and Shepherd, 2008; Antonovics and Knight, 2009; Price and Wolfers, 2010; Parsons et al., 2011). Most existing literature, however, examines only one-shot encounters; racial dynamics in repeated interactions remain understudied (Pager and Shepherd, 2008). Although some lab studies suggest that intergroup contact improves people's racial attitudes (Pettigrew and Tropp, 2006), this proposition has received only limited attention in fieldwork (e.g., Jackman and Crane, 1986; Sigelman and Welch, 1993; Chakravarti, Menon, and Winship, 2014). In the few survey studies on this topic, researchers generally failed to find a contact effect on racial bias, leading some to question whether intergroup contact has any impact on racial dynamics in the real world (Jackman and Crane, 1986).

But it is possible that contact could influence people's attitudes toward individuals and groups differently (Brown and Turner, 1981; Hewstone and Brown, 1986), especially if we view contact as a two-stage process. In the first stage, which I call *individuation*, repeated contact with an individual would bring familiarity, which may influence a person's bias toward that individual. In the second stage, which I call *generalization*, people may generalize their contact experience into future encounters with others and reduce bias toward groups as a whole. These processes involve different mechanisms, and there are theoretical reasons to expect that repeated interactions could lead to individuation but not generalization.

Individuation

Racial bias includes both preference for ingroups and prejudice against outgroups (Reskin, 2000, 2008). The similarity-attraction paradigm suggests that people are attracted to and prefer to be with similar others because they anticipate that their own values, attitudes, and beliefs will be reinforced (Byrne, 1971). In initial encounters, people generally feel more comfort, trust, and obligation toward members of their own race (Reskin, 2000; McPherson, Smith-Lovin, and Cook, 2001). But as repeated interaction allows both parties to establish familiarity and exchange individuating information, they should become more likely to perceive similarity based on deep-level traits such as personality, attitude, and values and should place less emphasis on surface-level traits such as race and ethnicity (Harrison, Price, and Bell, 1998; Harrison et al., 2002). Repeated contact should therefore reduce same-race preference.

Outgroup prejudice often rests on negative stereotypes that portray other-race individuals as less competent and less likeable (Bielby, 2000; Reskin, 2000, 2008). Research shows that stereotypes are automatically activated in initial interactions and distort people's perceptions of each other (Bielby, 2000; Reskin, 2000). But as repeated contact allows people to acquire in-depth information about each other that disconfirms the racial stereotype, they should be less likely to perceive outgroup members in the interaction as incompetent and

undesirable (Reskin, 2000). A number of studies have found that people are more likely to apply stereotypes to strangers than to familiar outgroup members (Flynn, Chatman, and Spataro, 2001; Quinn, Mason, and Macrae, 2009). Through learning, repeated interaction should reduce outgroup prejudice.

But both effects have important scope conditions. They assume that two parties make an effort to learn about each other, which is more likely when the parties share a common goal and need to collaborate. In addition, there should be no direct competition between the parties, which could reduce trust and promote negative stereotypes (Allport, 1954; Reskin, 2000; Pettigrew and Tropp, 2006). Thus, in a collaborative setting in which there is positive engagement between individuals, I propose:

Hypothesis 1 (H1): As the duration of interaction with a given individual increases, a decision maker's racial bias in the treatment of that individual decreases.

Generalization

Traditionally, the contact hypothesis is concerned with how repeated contact reduces people's bias toward a group in general (Allport, 1954). For example, sociologists have long tried to understand the effect of intergroup interactions on whites' attitudes toward blacks (Jackman and Crane, 1986; Sigelman and Welch, 1993). Theoretically, it is possible that frequent contact with outgroup members would reduce bias not only toward those individuals but also toward other members of the same outgroup. As people repeatedly interact with outgroup members, they can gain familiarity with the outgroup as a whole. To the extent that familiarity generates liking, repeated interaction could increase comfort and trust toward the outgroup and thereby reduce ingroup preference in general. Repeated contact with outgroup members also lets decision makers recognize and correct inaccurate racial stereotypes (Weber and Crocker, 1983; Hewstone and Brown, 1986; Pager and Karafin, 2009). These processes suggest that people could generalize their contact experiences into future encounters with others, which leads to the following hypothesis:

Hypothesis 2 (H2): As the duration of interaction with a given individual increases, a decision maker's bias in the treatment of that individual's racial group as a whole decreases.

It is also possible, however, that repeated contact does *not* lead to a generalized effect on people's treatment of groups as a whole. The theory of subtyping suggests that people tend to maintain their existing beliefs. Faced with disconfirming evidence, people often subtype, interpreting observations that do not conform to their expectations as exceptions and not generalizing these cases into future encounters (Weber and Crocker, 1983; Hewstone and Brown, 1986). Applying this theory to intergroup contact, people may simply treat unpleasant ingroup members and favorable outgroup members as exceptions and otherwise maintain their ingroup attraction and outgroup prejudice (Pager and Karafin, 2009). Such subtyping is a likely outcome of repeated contact, because learning about each other weakens the salience of race. Contrary to the conventional view, subtyping predicts that repeated interaction reduces a decision maker's bias toward individuals but not their groups.

When we could expect subtyping rather than generalization may depend on the type of racial bias, although it is beyond the scope of this article to explore conditions for these two processes. In general, racial preference is more difficult to change than stereotypes of racial competence: performance stereotyping, based on biased perceptions that certain groups are more qualified than others, could change as people encounter disconfirming evidence (Brown and Turner, 1981; Hewstone and Brown, 1986), while preference for ingroups over outgroups is considered an intrinsic characteristic, persistently documented across a wide range of experimental conditions (Brewer, 1999; Reskin, 2000). Therefore it is possible that when bias comes from performance stereotyping, interaction is more likely to shift one's biased perceptions about a racial group's competence, resulting in generalization. When bias comes from racial preference, however, one may be more likely to treat individuals in the interaction as exceptions and retain one's racial preference, which would lead to subtyping.

A Strategic Setting: The National Basketball Association

It is generally difficult to analyze racial bias in repeated interactions. Audit studies, as well as lab experiments that can isolate the effect of race from other confounders, are mostly confined to one-time interactions between strangers (Pager and Shepherd, 2008; Pager, Western, and Bonikowski, 2009). In observational studies, researchers measure bias using residual racial differences after controlling for individual characteristics (Tomaskovic-Devey, Thomas, and Johnson, 2005). But in practice it is often difficult to objectively observe and quantify an individual's quality, and there may be important individual differences that are visible to the decision makers but not to researchers (Pager and Shepherd, 2008; Charles and Guryan, 2011). Claims of racial bias in observational studies have therefore frequently been questioned (Farkas and Vicknair, 1996; Charles and Guryan, 2011). To minimize the issue of unobserved individual heterogeneity, I use a unique setting: the National Basketball Association. I focus on coaches' allocation of playing time, a scarce and valuable resource for most players (Staw and Hoang, 1995; Camerer and Weber, 1999). With more playing time, a player gets more opportunity to display his skills and receives more attention from fans and the media, which can lead to more endorsement income and a higher-paying contract.

Studying racial bias in the NBA offers two important advantages (Kahn, 1991; Price and Wolfers, 2010). First, it allows a robust empirical examination of racial bias by using individual fixed-effects models to rule out unobserved time-invariant individual characteristics (Halaby, 2004). Because both players and coaches move frequently from team to team, an individual fixed-effects model observes how a player's opportunities (that is, his playing time) change as he moves between white and black coaches. Second, the context offers objective measures of individual performance, as the NBA tracks each player's game statistics on a per-minute basis. These statistics are quite detailed, involving different components of a player's game, such as the ability to score, rebound, and assist. There is a tradition in the organizational literature of using basketball data to provide objective measures of individual performance (Kahn, 1991; Harder, 1992; Staw and Hoang, 1995; Camerer and Weber, 1999), and these measures should provide robust controls for time-variant individual characteristics.

Compared with other labor markets, the NBA is a highly transparent context in which a coach's every action is scrutinized by fans and the media. Research has proposed that outside scrutiny can significantly reduce bias, as decision makers may fear the consequences of discrimination (Parsons et al., 2011; Pope, Price, and Wolfers, 2013). Thus this study offers a conservative test of labor market bias: if racial bias occurs in a highly visible setting such as the NBA, it may occur in other labor market settings as well.

METHODS

My analysis draws on NBA regular-season data from the 1955–1956 season to the 1999–2000 season. I did not include data after 2000 because of the large number of foreign players in the league at that time, which could complicate racial dynamics. I downloaded the data from basketball-reference.com, a site commonly used in sports analyses (Kubatko et al., 2007). For each player, I recorded his performance statistics in every season in which he appeared on an NBA roster, including points, rebounds, assists, fouls, and minutes played per game. I also recorded his team, playing position, and number of years in the league. For each team-year observation, I entered the name of the head coach, the team's win–loss record, and its conference.

I manually collected information on the race of every player and coach in the sample. I conducted coding using online photos of players and coaches with the assistance of a colleague. We coded separately, with an intercoder reliability rate of 99 percent. Each player was coded as either "black" or "white." I excluded the few cases in which my colleague and I coded differently, as well as cases in which a player's picture was not available. The final sample consists of 163 head coaches, 2,360 players, and 11,618 player-year observations.

When a player changed teams midway through a season, I used statistics from the team on which he started for that season. As a robustness check, I also ran models that included separate observations for each team a player spent time with during that season; the results are substantially similar. When a team changed its head coach during the season, I used the head coach who spent the most time with the team during that season. As a robustness check, excluding teams with a midseason coach change does not substantively change the results.


Another potential issue is the presence of foreign players: those who grew up overseas and did not attend college in the U.S. In the post-2000 period, about half of all white players are foreign, but I identified only 73 foreign players in my sample, with a total of 264 player-year observations. I ran models both with and without these foreign players in the sample, and the results are substantively similar.

Table 1 gives the descriptive statistics of the sample. Overall, there are more black than white players in the NBA but more white than black coaches. Despite the predominance of white coaches, black players tend to receive more playing time, and the statistics suggest that a large part of this gap is because, on average, the black players are better.² Figure 1 presents the distribution of black and white players and coaches over time. The percentage of black players in the league increased from only 13 percent in 1956 to 80

² A statistical comparison of the performances of black and white players is available upon request.

Table 1. Distribution of Player-year Characteristics

Variable	Black players	White players	All players
Number of black coaches	1,275	430	1,705
Number of white coaches	6,637	3,276	9,913
Number of coaches	7,912	3,706	11,618
Minutes per game	21.95	19.43	21.14
Points per 48 minutes	19.54	18.67	19.26
Rebounds per 48 minutes	8.86	9.67	9.12
Assists per 48 minutes	4.50	4.19	4.40
Years in the league	4.78	4.62	4.73
Years of collaboration with the coach	1.94	1.99	1.96
Coach experience (in years)	6.28	5.60	6.06

Figure 1. Racial compositions in the NBA.

percent in 2000. The pattern for coaches is slightly different. NBA coaches were exclusively white throughout the 1950s and the early 1960s. Black coaches emerged in the late 1960s but have remained a minority. Because racial numerical composition can have implications for racial dynamics, I included year fixed effects to control for overall temporal changes and ran models broken down by time periods.

Fixed-effects Models

I used linear panel models with player fixed effects to account for time-invariant individual heterogeneity. By treating individual players as dummy variables, the models rule out any time-invariant individual-level traits such as personality and physical characteristics that may influence a coach's decision in allotting playing

time (Halaby, 2004). I also included fixed effects on year to account for differences across time periods. Fixed effects cannot account for time-varying within-cluster correlations, so I also clustered standard errors at the player level.

In examining the individuation effect, I used linear panel models with coach–player dyadic fixed effects as an alternative. In this approach, I created a dummy for each coach–player relationship, which controls for time-invariant characteristics of each coach–player dyad. This allowed me to observe the changes within each coach–player relationship over time.

Player fixed-effects models examine changes in individual players' playing time as they move from a black coach to a white coach or vice versa. In the sample, 685 players made such a transition, accounting for 5,253 player-year observations (45 percent of the sample). These players tended to have longer tenures in the league. After I controlled for tenure, however, they did not systematically differ from others in race or playing position. For coach–player dyadic fixed-effects models, the focus is changes in playing time within each coach–player dyad. There are 7,015 coach–player dyads in the sample; 2,520 have multiple years of observations, resulting in 7,123 player-year observations (61 percent of the sample). Players in these dyads, too, did not show any systematic difference from others in race or playing position after controlling for league tenure.

Dependent Variable

I used a player's playing time per game as the dependent variable (Staw and Hoang, 1995; Camerer and Weber, 1999). Playing time gives players opportunities to showcase their skills and improve their market value. An additional analysis, available upon request, found that playing time is positively correlated with a player's likelihood of staying on the team, becoming an All-Star, and getting a raise. Playing time is a more-appropriate outcome variable than salary. The NBA has a complex wage system that includes minimum and maximum caps and takes into account a player's tenure and contract length. Playing time, in contrast, is a straightforward value determined by the head coach, who can substitute a player pretty much whenever and as often as he desires. It was calculated as a player's total minutes played in a season divided by the number of games in which he participated in the season, which does not include games in which he was unable to play due to injury.

Independent Variables

The main independent variable is whether a player has a same-race head coach. In alternative models, I broke down this variable into two scenarios: white coach–white player and black coach–black player. As table 1 shows, in 4,551 (39 percent) of the player-year observations, players and coaches are of the same race: 3,276 involve white coaches and players, and 1,275 involve black coaches and players. Among the 7,067 (61 percent) observations with other-race coaches, 6,637 have white coaches and black players and 430 have black coaches and white players.

In analyzing individuation, I measured the duration of contact between a head coach and a player in two ways. First, I used the total number of years a coach and a player spent on any team together, including, in a few cases, years

in which the coach was an assistant coach before becoming a head coach.³ This variable ranges from 1 to 16 years, and I logged it to account for its skewed distribution. Second, I included dummy variables to indicate whether a coach-player collaboration is in its first year, second year, or beyond. This detailed breakdown allowed me to observe when a reduction in racial bias is more likely. In the sample, 6,412 coach-player dyads (55 percent) are first-year collaborations, 2,613 (23 percent) are second-year, and 2,593 (22 percent) are third-year or above. I included these two measures of contact duration in separate models to avoid collinearity.

To examine generalization, I measured the amount of interaction a coach had with other-race players in his NBA coaching career using two measures. The more straightforward one is a coach's tenure in the league, assuming that a longer tenure affords more opportunities to interact with other-race players. I also constructed a more-precise measure: the number of other-race players a coach has coached. The two measures are strongly correlated at .94. I logged them to account for skewed distributions and included them in separate models to avoid collinearity.

Controls

To account for changes in a player's performance over time, I included his standard performance statistics—points, rebounds, and assists per game—in each season. Because these are strongly correlated with playing time, I standardized them on a per-minute basis (Staw and Hoang, 1995; Kubatko et al., 2007; Price and Wolfers, 2010). I also took into account teammates' influence using several measures. First, I controlled for a team's average win-loss record in the past three seasons to account for intrateam competitiveness. For newer teams, I set their past record at .5, the average value in the league. Presumably, players would face stiffer competition for playing time on a high-performing team than on a low-performing team. Second, I controlled for each team's overall racial composition using both Blau's index of heterogeneity for racial diversity and the percentage of teammates of the same race as the focal player. The Blau index measures a group's overall diversity, with a higher value indicating greater racial heterogeneity. I normalized this variable to the range between 0 and 1. Third, to measure seniority, I ranked players based on their tenure on a team and reverse coded the rank by subtracting it from 25 so that a higher rank indicates longer tenure.

I included a number of other variables that may influence playing time. First, I measured the number of years a player has been in the league to control for his NBA experience. Second, because fouls can restrict a player's playing time, I included a player's fouls per minute in the previous season. Third, I constructed a variable that measures whether a player is oriented toward defense or offense each year, as offensively oriented players may receive more playing time. I ranked each player's statistics in terms of points, assists, rebounds, and blocks each year in the league, sorted by position. Given that points and assists are typically associated with offense and rebounds and blocks with defense, I categorized players each season as either offensively or defensively oriented

³ In constructing this variable, I traced the data back to the 1946–1947 season, when the NBA was first established. Although many other variables are missing in the early years, I was able to obtain each player's name and his coach's name.

Table 2. Summary Statistics

Variable	Mean	S.D.	Min.	Max.
1. Minutes per game	21.14	10.75	1	48.5
2. Same-race coach	39.17%	NA	0	1
3. Years of collaboration (log)	0.47	0.59	0	2.77
4. Coach experience (log)	1.40	0.93	0	3.30
5. Other-race players coached (log)	2.85	1.65	0	5.44
6. Points per 48 minutes	19.26	6.90	0	144
7. Rebounds per 48 minutes	9.12	4.74	0	96
8. Assists per 48 minutes	4.40	3.03	0	48
9. Team's past win-loss record	0.50	0.12	0.19	0.83
10. Team diversity	0.74	0.22	0	1
11. Percentage same-race teammates	0.63	0.22	0.05	1
12. Seniority on the team	17.55	4.14	3	24
13. Years in the league	4.73	3.46	1	21
14. Fouls per 48 minutes	5.48	1.95	0	26.67
15. Defensive-oriented player	49.12%	NA	0	1
16. Coach's past win-loss record	0.52	0.09	0.22	0.83
17. Player is also the coach	0.26%	NA	0	1
18. Coach is also the GM	12.21%	NA	0	1
19. Western Conference	49.63%	NA	0	1

based on whether they have a higher average ranking in points and assists or in rebounds and blocks. Fourth, I controlled for coaches' win-loss record in the past three seasons. For newer coaches, I set their past record at .5. Fifth, I included two dummies to indicate whether a player is also the coach and whether a coach is also the general manager (GM), both of which are relatively rare occurrences that took place mostly in earlier periods. Finally, I included a control for whether a team is in the Western Conference. Table 2 shows summary statistics for the variables and controls I used. The correlations are in table A1 in the Online Appendix (<http://journals.sagepub.com/doi/suppl/10.1177/0001839217705375>).

RESULTS

My models show that a given player gets more playing time with a coach of his race than with a coach of a different race, even with no difference in performance. This racial bias is moderated by the duration of the coach's relationship with the player but not by the coach's tenure in the league.

Model 1 of table 3 shows that having a same-race coach increases a player's playing time per game by 34 seconds (.56 minutes) on average. Though that appears to be a small gap, it can accumulate into a significant disparity in the course of an entire season. As a comparison, in the same model a player receives only 16 seconds (.27 minutes) of playing time for every point he scores per 48 minutes, which is the length of a full game. A player playing for an other-race coach needs to score two more points every 48 minutes to receive the same playing time as a player playing for a same-race coach. To put this into perspective, two points per 48 minutes is almost one-third of a standard deviation (S.D. = 6.9). Thus having a same-race coach gives a player a significant edge. In model 4 of table 3, I broke down same-race coach into two

Table 3. Predicting Playing Time: Player Fixed Effects*

Variable	Model 1	Model 2	Model 3	Model 4	Model 5
Same-race coach	.563*	.935**	.808*		
	(.259)	(.343)	(.345)		
Same-race coach × Years of collaboration (log)		-.730**			
		(.252)			
Same-race coach × Coach experience (log)		-.042			
		(.193)			
Same-race coach × Second-year collaboration			-.663*		
			(.265)		
Same-race coach × Third-year collaboration and above			-.873*		
			(.359)		
Same-race coach × Other-race players coached (log)			-.016		
			(.105)		
Black coach–Black player				.498	.841
				(.292)	(.431)
White coach–White player				.749	1.126
				(.559)	(.607)
Black coach–Black player × Years of collaboration (log)					-.492
					(.431)
White coach–White player × Years of collaboration (log)					-.830**
					(.281)
Black coach–Black player × Coach experience (log)					-.101
					(.285)
White coach–White player × Coach experience (log)					-.009
					(.223)
Points per 48 minutes	.269**	.269**	.270**	.269**	.269**
	(.040)	(.040)	(.040)	(.040)	(.040)
Rebounds per 48 minutes	-.002	.000	-.001	-.001	.001
	(.047)	(.047)	(.048)	(.047)	(.047)
Assists per 48 minutes	.421**	.418**	.417**	.421**	.417**
	(.076)	(.076)	(.076)	(.076)	(.076)
Team's past win–loss record	-4.779**	-4.792**	-4.673**	-4.781**	-4.783**
	(.931)	(.933)	(.932)	(.932)	(.934)
Team diversity	.278	.248	.224	.265	.225
	(.487)	(.489)	(.489)	(.489)	(.491)
Percentage same-race teammates	-.755	-.914	-.864	-.758	-.934
	(.803)	(.809)	(.807)	(.804)	(.810)
Seniority on the team	.260**	.259**	.268**	.260**	.259**
	(.023)	(.023)	(.023)	(.023)	(.023)
Years in the league	.213	.217	.243	.211	.213
	(.212)	(.213)	(.213)	(.212)	(.213)
Fouls per 48 minutes	-.381**	-.381**	-.382**	-.382**	-.381**
	(.061)	(.061)	(.061)	(.061)	(.061)
Defensive-oriented player	-2.900**	-2.915**	-2.921**	-2.899**	-2.916**
	(.226)	(.227)	(.227)	(.226)	(.227)
Coach's past win–loss record	-2.606*	-2.664*	-2.402*	-2.609*	-2.662*
	(1.119)	(1.121)	(1.116)	(1.119)	(1.122)
Player is also the coach	-4.126**	-4.204**	-4.217**	-4.115**	-4.250**
	(1.281)	(1.308)	(1.322)	(1.279)	(1.320)
Coach is also the GM	-.014	-.010	-.020	-.015	-.009
	(.272)	(.274)	(.273)	(.272)	(.274)
Western Conference	-.019	-.015	-.015	-.020	-.019
	(.212)	(.212)	(.212)	(.212)	(.212)
Years of collaboration (log)	1.185**	1.450**		1.186**	1.454**
	(.161)	(.187)		(.160)	(.188)

(continued)

Table 3. (continued)

Variable	Model 1	Model 2	Model 3	Model 4	Model 5
Second-year collaboration			.946** (.181)		
Third-year collaboration and above			1.842** (.257)		
Coach experience (log)	-.465** (.093)	-.442** (.113)		-.468** (.094)	-.446** (.113)
Other-race players coached (log)			-.256** (.062)		
Constant	37.364** (5.344)	37.402** (5.351)	37.547** (5.354)	37.275** (5.338)	37.255** (5.344)
R-squared	.226	.227	.226	.226	.227

• $p < .05$; ** $p < .01$.

* Standard errors are in parentheses. In all models, the number of players is 2,360, and the number of observations is 11,618. Player and year fixed effects are included in all models.

types: white coach–white player and black coach–black player. They show comparable coefficients, suggesting that this same-race effect applies to both black and white players.

At the same time, same-race bias decreases as coaches and players spend more years together. Model 2 of table 3 shows that the length of the collaboration significantly reduces the effect of having a same-race coach. In model 3, same-race bias is 48 seconds (.81 minutes) in first-year coach–player collaborations but only 9 seconds (.15 minutes) in second-year collaborations and essentially disappears afterwards. In model 5, I broke down the same-race coach category into white coach–player pairs and black coach–player pairs and found that collaboration reduces the same-race effect in both cases, although its moderating effect on black players is not statistically significant. Overall these findings support hypothesis 1 that repeated interaction reduces a coach's biases toward the players he coaches.

But table 3 shows that a coach's overall experience does not significantly moderate same-race bias. In models 2 and 5, the interactions between same-race coach and the coach's experience have small coefficients that are not statistically significant, showing no evidence that a more-experienced coach exercises less racial bias than a less-experienced coach. Similarly, in model 3, I interacted same-race coach with the total number of other-race players a coach has coached in his career; this interaction, too, has a small coefficient that is not statistically significant. Together, these results do not support hypothesis 2, that repeated interaction with a large number of other-race players reduces a coach's bias in general. Models 2, 3, and 5 include multiple interaction terms; to account for possible collinearity, I also ran these interaction effects separately, and the results are substantially similar.

Selection Issues

In examining individuation, selection could be an issue. Though a team's general manager is responsible for signing players and coaches, his decision may

be influenced by a player's relationship with the coach, and therefore the length of each coach–player relationship may be endogenously determined. This selection could lead to what appears to be a reduction in racial bias, as the less-favorable coach–player relationships are more likely to dissolve over time. To account for this scenario, I conducted two additional analyses.

First, table 4 uses coach–player dyadic fixed-effects models to more robustly examine temporal changes in each coach–player dyad. By fixing each dyad, the effect of same-race coach is omitted. Consistent with table 3, model 2 of table 4 shows that the length of coach–player collaboration significantly reduces the effect of having a same-race coach. In model 3, the same-race effect is .24 minutes less in a second-year collaboration than in a first-year collaboration and is .94 minutes less in the third year and beyond. By showing a significant reduction in racial bias in each coach–player relationship, these models give more-robust evidence of the individuation effect predicted by hypothesis 1.

Second, I ran an analysis on first-year head coach–player collaborations to examine whether a coach who previously collaborated with a player as an assistant coach has less same-race bias in the first year of working with that player as a head coach. A team typically has several assistant coaches, and they have limited influence on a player's playing time. Because it is highly unlikely that the general manager would make roster decisions based on a player's relationship with one of them, collaboration between an assistant coach and a player is less subject to selection concerns. At the same time, selection should also have less influence on the formation of first-year collaborations between a head coach and a player, as the general manager has not yet observed how well they work together. This model, therefore, is more robust to selection issues. Table 5 shows that among first-year coach–player relationships, coaches who have previously interacted with players as assistant coaches show significantly less same-race bias than those who have not. This is also consistent with the individuation effect proposed in hypothesis 1.

Additional Analyses

I ran three additional analyses to further substantiate the main findings. First, to entertain the possibility that players are more motivated to perform under a same-race coach—and assuming that increased motivation leads to better performance—I ran a set of player fixed-effects models that examine whether having a same-race coach leads to better performance. I found that having a same-race coach has no significant effect on a player's points, rebounds, and assists per 48 minutes and defensive orientation, suggesting no significant differences in a player's behavior under coaches of different races (see table A2 in the Online Appendix). Of course, it is possible that players may exhibit different off-court behaviors that I could not fully capture, but at a minimum, coaches appear to favor same-race players in ways not warranted by their playing.

Second, I considered the possibility that coaches' exposure to other-race members prior to becoming NBA coaches influences their same-race biases. According to hypothesis 2, if people generalize their contact experiences into future encounters with others, then having frequent other-race contact should reduce one's racial bias in general. Coaches may therefore exhibit different levels of bias depending on their outgroup contact prior to becoming NBA coaches. I considered two scenarios. In the first, some coaches were previously

Table 4. Predicting Playing Time: Coach–Player Fixed Effects*

Variable	Model 1	Model 2	Model 3
Same-race coach × Years of collaboration (log)		-.910** (.346)	
Same-race coach × Second-year collaboration			-.244 (.277)
Same-race coach × Third-year collaboration and above			-.944* (.457)
Points per 48 minutes	.307** (.051)	.304** (.051)	.305** (.051)
Rebounds per 48 minutes	.001 (.087)	.003 (.087)	.001 (.087)
Assists per 48 minutes	.408** (.102)	.406** (.102)	.406** (.103)
Team's past win–loss record	2.795 (1.806)	2.618 (1.802)	2.672 (1.798)
Team diversity	.451 (.690)	.328 (.692)	.342 (.690)
Percentage same-race teammates	.639 (1.128)	.038 (1.146)	.256 (1.137)
Seniority on the team	.274** (.040)	.271** (.040)	.279** (.040)
Years in the league	1.545* (.696)	1.460* (.695)	1.558* (.686)
Fouls per 48 minutes	-.019 (.074)	-.015 (.073)	-.013 (.074)
Defensive-oriented player	-1.799** (.280)	-1.813** (.279)	-1.810** (.279)
Coach's past win–loss record	4.741* (1.981)	4.680* (1.969)	4.879* (1.983)
Player is also the coach	-10.316** (1.324)	-11.536** (1.398)	-11.561** (1.169)
Coach is also the GM	.332 (.668)	.361 (.673)	.378 (.672)
Western Conference	.404 (.919)	.348 (.913)	.359 (.912)
Years of collaboration (log)	.537 (.383)	.923* (.403)	
Second-year collaboration			.322 (.227)
Third-year collaboration and above			.920* (.390)
Constant	61.543** (16.034)	59.617** (15.994)	60.162** (16.040)
R-squared	.138	.140	.139

* $p < .05$; ** $p < .01$.

* Standard errors are in parentheses. In all models, the number of coach–player dyads is 7,015, and the number of observations is 11,618. Coach–player dyad and year fixed effects are included in all models.

NBA players and may therefore already have had extensive interactions with other-race players. In the second, coaches came from various racial environments. Presumably, coaches who grew up in more-diverse neighborhoods had more opportunities to interact with other-race individuals. To take these two possibilities into account, I created two additional variables. The first indicates

Table 5. Predicting Playing Time: First-Year Head Coach–Player Collaborations Only*

Variable	Model 1	Model 2
Same-race coach	.677* (.271)	.815** (.278)
Same-race coach × Years of assistant coach–player collaboration (log)		−1.126* (.562)
Points per 48 minutes	.242** (.052)	.242** (.052)
Rebounds per 48 minutes	.009 (.051)	.008 (.051)
Assists per 48 minutes	.510** (.087)	.509** (.087)
Team’s past win–loss record	−6.186** (1.040)	−6.195** (1.040)
Team diversity	.049 (.661)	.048 (.662)
Percentage same-race teammates	−.551 (1.028)	−.570 (1.030)
Seniority on the team	.251** (.027)	.251** (.027)
Years in the league	.289 (.220)	.293 (.220)
Fouls per 48 minutes	−.147* (.067)	−.146* (.067)
Defensive-oriented player	−3.084** (.294)	−3.087** (.294)
Coach’s past win–loss record	−6.620** (1.423)	−6.601** (1.424)
Player is also the coach	−.898 (2.635)	−.940 (2.578)
Coach is also the GM	−.387 (.325)	−.395 (.325)
Western Conference	−.090 (.231)	−.085 (.232)
Years of assistant coach–player collaboration (log)	.685* (.328)	1.113** (.389)
Constant	32.444** (5.846)	32.454** (5.847)
R-squared	.183	.184

* $p < .05$; ** $p < .01$.

* Standard errors are in parentheses. In both models, the number of players is 2,342, and the number of observations is 6,961. Player and year fixed effects are included in all models.

whether a coach had been an NBA player, and the second measures the racial heterogeneity of a coach’s hometown. I identified each coach’s hometown from basetkall-reference.com and collected data from the U.S. Census to calculate each city’s average racial diversity in the past 50 years, based on Blau’s index of heterogeneity. I was able to identify hometown racial diversity for 100 of the 163 coaches in my sample. I interacted these two variables with same-race coaches to see if their backgrounds influenced their same-race biases and found these interaction terms to be statistically insignificant (see table A3 in

Table 6. Predicting Playing Time: The Moderating Role of Competition*

Variable	Model 1 Playoffs	Model 2 Regular season	Model 3 Regular season
Same-race coach	.173 (.731)	1.270** (.370)	4.462** (1.609)
Same-race coach × Percent close games			-6.193* (3.105)
Points per 48 minutes	.780** (.071)	.171** (.045)	.173** (.045)
Rebounds per 48 minutes	.269 (.157)	-.027 (.051)	-.029 (.051)
Assists per 48 minutes	.715** (.173)	.379** (.114)	.379** (.114)
Team's past win-loss record	-.814 (2.214)	-4.198** (1.275)	-4.080** (1.277)
Team diversity	-2.141 (1.200)	.780 (.667)	.852 (.663)
Percentage same-race teammates	.348 (2.097)	1.723 (1.343)	1.758 (1.343)
Seniority on the team	.340** (.057)	.232** (.035)	.232** (.035)
Years in the league	.620 (.780)	-.032 (.404)	-.050 (.405)
Fouls per 48 minutes	-.887** (.163)	-.293** (.090)	-.293** (.089)
Defensive-oriented player	-.280 (.576)	-1.972** (.367)	-1.945** (.367)
Coach's past win-loss record	-7.235** (2.240)	-4.175** (1.442)	-4.218** (1.440)
Coach is also the GM	1.507 (.863)	.467 (.399)	.458 (.398)
Western Conference	-1.096 (.575)	.050 (.326)	.048 (.326)
Years of collaboration (log)	.771* (.342)	.719** (.254)	.713** (.253)
Percent close games		-1.415 (1.507)	.353 (1.837)
Constant	12.193** (4.319)	19.897** (2.222)	18.864** (2.326)
Observations	3,403	4,157	4,157
R-squared	.224	.155	.156
Number of players	905	1,011	1,011

* $p < .05$; ** $p < .01$.

* Standard errors are in parentheses. Player and year fixed effects are included in all models.

the Online Appendix). These results are consistent with my other findings and do not support hypothesis 2.

Finally, I broke down the sample across three time periods to observe how coaches' racial bias changed over time. I found that having a same-race coach had a stronger effect on playing times in the 1990s than in the earlier periods (see table A4 in the Online Appendix). The higher level of bias in the 1990s is likely due to the increasingly competitive talent pool. As the NBA became more competitive, the margin between players narrowed. In supplementary analyses

I found that in the 1990s, variance among players' points per game significantly declined, while an individual player experienced significantly more variance in his playing time. With more players of a similar quality to choose from, coaches are more likely to resort to their personal preferences, which could explain the increasing bias in the 1990s.

Underlying Mechanisms

I also examined whether the observed bias comes from coaches' racial preference, performance stereotyping, and/or true statistical discrimination. I expected racial preference to be more likely in the NBA, as both performance stereotyping and statistical discrimination tend to occur when there is limited information on an individual's quality (Altonji and Pierret, 2001; Rubineau and Kang, 2012). In the NBA, every player's performance can be objectively observed, so it seems less likely that coaches would strongly associate a player's competence with his race.

I examined the effect of competition on this bias, because competitive pressure should discourage coaches from indulging in their personal preferences and push them to use the best players available, thus reducing preference-induced bias (Becker, 1957). In the NBA, competition tends to be most intense in playoff games and close games. I used players' average playoff minutes as outcomes in model 1 of table 6, which shows that having a same-race coach has a small, statistically insignificant effect. I then included a variable counting the number of closely contested games a team played each regular season. I considered a game to be close if the final difference in score is less than 10 points, and I counted the percentage of close games each team had in each regular season since 1990–1991. On average, 52 percent of the games are close. As model 2 of table 6 shows, the percentage of closely contested games each season significantly reduces a coach's same-race bias. These two analyses show that coaches' same-race biases decrease in competitive settings, a result more consistent with preference-based bias.

DISCUSSION

Does repeated interaction reduce racial bias? Using NBA player fixed effects and objective performance metrics, I found that coaches give same-race players more playing time than other-race players, but this racial bias significantly declines when coaches and players spend more time on the same team, suggesting that repeated interaction minimizes coaches' bias toward their players. At the same time, coaches' interaction with other-race players has little impact on their same-race bias in general. A coach who has spent years working with other-race players exhibits as strong a racial bias as he did upon first entering the league.

These findings contribute to theories on intergroup contact. Drawing on the contact hypothesis, researchers generally assume that interaction can help people develop favorable views of outgroups by facilitating knowledge and familiarity (Allport, 1954; Jackman and Crane, 1986; Sigelman and Welch, 1993). My findings show, however, that collaborative interaction could have many of the theorized effects—weakening ingroup attraction and outgroup

prejudice—while racial bias persists, because the contact changes a person's attitude only toward an individual, not toward that individual's group as a whole. This finding supports the proposition that two separate processes underlie the contact effect. Individuation changes decision makers' attraction and attitude toward individuals with whom they have interacted, and generalization supposedly allows them to generalize their past interactions to future ones with other individuals. Drawing on the theory of subtyping, however, I suggest that generalization may be difficult to achieve when bias comes from racial preference. In these cases, while repeated interaction can still reduce bias, its moderating effects would be limited to one's treatment of individuals rather than of entire racial groups.

This distinction between individuals and groups helps explain the effect of racial integration in organizations. The conventional view assumes that repeated contact changes people's racial attitudes in general and therefore predicts that bias will be minimized in integrated workplaces (Reskin, 2000; Kalev, 2009; Chakravarti, Menon, and Winship, 2014). But research has found that even in highly integrated organizations, racial bias is still prevalent (Pager and Shepherd, 2008). My study suggests that in integrated settings, people could still maintain significant racial bias toward those with whom they are less familiar. This effect is clearly illustrated in the NBA context, a highly integrated work setting in which coaches and players of different races collaborate daily. Racial bias is prevalent even in such a setting largely because contact does not change coaches' racial attitudes in general, so their treatment of players with whom they have not previously collaborated is still strongly biased.

This study also shows that the contact effect takes time. Although the contact hypothesis does not specify a timeline, the literature often implicitly assumes that contact can reduce bias fairly quickly. Most lab studies are based on contacts that take place only briefly, usually involving days or weeks of interaction at most (Amir, 1969; Pettigrew and Tropp, 2006). I found instead that reducing bias is a slow process. Even in the highly interactive NBA context, it takes a long time to fully change coaches' racial preferences toward their players. After years of interaction, additional contact still has an impact. Dyadic fixed-effects models show that there is a significant reduction in coaches' bias toward their players even going from the second year of collaboration to the third year and beyond.

This study contributes to the literature on bias and discrimination in two ways. First, it has implications for reducing bias in organizations. Some researchers have suggested that team-based work could help reduce racial bias by providing more intergroup contact (Kalev, 2009), but there has been limited empirical evidence supporting this claim, and many remain skeptical (Kalev, 2009). This study, in which working together on a team reduces coaches' racial bias toward their players, offers empirical support.

Second, I depart from existing work on organizational inequality by focusing on the allocation of opportunities. The literature tends to analyze distinct stages of employment—such as hiring, promotion, and layoffs—as mechanisms of inequality but seldom examines opportunity allocation as one (Petersen and Saporta, 2004; Pager and Shepherd, 2008). Opportunities to work on important and visible tasks can influence career success. Briscoe and Kellogg (2011) found that employees do significantly better in their careers when they are

initially assigned to powerful supervisors, largely because it brings access to reputation-building project opportunities. My study suggests the importance of opportunity allocation as a source of racial inequality. NBA players experience racial disparities in their opportunities to perform in each game, which can affect their careers and long-term earnings. Similar dynamics could take place in other organizations. For instance, managers could assign more-favorable portfolios and more-important clients to same-race workers. Compared with bias in hiring and promotion, bias in opportunity allocation could be more difficult to prevent because there is less oversight (Petersen and Saporta, 2004). Even if such an opportunity gap is not significant on any given day, it can accumulate over time into significant racial disparities in an organization (Castilla, 2008).

Some limitations in this study provide opportunities for future research. First, it is not clear whether the observed bias comes from a decision maker's same-race attraction, other-race prejudice, or both. It is often assumed that these two processes are reciprocally related, but recent work suggests that they could be distinct (Brewer, 1999). Unfortunately, the NBA context is not well suited to tease them apart. Playing time is a zero-sum game, so coaches who favor same-race players are penalizing other-race players, if only inadvertently. Second, it is also difficult with this dataset to distinguish whether the bias comes from white coaches, black coaches, or both. In within-player fixed effects, a player's playing time under a same-race coach serves as a baseline for his playing time under an other-race coach and vice versa. Asking which coaches are exercising bias requires establishing a no-bias baseline, which is impossible with a player fixed-effects model. Using a between-individual model would raise the issue of unobserved individual heterogeneity, as there could be important unobserved differences between players (Kahn, 1991; Halaby, 2004; Price and Wolfers, 2010). Thus although the NBA context offers a robust empirical setting to identify racial bias, it also leaves some questions unanswered. Perhaps future studies could use experiments to clarify the underlying processes in more detail.

As organizations become more racially diverse, recognizing racial inequality in the workplace is increasingly important. Although racial bias has been well studied in one-time encounters, there is limited understanding of it in contexts that involve frequent contact, such as workplaces. By showing that contact reduces bias toward individuals, this study highlights the importance of further research to explore the role of repeated interactions in mitigating racial inequality in the workplace.

Acknowledgments

I gratefully acknowledge Bart Bonikowski, Frank Dobbin, Alexandra Killewald, and Elena Obukhova, who have given me tremendous support throughout this project. I also thank Julie Battilana, Jason Beckfield, Daniel Brown, Nancy DiTomaso, Roberto Fernandez, Simo Goshev, Alexandra Kalev, Kathleen McGinn, Devah Pager, Lakshmi Ramarajan, Amy Tsang, Daniel Wu, Yanhua Zhou, and members of Dobbin Research Group for their feedback on various drafts. Finally, I thank Associate Editor John Wagner and three anonymous reviewers for their comments and Linda Johanson and Joan Friedman for copyediting.

REFERENCES

- Aigner D. J., and G. Cain**
1977 "Statistical theory of discrimination in labor markets." *Industrial and Labor Relations Review*, 30: 175–187.
- Allport, G. W.**
1954 *The Nature of Prejudice*. Cambridge, MA: Addison-Wesley.
- Altonji, J. G., and C. R. Pierret**
2001 "Employer learning and statistical discrimination." *Quarterly Journal of Economics*, 116: 313–350.
- Amir, Y.**
1969 "Contact hypothesis in ethnic relations." *Psychological Bulletin*, 71: 319–342.
- Antonovics, K., and B. G. Knight**
2009 "A new look at racial profiling: Evidence from the Boston Police Department." *Review of Economics and Statistics*, 91: 163–177.
- Becker, G. S.**
1957 *The Economics of Discrimination*. Chicago: University of Chicago Press.
- Bertrand M., and S. Mullainathan**
2004 "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *American Economic Review*, 94: 991–1013.
- Bielby, W. T.**
2000 "Minimizing workplace gender and racial bias." *Contemporary Sociology*, 20: 120–129.
- Brewer, M. B.**
1999 "The psychology of prejudice: Ingroup love and outgroup hate?" *Journal of Social Issues*, 55: 429–444.
- Briscoe, F., and K. C. Kellogg**
2011 "The initial assignment effect: Local employer practices and positive career outcomes for work–family program users." *American Sociological Review*, 76: 291–319.
- Brown, R. J., and J. C. Turner**
1981 "Interpersonal and intergroup behaviour." In J. C. Turner and H. Giles (eds.), *Intergroup Behaviour*: 33–65. Oxford: Blackwell.
- Byrne, D. E.**
1971 *The Attraction Paradigm*, vol. 11. Cambridge, MA: Academic Press.
- Camerer, C. F., and R. A. Weber**
1999 "The econometrics and behavioral economics of escalation of commitment: A re-examination of Staw and Hoang's NBA data." *Journal of Economic Behavior and Organization*, 39: 59–82.
- Castilla, E. J.**
2008 "Gender, race, and meritocracy in organizational careers." *American Journal of Sociology*, 113: 1479–1526.
- Chakravarti, A., T. Menon, and C. Winship**
2014 "Contact and group structure: A natural experiment of interracial college roommate groups." *Organization Science*, 25: 1216–1233.
- Charles, K. K., and J. Guryan**
2011 "Studying discrimination: Fundamental challenges and recent progress." *Annual Review of Economics*, 3: 479–511.
- Copus, D., R. S. Ugelow, and J. Sohn**
2005 "A lawyer's view." In F. J. Landy (ed.), *Employment Discrimination Litigation: Behavioral, Quantitative, and Legal Perspectives*: 450–502. San Francisco: Jossey-Bass.
- Correll, S. J., and S. Benard**
2006 "Biased estimators? Comparing status and statistical theories of gender discrimination." In S. R. Thye and E. J. Lawler (eds.), *Social Psychology of the Workplace*, vol. 23: *Advances in Group Processes*: 89–116. Bingley, UK: Emerald Group.

- Donohue, J. J., III, and S. D. Levitt**
2001 "The impact of race on policing and arrests." *Journal of Law and Economics*, 44: 367–394.
- England, P.**
1992 *Comparable Worth: Theories and Evidence*. London: Transaction.
- Farkas, G., and K. Vicknair**
1996 "Appropriate tests of racial wage discrimination require controls for cognitive skill: Comment on Cancio, Evans, and Maume." *American Sociological Review*, 61: 557–560.
- Flynn, F. J., J. A. Chatman, and S. E. Spataro**
2001 "Getting to know you: The influence of personality on impressions and performance of demographically different people in organizations." *Administrative Science Quarterly*, 46: 414–442.
- Halaby, C. N.**
2004 "Panel models in sociological research: Theory into practice." *Annual Review of Sociology*, 30: 507–544.
- Harder, J. W.**
1992 "Play for pay: Effects of inequity in a pay-for-performance context." *Administrative Science Quarterly*, 37: 321–335.
- Harrison, D. A., K. H. Price, and M. P. Bell**
1998 "Beyond relational demography: Time and the effects of surface- and deep-level diversity on work group cohesion." *Academy of Management Journal*, 41: 96–107.
- Harrison, D. A., K. H. Price, J. H. Gavin, and A. T. Florey**
2002 "Time, teams, and task performance: Changing effects of surface- and deep-level diversity on group functioning." *Academy of Management Journal*, 45: 1029–1045.
- Hewstone, M. E., and R. E. Brown**
1986 *Contact and Conflict in Intergroup Encounters*. Oxford: Basil Blackwell.
- Jackman, M. R., and M. Crane**
1986 "'Some of my best friends are black . . .': Interracial friendship and whites' racial attitudes." *Public Opinion Quarterly*, 50: 459–486.
- Kahn, L. M.**
1991 "Discrimination in professional sports: A survey of the literature." *Industrial and Labor Relations Review*, 44: 395–418.
- Kalev, A.**
2009 "Cracking the glass cages? Restructuring and ascriptive inequality at work." *American Journal of Sociology*, 114: 1591–1643.
- Kubatko, J., D. Oliver, K. Pelton, and D. T. Rosenbaum**
2007 "A starting point for analyzing basketball statistics." *Journal of Quantitative Analysis in Sports*, 3 (3): 1–22.
- Ladd, H. F.**
1998 "Evidence on discrimination in mortgage lending." *Journal of Economic Perspectives*, 12: 41–62.
- Landy, F. J.**
2008 "Stereotypes, bias, and personnel decisions: Strange and stranger." *Industrial and Organizational Psychology*, 1: 379–392.
- McPherson, M., L. Smith-Lovin, and J. M. Cook**
2001 "Birds of a feather: Homophily in social networks." *Annual Review of Sociology*, 27: 415–444.
- Pager, D., and D. Karafin**
2009 "Bayesian bigot? Statistical discrimination, stereotypes, and employer decision making." *Annals of the American Academy of Political and Social Science*, 621 (1): 70–93.

Pager, D., and H. Shepherd

2008 "The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets." *Annual Review of Sociology*, 34: 181–209.

Pager, D., B. Western, and B. Bonikowski

2009 "Discrimination in a low-wage labor market: A field experiment." *American Sociological Review*, 74: 777–799.

Parsons, C. A., J. Sulaeman, M. C. Yates, and D. S. Hamermesh

2011 "Strike three: Discrimination, incentives, and evaluation." *American Economic Review*, 101: 1410–1435.

Petersen, T., and I. Saporta

2004 "The opportunity structure for discrimination." *American Journal of Sociology*, 109: 852–901.

Pettigrew, T. F., and L. R. Tropp

2006 "A meta-analytic test of intergroup contact theory." *Journal of Personality and Social Psychology*, 90: 751–783.

Phelps, E. S.

1972 "The statistical theory of racism and sexism." *American Economic Review*, 62: 659–661.

Pope, B. R., and N. G. Pope

2015 "Own-nationality bias: Evidence from UEFA Champions League football referees." *Economic Inquiry*, 53: 1292–1304.

Pope, D. G., J. Price, and J. Wolfers

2013 "Awareness reduces racial bias." NBER Working Paper No. 19765.

Price, J., and J. Wolfers

2010 "Racial discrimination among NBA referees." *Quarterly Journal of Economics*, 125: 1859–1887.

Quinn, K. A., M. F. Mason, and C. N. Macrae

2009 "Familiarity and person construal: Individuating knowledge moderates the automaticity of category activation." *European Journal of Social Psychology*, 39: 852–861.

Reskin, B. F.

2000 "The proximate causes of employment discrimination." *Contemporary Sociology*, 29: 319–328.

Reskin, B. F.

2008 "Including mechanisms in our models of ascriptive inequality." In L. B. Nielsen and R. L. Nelson (eds.), *Handbook of Employment Discrimination Research*: 75–97. New York: Springer.

Rissing, B. A., and E. J. Castilla

2014 "House of green cards: Statistical or preference-based inequality in the employment of foreign nationals." *American Sociological Review*, 79: 1226–1255.

Rubineau, B., and Y. Kang

2012 "Bias in white: A longitudinal natural experiment measuring changes in discrimination." *Management Science*, 58: 660–677.

Sigelman, L., and S. Welch

1993 "The contact hypothesis revisited: Black–white interaction and positive racial attitudes." *Social Forces*, 71: 781–795.

Staw, B. M., and H. Hoang

1995 "Sunk costs in the NBA: Why draft order affects playing time and survival in professional basketball." *Administrative Science Quarterly*, 50: 474–494.

Tomaskovic-Devey, D., M. Thomas, and K. Johnson

2005 "Race and the accumulation of human capital across the career: A theoretical model and fixed-effects application." *American Journal of Sociology*, 111: 58–89.

Weber, R., and J. Crocker

1983 "Cognitive processes in the revision of stereotypic beliefs." *Journal of Personality and Social Psychology*, 45: 961–977.

Yinger, J.

1996 "Discrimination in mortgage lending: A literature review." In J. M. Goering (ed.), *Mortgage Lending, Racial Discrimination and Federal Policy*: 29–74. Washington, DC: Urban Institute Press.

Author's Biography

Letian Zhang is a doctoral student in the Department of Sociology, Harvard University, 33 Kirkland Street, Cambridge, MA 02138 (e-mail: letian.lt.zhang@gmail.com). His current research examines how an organization's diversity influences both (1) its allocation of opportunities and (2) its evaluation by stakeholders. He received his B.S. in mathematics from Stanford University.