Journal of Economic Behavior and Organization 175 (2020) 98–113

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

‘Overattention’ to first-hand experience in hiring decisions: Evidence from professional basketball☆

Michael Dalton[a], Peter Landry[b,*]

[a] Bureau of Labor Statistics, United States
[b] University of Toronto, Mississauga, and Rotman School of Management, University of Toronto, St. George, Canada

ARTICLE INFO

Article history:
Received 4 September 2019
Revised 11 April 2020
Accepted 13 April 2020
Available online 17 May 2020

Keywords:
Inattention
Behavioral labor economics
Hiring decisions

ABSTRACT

We provide evidence from a real-world, high-stakes, and empirically-advantageous labor market — the market for NBA basketball players — that employers' hiring decisions rely too heavily on first-hand experiences with job candidates. Specifically, we find that employers are biased in favor of acquiring players with better-than-usual performances when the employer's team was playing or preparing to play the player's original team, with performance information receiving approximately 1.8 times more weight in hiring decisions if it is conveyed through such first-hand experiences. These effects are not predicted by leading behavioral learning theories used to explain similar effects observed in other domains. Instead, our findings point to overattention as a key mechanism through which first-hand experience biases can arise.

1. Introduction

Employers typically learn about job candidates from a variety of information sources. Some, such as a college transcript, a reference letter, or a performance evaluation from another job, are second-hand in nature. Others, such as an interview, an interaction at a networking event, or an opportunity to observe on-the-job behavior, may instead be described as first-hand experiences. A familiar hiring dilemma can arise when second-hand information and first-hand experiences point to different conclusions — one candidate may look better "on paper" while another seems better "in person." In this paper, we investigate whether employers optimally balance these information sources in real-world hiring decisions.

Past research suggests that hiring managers tend to be over-reliant on personal interviews (Highhouse, 2008; Dana et al., 2013) and under-reliant on the (inherently second-hand) input of "algorithmic" hiring aids (Kuncel et al., 2013; Hoffman et al., 2018). Such findings naturally raise the possibility that employers may systematically overweight first-hand experience relative to second-hand information in hiring decisions. However, these tendencies could instead reflect motives of hiring managers that go beyond hiring the best candidate for the job. For instance, personal interviews could allow hiring

☆ We are grateful for helpful comments from many who attended presentations of this work (at the University of Toronto and at the Society of Labor Economists) and from the review team. We also benefit greatly from discussions with Attila Ambrus, Peter Arcidiacono, Joshua Gans, Daniel Goetz, Avi Goldfarb, Mitch Hoffman, Tanjim Hossain, Nico Lacetera, Matt Mitchell, Matt Osborne, Avni Shah, and especially Patrick Coate. Ian Weaver provided excellent research assistance. The views and opinions expressed in this paper are our own and do not reflect the views of the BLS or the Federal Government. The authors declare that they have no relevant interests (financial or otherwise) to disclose. The appendix to this paper is available online at Supplementary material.

* Corresponding author.
E-mail addresses: dalton.michael@bls.gov (M. Dalton), peter.landry@rotman.utoronto.ca (P. Landry).

https://doi.org/10.1016/j.jebo.2020.04.015
0167-2681/

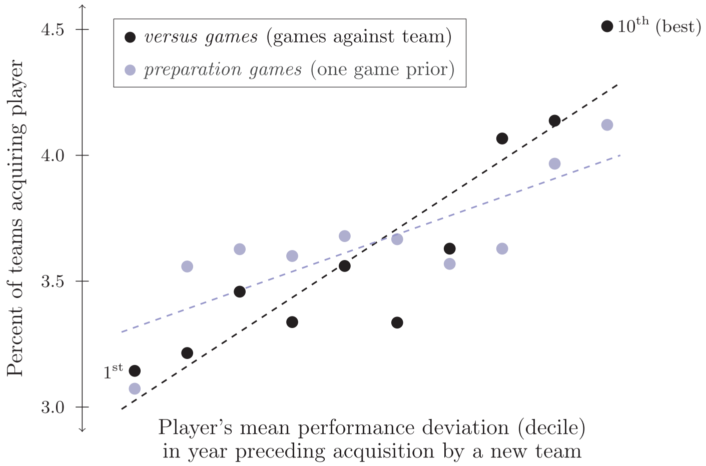

Fig. 1. Teams are more likely to acquire players who played relatively well (compared to their mean performance) when the team was playing or preparing for the player's original team.

managers to tilt the scales in favor of job candidates with whom they would personally like to affiliate. Consistent with this idea, interviewers' evaluations are known to be swayed by a candidate's physical attractiveness (Ruffle and Shtudiner, 2014), relationship status (Rivera, 2017), and whether they enjoy similar leisure activities (Rivera, 2012). Moreover, a hiring manager could be motivated to over-ride the recommendations of algorithmic hiring aids considering deference to such technologies may send a signal that their own input is unimportant (Goldsmith, 2000; Highhouse, 2008).

It is also not clear whether an over-reliance on personal interviews and under-reliance on algorithmic hiring aids would generalize to other forms of first-hand experience and second-hand information. Even if we put aside any doubts regarding hiring managers' personal motives, these findings could merely reflect a bias in favor of tradition. After all, interviews have historically been (and continue to be) the most widely-used tool for evaluating job candidates (Buckley et al., 2000), while algorithmic hiring aids are a relatively recent innovation that (as with other new and unfamiliar technologies) decision-makers may naturally be reluctant to adopt (see Highhouse, 2008).

To investigate whether employers do in fact systematically and suboptimally overweight first-hand experience relative to second-hand information in real-world hiring decisions, the National Basketball Association (NBA) offers a near-ideal testing ground. NBA teams can learn about the quality of another team's player — a potential future employee — based on his current job performance, which is readily quantifiable and freely accessible through widely-published performance statistics (e.g., points per game). While these statistics efficiently convey the player's overall performance in all games played during a given period of time, some of these performances — such as those occurring in games against the team evaluating the player — are also experienced first-hand. To the analyst, this feature is especially helpful for isolating the impact of first-hand experiences on employers' hiring decisions.

As a first look at the effect of first-hand experience, consider the black data points in Fig. 1. The horizontal axis indicates, by decile, the difference between a player's mean performance against a team and his mean performance against all teams in the year prior to the player joining a new team.[1] The vertical axis indicates the percentage of cases for which the player joins that particular team. As apparent from the upward-sloped trendline, teams are more likely to acquire players who played well against them in the past year.

As another source of first-hand experience, it is common practice for NBA teams (that is, certain personnel involved in player evaluation and acquisition decisions) to prepare for an upcoming game by observing their opponent's *previous* game.[2] The gray trendline in Fig. 1 shows that teams are also more likely to acquire players who perform better in games immediately preceding games against their own team.

The results of our controlled multinomial regressions reaffirm these relationships, as we find that a team's likelihood of acquiring a player independently increases with the player's (relative) performance in "versus games" and in "prepara-

[1] The performance measure we use, known as the 'efficiency' metric, is a widely-used (including by the NBA for official statkeeping purposes) and freely-available composite of the statistics reported in a standard NBA box score. Specifically, it is the sum of the good (points, rebounds, assists, steals, blocks) minus the bad (missed shots, turnovers) statistics. The ranges in players' mean performance deviations (according to this measure) for each decile in Fig. 1 are provided in Appendix Table A1.

[2] Namely, NBA teams routinely send a 'scout' to an opponent's previous game, who then communicates opposing players' strengths and weaknesses to coaches and managers, and work with team-employed video specialists to create video clips of key plays for coaches to review and show their own players. Waiting until an opponent's immediately-preceding game before scouting them allows teams to acquire the most up-to-date information on an opponent's plays and play-calling signals. As one long-time NBA coach explains, "we prefer to scout an opponent as close to the date of our game as possible including in its last game prior to playing us" (Dunleavy and Eyen, 2009). This practice is reinforced by seating arrangements at NBA arenas, which generally hold seats specifically designated for scouts representing each of the two participating team's next opponents. According to one veteran scout: "Each arena gives you two seats I can't really do my job unless I'm the next opponent, because only the next opponent gets the seats" (Agnessm, 2016).

tion games" (i.e. games when the team is playing, and games when the team is preparing to play the player's original team, respectively). We argue that these relationships, which we refer to as the *versus* and *preparation effects*, reflect a bias of *over-attention* to first-hand experience whereby teams are overly influenced by performances in versus and preparation games. That is, while players' performances in all games are presumably conveyed (and objectively weighted) through second-hand summary statistics, teams naturally allocate extra attention to games they experience first-hand — while evidently assigning meaning to the redundant performance signals provided in such games. As for the magnitude of this alleged bias, our estimates indicate that teams overweight performances in games they experience first-hand roughly by a factor of 1.8, causing 2.3 percent of players to be matched to the "wrong" team.

Certainly the versus effect (as described thus far) lends itself to many alternate explanations. For example, a player may perform better against teams that lack — thus having greater existing demand for — players with his particular skillset. Or perhaps players perform better against worse teams, which are already more prone to acquire new players in order to improve their roster. While we use a variety of control variables to address these and several other alternate hypotheses, we can already see a problem with any hypothesis linking better performances to a higher pre-existing likelihood of being acquired by the opposing team: they do not account for the preparation effect. Indeed, with comprehensive game-by-game and aggregate performance statistics freely accessible through second-hand sources, it is particularly unclear why a team would logically place extra weight on a player's performance in a game that happens to precede a game against their own team.

The rest of the paper proceeds as follows. Section 2 provides relevant institutional background and describes the data. Section 3 describes our model of overattention. Section 4 presents our main empirical results. Section 5 addresses alternate explanations. Section 6 assesses the magnitude of the apparent first-hand experience bias. Section 7 concludes. All appendices are available online.

2. Background and data

2.1. NBA player transitions

The NBA is a professional basketball league that currently employs roughly 500 players on 30 teams. NBA players tend to have short careers (approximately 99 percent retire before the age of 40) during which they change teams relatively often (roughly once every two years). There are two main ways NBA players can change teams. First, a team may sign a player to a new contract, provided that player is no longer under contract with another team. Second, teams may trade for another team's player (or players) on their existing contract(s), usually in exchange for a player (or players) on their own team.[3]

All else equal, a team's decision to sign or trade for a new player suggests that the team values the player more than a team that did not acquire the player. In practice, however, both types of decisions are imperfect indicators (albeit for somewhat different reasons) of the team's valuation of the player as a prospective member of their team. To start, since a player must agree to sign a new contract, a team's decision to sign a new player presumably reflects the player's interest in joining that team as well as the team's interest in acquiring the player. While trades generally do not require the player's consent, they can still reflect other considerations beyond the team's perception of his value as a member of the team. For instance, there are many rules that restrict the types of trades teams can make, which can cause teams interested in trading one player for another to include other, non-targeted players in the deal simply as a means to bring the trade in compliance with league rules.[4]

There are two 'seasons' during which NBA teams may acquire new players: the 'offseason' (roughly June through October of each year) when no games are played, and the 'regular season' (November through mid-April of the following year) during which each team plays 82 games according to a predetermined schedule. In our sample (described shortly), a little more than half of new player acquisitions occur in the offseason (54.2%, including 56.5% for signed players and 52.7% for traded players) with the rest occurring in the regular season.

2.2. Team personnel involved in player-acquisition decisions

Our understanding of team decision-making is invariably limited by a lack of available data as to which personnel actually watch relevant games, how information is aggregated and transmitted among relevant personnel, and which personnel actually have a say in teams' final decisions. From what we do know, some aspects of how teams learn about other teams' players are fairly standardized, yet there does not appear to be a single, generalizable template at higher levels of decision-making. For this reason, we will (throughout the text) generally treat teams as if they were individual decision-makers, abstracting from more refined descriptions of what goes on inside the black box of team decision-making.

[3] Trades can also include cash, the rights to acquire a new player who has not previously played in the NBA, and, on occasion, coaches. Related to this, teams may also claim a player who is placed 'on waivers' by his original team, which gives other teams the opportunity to acquire the player on his existing contract — in effect, the opportunity to trade for the player in exchange for nothing. If a player is not claimed off waivers, he is then eligible to sign a new contract with another team.

[4] Appendix B discusses some other rules that restrict teams' player-acquisition decisions, with a particular focus on rule changes that occurred during our sample period (though empirical results presented in this appendix suggest these rule changes did not meaningfully affect teams' decision-making as it relates to their weighting of first-hand experience relative to second-hand information).

Table 1
Relevant personnel in NBA team decision-making.

	Typically watch versus games?	Typically watch prep. games?	Final say in team decisions?
Coaches	Yes	Yes (on film)	Some*
Scouts	No	Yes	No
Owners	Some	No	Some
Other executives	Some	Some*	Some

"Other executives" may includes team presidents, managers, and directors. See Appendix C for a detailed discussion of the roles of each type of personnel in team decision-making and the evidence on which these classifications are based.

* Some coaches also serve as a team executive. In these cases, the individual would typically observe preparation games (in their capacity as coach) and may also have the final say in player-acquisition decisions (in their executive role). However, it is less likely team executives who do not coach would routinely observe preparation games, or that coaches without an executive role would have authority in player-acquisition decisions.

Table 2
Summary statistics.

	Mean	S.D.	Percentiles			Obs.
			10 th	50 th	90 th	
Games in year before switch	62.66	23.72	22.0	75.0	82.0	4,567
Games per opponent	2.02	0.66	1.07	2.06	2.83	4,567
Mean performance per game	7.79	5.16	2.17	6.63	15.2	4,567
Std. Dev. performance (across games)	5.81	2.5	2.38	5.89	8.95	4,567

With these caveats in mind, Table 1 lists relevant team personnel and their potential roles in teams' hiring decisions. Here, it should also be noted that teams often employ several individuals in each of the four listed categories (see Appendix Table C1). Therefore, a 'Yes' or 'Some' designation in Table 1 may apply for at least one (but not necessarily all) individual(s) in that role. Furthermore, the designations only reflect an interpretation based on available information (discussed at length in Appendix C).

2.3. Our data

Our analysis uses box score data (collected from basketballreference.com) for regular season NBA games between October 1983 and April 2016. For each game, the box score provides the date, the location, the number of fans in attendance, the two teams that competed, the players who played for each team, and various performance statistics for each of these players. Since players are identified by name, we can also identify cases in which a player changes teams from the box score data.[5]

Our sample contains 4,567 cases in which a player changes teams. On average, players in the sample play in 62.66 games during the year preceding their first game on their new team, with 2.02 games played against each of the 22 to 29 other teams in the league at that time. Using the efficiency metric that adds up the good performance statistics and subtracts the bad performance statistics (see footnote 1), players' average per-game performance during the year preceding their transition is 7.79 while the average standard deviation (standard deviation across games, average across player-team switches) in performance is 5.81.[6] See Table 2 for these and other summary statistics.

3. Model

Our model is developed in three steps. In Section 3.1, we present a simple learning model that characterizes how a team balances second-hand information with first-hand experience in evaluating a player on another team. This learning model will serve as a building block for a multinomial model of teams' hiring decisions, which we present in Section 3.2. We then operationalize the model for empirical estimation in Section 3.3.

Before proceeding, we do note that our model is quite stylized. Most notably, the model treats teams as if they are individual decision-makers and does not distinguish between versus and preparation games. As mentioned in Section 2.2 (and discussed at greater length in Appendix C), realistically team decision-making involves input from scouts, coaches, owners, and other team executives. In addition, first-hand experiences in versus games and in preparation games may inform

[5] Our box score data is supplemented by other forms of data from the same source, including data on players' birth states and birth dates, data indicating each team's 'conference' and 'division' according to the NBA's (primarily geographical) categorizations, and data indicating whether a player who changes teams signed a new contract upon joining their new team or was traded on an existing contract. Also note, we can only identify a player as having been a member of a team if that player shows up in at least one box score for that team. Thus, if a player is traded from some team A to some team B, but never actually plays for team B because he is immediately traded from team B to some team C (who the player does end up playing for), this appears in our data (and is treated in our analysis) as a direct transition from team A to team C.

[6] For comparison, the average per-game performance for all players in our data (regardless of whether they transition to a new team) is 7.82, while the average standard deviation is 6.33. As addressed in footnote 11, this suggests that players with moderate performance levels are disproportionately likely to change teams.

decision-making through somewhat different channels involving different personnel (see Table 1). With that said, we will make the distinction between versus and preparation games in some of our empirical specifications.

3.1. Team learning

A team's assessment of the overall performance of a player on another team is presumed to be based on information concerning the player's performances, denoted by $\mathrm{P}(g)$, for each game $g \in \mathcal{G}$ that the player competes in during a given evaluation period. We assume that the team receives *second-hand information* concerning the player's performance in *all* games and that, for a subset of these games, the team receives an additional (i.e. redundant) signal of the player's performance through *first-hand experience*.[7]

Letting $\mathcal{G}^{\mathrm{FH}} \subset \mathcal{G}$ denote the subset of games for which the team attains first-hand experience (implying its rank, $|\mathcal{G}^{\mathrm{FH}}|$, denotes the number of such games), the mean first-hand performance signal is then $\bar{\mathrm{P}}^{\mathrm{FH}} \equiv |\mathcal{G}^{\mathrm{FH}}|^{-1} \sum_{g\in\mathcal{G}^{\mathrm{FH}}} \mathrm{P}(g)$. Meanwhile, the mean second-hand performance signal is simply the mean performance in all games: $\bar{\mathrm{P}} \equiv |\mathcal{G}|^{-1} \sum_{g\in\mathcal{G}} \mathrm{P}(g)$.

While the second-hand performance signal $\bar{\mathrm{P}}$ can be understood as an objective measure of the player's overall performance, we allow for the possibility of a bias in the team's subjective assessment arising from 'overattention' to first-hand experience. The *overattention parameter* $\omega \geq 0$ captures the extent of the potential bias — or more specifically, the extent to which the team weights first-hand experience in relation to second-hand information. The team's subjective assessment is then given by

$$\widehat{\mathrm{P}} \equiv \frac{\bar{\mathrm{P}} + \omega \cdot \bar{\mathrm{P}}^{\mathrm{FH}}}{1+\omega}. \tag{1}$$

3.2. Player acquisition

We now adapt the simple learning model developed thus far to a multinomial framework describing how a player is matched to one of multiple potential new teams. Here, $\mathcal{J}$ denotes the set of potential new teams and $j \in \mathcal{J}$ denotes a team in this set. In this multi-team setting, it will be useful to separate a team's assessment as given in (1) (except now indexed by j) into a component that is common to all teams and a team-specific component as

$$\widehat{\mathrm{P}}_j \equiv \bar{\mathrm{P}} + \left(\frac{\omega}{1+\omega}\right) \cdot \overline{\mathrm{PD}}_j, \tag{2}$$

where $\overline{\mathrm{PD}}_j \equiv \bar{\mathrm{P}}^{\mathrm{FH}}_j - \bar{\mathrm{P}}$ is the player's *mean performance deviation* in games for which team j attains first-hand experience. Team j's overall valuation of the player is presumed to depend (though not exclusively) on its assessment of the player's performance as follows:

$$v_j = f(\widehat{\mathrm{P}}_j) + \mu_j, \tag{3}$$

where μ_j simply encapsulates all other factors that affect team j's valuation of the player besides its assessment of the player's performance. Here it may be implicitly presumed that the team's valuation of the player is based on an underlying profit-maximization objective and that, all else equal, a team's profits increase with the performance of its players.[8]

Taking f to be linear ensures that the additive separability of $\bar{\mathrm{P}}$ and $\overline{\mathrm{PD}}_j$ in team j's assessment of performance $\widehat{\mathrm{P}}_j$ is maintained in team j's overall valuation, which is now:

$$v_j = \alpha + \beta\,\bar{\mathrm{P}} + \gamma\,\overline{\mathrm{PD}}_j + \mu_j. \tag{4}$$

We may note the overattention parameter can be expressed in terms of the coefficients β and γ as $\omega = \frac{\gamma}{\beta-\gamma}$, implying that teams are unbiased (i.e. $\omega = 0$) if $\gamma = 0$.

We assume that the player joins the team with the highest valuation. Abstracting from the possibility of a tie, team $j \in \mathcal{J}$ thus acquires the player if and only if $j = \arg\max_{k\in\mathcal{J}}\{v_k\}$. This matching rule captures the intuitive notion that a team with a higher valuation of a player is more likely to acquire that player. Alternatively, v_j can be interpreted as team j's bid, in which case the rule could be thought of as representing an auction in which the player is matched to the highest bidder.

In practice, there are multiple (and often complex) processes by which NBA players are matched to a new team. As discussed in Section 2.1, some players switch teams by signing a contract with that team while others are traded on an existing contract. Certainly it is not clear that an auction would provide a reasonable description of a player who is traded. Even for those who sign a new contract, the negotiation process is rather opaque, and affected by many rules constraining

---

[7] This characterization is compatible with two potential ways in which a team might attain second-hand information. First, the team may learn from the player's performance statistics on a game-by-game basis, as if reading the box scores, which are widely published in newspapers and online. Second, the team may learn from the player's cumulative performance statistics, which are also widely published.

[8] Higher-performing players may increase team profits by increasing the demand for tickets or team merchandise, or by helping a team qualify for — and thus sell additional tickets during — the playoffs (an annual tournament after the regular season that determines a league champion). Alternatively, a team's objective may be to maximize the proportion of games it wins or its likelihood of winning a championship, though higher-performing players would presumably serve these objectives as well.

the contracts teams can offer players.[9,10] Lastly, it is possible that only a subset of a player's potential new teams actually consider the player for acquisition (conversely, a team may only consider a subset of all available players). While such considerations would complicate modeling the matching mechanism with greater specificity, they can be crudely understood as entering our model through μ_j.

3.3. Adaptations for empirical estimation

We now adapt our model for empirical estimation. To reflect the relevant unit of observation in our data, we will now use i to denote a particular player and t to denote the time at which a particular player changes teams. Thus, μ_{ijt} represents the component of team j's valuation of player i at time t that does not depend on the team's subjective assessment of the player's performance. Decomposing $\mu_{ijt} = m_{ijt} + \epsilon_{ijt}$ into an observed component, m_{ijt}, and an unobserved component, ϵ_{ijt}, team j's valuation of player i at time t is then

$$v_{ijt} = \alpha + \beta\,\bar{\mathrm{P}}_{it} + \gamma\,\overline{\mathrm{PD}}_{ijt} + m_{ijt} + \epsilon_{ijt}. \tag{5}$$

In our empirical estimations, m_{ijt} will include a variety of control variables (described in the next section). We assume ϵ_{ijt} is independently and identically distributed according to the type I extreme value distribution. This assumption allows us to estimate our model as a pure conditional logit regression where the probability that team $j \in \mathcal{J}_{it}$ acquires player i at time t can be expressed in closed-form as:

$$\Pr[j = \arg\max_{k \in \mathcal{J}_{it}} \{v_{ikt}\}] = \frac{e^{\gamma\,\overline{\mathrm{PD}}_{ijt} + m_{ijt}}}{\sum_{k \in \mathcal{J}_{it}} e^{\gamma\,\overline{\mathrm{PD}}_{ikt} + m_{ikt}}}. \tag{6}$$

Note, we were able to eliminate $\alpha + \beta\bar{\mathrm{P}}_{it}$ from this expression because it is common to all of player i's potential new teams.

4. Results

In this section, we report empirical estimates capturing the effect of first-hand experience on teams' player-acquisition decisions. Here and in subsequent discussions, a "player" will (unless otherwise noted) refer to any player who changes teams and a "team" will refer to any of the player's potential new teams (i.e. any team *except* for the player's pre-transition team).[11] The relevant evaluation period is taken to be the one-year period that precedes the player's transition to a new team.

With our pure conditional logit matching model, the coefficients on control variables that would be the same for all teams are not identified. This includes the coefficient β in Eq. (5) on the player's mean performance in all games $\bar{\mathrm{P}}_{it}$ (i.e. the second-hand performance signal). However, we do include several team-varying control variables: the team's winning percentage during the year; the mean performance of all players that played against the team; the player's mean utilization (in minutes per game) in versus games and in preparation games; the numbers of potential and actual versus games; the numbers of potential and actual preparation games; dummies indicating whether any actual versus games and any actual preparation games occurred; and dummies indicating whether the team is located in the player's birth state, whether the team is located in a state that borders the player's birth state, whether the team and the player are in the same conference, and whether the team and the player are in the same division.[12]

There was at least one versus game for 75% of player-team observations, while the same proportion had at least one preparation game. However, only 22% of players had at least one versus game and at least one preparation game for all of their potential new teams. Whenever there was no versus and/or preparation game, the associated mean performance and mean utilization variables were set to the player's mean in all games (implying the mean performance deviation was set to zero), while the dummy variable(s) indicating whether or not at least one such game occurred allowed us to control for these undefined measures. Removing these observations from the sample did not meaningfully impact the magnitude or interpretation of our results (see Appendix Tables A2 and A3).

4.1. Main results

The first two columns of Table 3 provide estimates (with and without control variables) of the overall effect of first-hand experience, as captured by the coefficient γ on the player's mean performance deviation in versus and preparation games

---

[9] As examples of such constraints (which have changed over time), the NBA imposes both maximums and minimums on players' salaries, contract lengths, total team salaries, and the number of players per team.

[10] The potential role of player preferences will be addressed at length in Section 5.

[11] Recall, our main sample excludes players who do not change teams and players who leave their team, but remain unmatched thereafter. Since higher-performing players tend to change teams less often and lower-performing players are more likely to remain unmatched, players with moderate performance levels are overrepresented. With that said, unmatched players and players who do not change teams are considered in Section 6.1 and in Appendix D, respectively. In both cases, the estimates of interest are either statistically indistinct or greater than those reported in Table 3, which suggests that, if anything, the overrepresentation of moderate-performing players leads us to underestimate the magnitude of first-hand experience effects.

[12] In all of our analyses, standard errors are clustered around each subset of player observations.

Table 3
Estimates of conditional logit player-team matching model.

Mean performance deviation in versus and prep. games	0.033	0.028		
	(.004)	(.005)		
	[<.001]	[<.001]		
in versus games only			0.023	0.017
			(.004)	(.004)
			[<.001]	[<.001]
in preparation games only			0.012	0.013
			(.003)	(.004)
			[<.001]	[.001]
Controls?	No	Yes	No	Yes
Observations	126,404	126,404	126,404	126,404

Standard errors are in parentheses and p-values are in square brackets. See Appendix Table A4 for an expanded version of this table with estimates for all control variables (where applicable).

(see Eq. (4)), while the last two columns provide separate estimates for versus games and for preparation games. In all four specifications, the estimated coefficient(s) of interest are positive and statistically significant. This indicates that teams have a tendency to acquire players with better-than-usual performances in versus and preparation games, as we would expect if teams are 'overattentive' to first-hand experience (as modeled in Section 3).[13,14]

As detailed in Appendix F, if we use the total (instead of mean) performance deviation across applicable games or the normalized mean performance deviation (so that its unit is the standard deviation of the player's mean performance deviations across teams) as our measure of relative performance, the estimated coefficient(s) of interest remain positive and statistically significant (see, in particular, Appendix Tables F1 and F2). The same is true when using a variety of alternatives to the efficiency metric as our underlying measure of performance (Appendix Tables F3 through F9).

To provide a sense of the overall impact of first-hand experience, suppose teams j and k initially have the same valuation of a given player and are thus equally likely to acquire that player. A one-unit increase in the player's mean performance deviation in versus and preparation games with team j would then, all else equal, make team j 2.8% more likely than team k (in terms of the odds ratio) to acquire the player, based on the estimate reported in the second column of Table 3.[15] As seen in the fourth column of Table 3, this overall effect can be separated into a versus effect and a preparation effect of comparable strengths (while the magnitude of the versus effect is larger, the difference is not statistically significant).[16]

There were three control variables for which the estimates were highly significant in both controlled specifications: (i) the birth state dummy, which had the largest coefficient (.389 in the regression reported in the second column of Table 3, with $p < .001$); (ii) team winning percentage ($-.259$, $p < .016$); and (iii) mean performance of all players against the team (.106, $p < .001$). A more detailed look at how the sizes of these effects compare to the effect of first-hand experience is provided in Appendix G.

The results of two placebo tests are presented in Table 4. Besides re-estimating the impact of a player's performance in a team's versus and preparation games, here we also estimated the impact of performances in games that the team does not experience first-hand, yet are temporally close to those that do. Specifically, our first placebo test includes an estimate for the game after the versus game, while the second test includes estimates for all games within a 2-game window of the versus game.[17,18] The estimated magnitude and statistical significance of the versus and preparation effects are nearly identical to those without placebo games, while the placebo game estimates are all statistically indistinct from zero.

[13] Furthermore, the effect does not appear to be confined to a small subset of teams in our sample. For instance, when separately estimated for each team, the coefficients on relative performance in versus and preparation games are positive for 28 out of the 30 teams (Appendix Table A6). While most of these team-level estimates are, on their own, not statistically significant — which is not surprising considering they are (on average) identified based on 1/30th of the observations in our sample — they are collectively distinct from zero ($p < .001$) and indistinct from each other ($p = .21$).

[14] As addressed in Appendix E, the effect also does not appear to be driven by teams responding to serious injuries to key player(s). This is noteworthy because injuries are largely unexpected, and may compel a team to acquire a player to help replace the injured player(s) with less planning and deliberation than usual.

[15] Note, this does not mean that team j's probability of acquiring the player increases by 2.8 percentage points. Instead, the ratio between the probability that team j acquires the player (p_j) and the probability that team k acquires the player (p_k) increases from 1 to 1.028. This can be seen from the conditional logit player-acquisition probabilities given in (6), as (suppressing the i and t subscripts) the marginal effect of an increase in a player's mean performance deviation on the odds ratio is given by $\partial(p_j/p_k)/\partial\,\overline{\mathrm{PD}}_j = \partial(e^{\gamma\,\overline{\mathrm{PD}}_j+m_j}/e^{\gamma\,\overline{\mathrm{PD}}_k+m_k})/\partial\,\overline{\mathrm{PD}}_j = \gamma \cdot e^{\gamma\,\overline{\mathrm{PD}}_j+m_j}/e^{\gamma\,\overline{\mathrm{PD}}_k+m_k} = \gamma \cdot p_j/p_k = \gamma \approx .028$, given $p_j = p_k$.

[16] Without controls, the difference — in which the versus game coefficient is roughly 90 percent larger than the preparation game coefficient — is statistically significant. The extent of this difference is on par with the difference as depicted in Fig. 1, where the slope of the linear best fit is roughly 85 percent larger for versus games than for preparation games. We do note, however, that Fig. 1 only includes observations for which there was at least one game of the corresponding type (unlike our present analysis).

[17] For each type of placebo game, we also included controls for the numbers of actual and potential games played, a control for the player's mean utilization, and a dummy indicating whether any such games were played (analogous to the existing controls defined for versus and preparation games).

[18] While it may be natural to wonder if teams scout opponents two games prior to the versus game, in which case it would be dubious to classify such games as placebos, it is only standard practice for NBA teams to scout an opponent's last game before playing their own team. As discussed in footnote 2, for example, NBA arenas only reserve seats for scouts representing the participating teams' *next* opponents.

Table 4
Placebo test results.

Mean performance deviation in games that take place		
2 games before versus games		0.000
		(.004)
		[.939]
1 game before versus games (in prep. games)	0.013	0.013
	(.004)	(.004)
	[<.001]	[.001]
in versus games	0.018	0.018
	(.004)	(.004)
	[<.001]	[<.001]
1 game after versus games	-0.003	-0.003
	(.004)	(.004)
	[.502]	[.438]
2 games after versus games		0.001
		(.004)
		[.712]
Observations	126,404	126,404

Both regressions include control variables. Standard errors are in parentheses and p-values are in square brackets. Baseline estimates of the conditional logit player-team matching model are provided in Table 3.

Table 5
AR(1) model estimates.

Dependent variable: mean perf. deviation in	games versus future team	1 game before vs. future team	1 game after vs. future team	any game
Mean perf. deviation 1 game before game in dependent variable	0.083	0.062	0.096	0.079
	(.014)	(.014)	(.014)	(.003)
	[<.001]	[<.001]	[<.001]	[<.001]
Observations	6,384	6,384	6,384	172,036

Standard errors are in parentheses and p-values are in square brackets.

4.2. The validity of the preparation effect

The placebo test results also help alleviate any potential concerns that the preparation effect — which will prove especially useful for addressing alternatives (discussed at length in Section 5) to our overattention hypothesis — may simply be a byproduct of the versus effect. Namely, if the preparation effect merely arises because preparation games are temporally adjacent to versus games, we would naturally expect an analogous “next game effect.” However, no such effect is observed.

The validity of the preparation effect is reinforced by the results of AR(1) regressions estimating serial correlations between a player’s relative performance over consecutive games of interest. As shown in Table 5, the estimated autoregressive coefficient between a player’s relative performance in his future team’s preparation and versus games is positive and statistically significant (column 1). However, this estimate is on par with (and statistically indistinct from) the estimated autoregressive coefficients between each of these games and the adjacent placebo games (columns 2 and 3), as well as between any two consecutive games (column 4). This suggests that the positive relationship between a player’s relative performance in preparation and versus games involving his future team is not responsible for the preparation effect.[19]

Described at length in Appendix H, a final check on the validity of the preparation effect assesses whether a proxy for a team’s *level* of preparation predicts the strength of the preparation effect. In this exercise, we used data indicating the ‘odds on favorite’ (i.e. the expected winner) for regular season games to identify cases of team overachievement and underachievement — that is, winning despite being expected to lose and losing despite being expected to win — in versus games as proxies for high and low levels of preparation (respectively). Here, the idea is that a team that is more attentive to the preceding preparation game would be more prepared and hence more likely to overachieve (and less likely to underachieve) in the versus game. Presuming a player’s performance in a preparation game is more salient when the (preparing) team is more attentive to that game, our notion of team overattention naturally suggests that the preparation effect (but not the versus effect) would be stronger when the team overachieves in the versus game and weaker when the team underachieves, which is exactly what we find.

[19] As seen in Appendix Table A5, there also is no apparent relationship between a player’s relative performance and relevant characteristics of the team that is preparing for the player’s team (i.e. the player’s next opponent), such as the team’s winning percentage, the new-player acquisition rate, and being located in the player’s birth state. Thus, it does not appear that players are simply more prone to better performances in games that precede games against teams with a greater predisposition to acquire them, nor do we see any indication that nonrandom scheduling plays a role.

4.3. A systematic bias?

According to the overattention hypothesis, the observed first-hand experience effects reflect a systematic and suboptimal bias in teams' player-acquisition decisions. To further test this idea, it helps to consider whether a team's prior first-hand experiences with a newly-acquired player relate to the realized value of that player to the team. If teams are *not* biased, we might expect a player's performance in games providing first-hand experience to his future team to be positively related to measures that reflect the player's value to the team — such as the player's overall performance on his new team, the team's winning percentage after acquiring the player, and the team's utilization of the player, and the length of the player's tenure (number of games played) as a member of that team. While all of these measures are positively and statistically related to the second-hand performance signal capturing the player's overall performance prior to the transition, they are not statistically related to the player's relative performance in games providing first-hand experience to his future new team (see Appendix Tables A7, A8, A9, and A10).[20] This suggests that first-hand experiences are not genuinely linked to the realized value of the player to the team.

The notion of a *systematic* bias in team-decision-making is also supported by additional results indicating that the observed effect of first-hand experience is not limited to specific situations within our domain. In particular, we find that first-hand experience has a positive and statistically significant effect on teams' player-acquisition decisions regardless of whether a player is acquired through a trade or signed to a new contract (see Appendix Table A11); whether the acquisition happens in the regular season or in the offseason (Appendix Table A12); whether the player's overall performance level is relatively high or relatively low (Appendix Table A13); whether the team's first-hand experience is attained at 'home' or on the 'road' (Appendix Table A14); and whether the transaction occurs in the earlier years or in the later years covered by our sample (Appendix Table A15).[21]

5. Alternate hypotheses

We now address alternate explanations for the observed first-hand experience effects.

5.1. Player agency

Players may have some degree of control in determining their new team along with a tendency to perform better when being observed by a team they would like to join. Addressing the second part, a player may naturally be excited to perform against a preferred future team due to the presence of that team's current players (desired future co-workers) and fans. In addition, if players prefer to join teams in or near cities where friends and family reside, players' friends and family may also attend such games. This idea is partially addressed by the dummy indicating whether the team is located in the player's birth state. Its effect is positive and statistically significant, suggesting players are more likely (perhaps due to their agency in determining their next team) to be acquired by teams near their place of birth. Still, players' preferences may vary due to other idiosyncratic reasons not captured here.

While there may be doubts as to whether it would explain the preparation effect, there is arguably a more compelling reason to reject the player agency hypothesis.[22] Namely, players who are traded generally have little or no say over their new team — with the exception of some high-profile "stars."[23] Thus, according to the player agency hypothesis, we would expect the effect (if any) of first-hand experience among traded players to be negligible compared to the effect among players who sign with a new team. However, the estimated effect among traded players is still positive, statistically significant, and indistinct from the effect for players who sign with a new team. Furthermore, the effect among traded players who are not "stars" — i.e. those who do not earn "all-star" status at any point in their career, and for whom agency in trading decisions is particularly unlikely — is nearly identical to the effect among "stars" who are traded (see Appendix Tables A11 and A18).

[20] While a player's relative performance in games providing first-hand experience to his future new team is not statistically related to the number of games played by that player during his full tenure as a member of that team, it *is* statistically (and positively) related to the number of games played during his first year on the team (Appendix Table A10). This finding suggests that past first-hand experiences may continue to influence teams' valuations of newly-acquired players in the short-run, while teams may eventually learn to objectively value such players in the long-run. We will revisit this idea in Section 7.1.

[21] Of note, the test in Appendix Table A15 does not provide evidence of a meaningful time trend, as the estimated effect of first-hand experience in the first half of our same is the same as the estimated effect in the second half of our sample (based on a rough median split). Similarly, we find no evidence of a linear time trend (Appendix Table A16). With that said, in Appendix B we consider various regime tests to explore the possibility that the effect may have varied with changes in the rules that govern teams' player-acquisition decisions or as a result of the so-called "analytics movement." The results of these tests raise the possibility that the effect may have diminished in the final years of our sample, but are also compatible with the possibility that the effect did not change. See Appendix B for details.

[22] In preparation games, the preferred team would generally only be represented at the game by its (relatively inconspicuous) scout sitting in the stands. Certainly, however, a player's preference to join a particular team could naturally translate to better performance in preparation games if the player is strategic. We will consider this possibility in Section 5.2.

[23] As discussed at length in Appendix I, players who are not considered "stars" often emphasize their lack of control over when, whether, and where they will be traded.

5.2. Player auditions

Players may treat games observed by a desired future team as "auditions." As in the player agency hypothesis, the player's preferences may create a correlation between his performances and his future match, except in this story the player strategically allocates effort to impress the team, thus increasing the likelihood that the team decides to acquire him.

If a player auditioned for his future team by exerting extra effort in preparation and versus games, we would naturally expect higher performances in consecutive preparation and versus games followed by lower performance in the next game. As a result, a player's performance in a versus game against his future team would tend to be closer to his performance in the preceding preparation game than in the next game. However, a player's performance against his future team is, on average, slightly closer to his next game performance than to his preparation game performance (mean absolute differences of 5.41 and 5.40, respectively).

The lack of a special relationship between a player's performances in consecutive preparation and versus games with his future team is reinforced by the AR(1) estimates in Table 5. Unlike what we would expect if players treated these games as auditions, the estimated autoregressive coefficient between a player's performance deviations in consecutive preparation and versus games with his future team is statistically indistinct from the coefficients between each of these games and the adjacent placebo games as well as the coefficient between any two consecutive games in the year preceding the transition.[24]

5.3. Team showcasing

A team's interest in a player may be known to the player's current team. In response, the player's current team may use the interested team's versus and preparation games as opportunities to "showcase" the player by giving him more or better opportunities to perform in such games, with the intent of increasing the interested team's valuation and thus strengthening the current team's bargaining power in a potential trade. A team's pre-existing interest in a player — which would naturally correlate with a higher likelihood of eventually acquiring that player — may therefore also correlate with that player's performance in versus and preparation games as a result of efforts by the player's current team to increase the value they could attain in return for that player if traded to the interested team.

To address this "showcasing" hypothesis, it may be useful to consider *how* a team could showcase a player. One likely possibility is through increased utilization (playing time), though this channel is controlled for by our controls on a player's mean utilization in versus and preparation games.[25] A team could also showcase a player by running more offensive plays for that player (thus providing more opportunities to score) without increased utilization. That said, if we exclude scoring-based statistics from our performance measure, we still see a positive and significant effect of first-hand experience.[26]

While the showcasing hypothesis suggests that a first-hand experience effect may, in contrast to the audition hypothesis, be driven by strategic actions of the player's current team as opposed to the player himself, it too would naturally suggest that a player's (on average, elevated) performance in a versus game against his future new team would tend to be closer to his (also elevated) performance in the preceding preparation game than in the game after the versus game when showcasing has ended. As noted earlier, however, a player's performance against his future team is, on average, slightly closer to his next game performance than to his previous game performance.

The showcasing hypothesis is also challenged by our finding that the effect of first-hand experience is as large for players who sign a contract with their new team as it is for players who are traded (Appendix Table A11). After all, a team gains nothing by showcasing a player to a team that eventually signs the player. With that said, a team may showcase a player to improve their bargaining position in a potential trade, but if a trade does not materialize the prospective trade partner may still have an elevated interest in (and hence, likelihood of signing) the player during free agency. However, a team would presumably have little incentive to showcase a player *after* the annual "trade deadline" when trades are no longer possible. Nonetheless, whether estimated for signed players or for all players in our sample, the estimated effect of first-hand experience in games played after the trade deadline remains positive and statistically significant, and is statistically indistinct from the effect in games played before the trade deadline (see Appendix Table A17).

[24] The audition hypothesis also cannot simultaneously explain (a) the lack of a long-term impact of first-hand experience (i.e. beyond one year; see Appendix J), and (b) the observation that the estimated first-hand experience effect is as large for traded players as it is for players who sign with a new team (Appendix Table A11). That is, since player contracts usually span multiple years, players who sign with a new team upon the expiration of a previous contract will have anticipated the possibility of switching teams at that time. However, traded players often have one or more years remaining on their existing contract, making them less likely to have anticipated switching teams. Thus, the observed first-hand experience effect among traded players would, according to the audition hypothesis, indicate that players are auditioning well before their current contracts end. Assuming players' preferences for teams are serially correlated, higher relative performances from 'auditions' occurring one to two years before switching teams would thus also predict a higher likelihood of being acquired by that team, but this prediction is not supported by the data.

[25] The estimated coefficient on versus game utilization is positive and borderline significant (with $p = .084$ in the regression that separately estimates the versus and preparation effects and $p = .027$ in the regression that estimates the overall effect), offering some support for the idea that showcasing may contribute to the versus effect as estimated without controls. The estimated coefficient on preparation game utilization is negative and statistically insignificant (in both controlled regressions), which suggests that teams do not showcase players through increased playing time during an interested team's preparation games.

[26] Here, "scoring-based statistics" refer to points, attempted and made field goals, and attempted and made free throws. If we also exclude assists (passes leading to a made field goal) from our performance measure, the estimates of interest are still positive and statistically significant. See Appendix Tables F5 and F6.

5.4. Fan overattention

Perhaps it is not the team, but its fans who overweight first-hand experiences with other teams' players. Teams may then be inclined to acquire players who perform relatively well in games that are experienced first-hand because these players are over-valued by their fans, which could then lead to higher profits from ticket and merchandise sales. However, fans do not typically watch their team's preparation games, which rarely involve their own team.[27] Thus, "fan overattention" does not explain the preparation effect.

In addition, the fan overattention hypothesis suggests that a team's fans would (all else equal) have more interest in watching a newly-acquired player with better relative performance in past first-hand experiences with the team. As a result, the team's utilization of the player would naturally be higher — especially in home games — than it otherwise would; fan attendance at the team's home games may also be higher. However, higher relative performance in games providing first-hand experience to a player's future new team does not predict higher overall utilization, disproportionate utilization in home games, or higher home game attendance for his new team (see Appendix Tables A9, A19, and A20).

5.5. Private learning

Perhaps second-hand performance statistics do not tell the full story, while first-hand experiences allow a team to learn about a player's unreported attributes, such as his ability to defend. A first-hand experience effect could then arise if a player's relative performance in versus and preparation games was positively correlated with the amount of private information conveyed in such games.[28] Such a correlation could naturally arise through the player's utilization, as higher utilization would allow a player to accumulate higher performance statistics while also giving teams more time to learn about unreported attributes. Our controls for the player's mean utilization in versus and preparation games help account for these potential effects. In addition, our controls for the numbers of versus and preparation games help account for the number of such learning opportunities.

Our results suggest that any effect of private learning stemming from a player's utilization or from the number of versus and preparation games is negligible, as we fail to reject the hypothesis that these four coefficients are jointly equal to zero ($p = .22$ for the model in Table 3, column 2). Even if we accept the coefficients at face value, the estimates imply that the average effect of an extra game providing first-hand experience is just one-seventh the size of the effect of a one-unit increase in the player's mean performance deviation, thus undermining the notion that teams' evaluations are meaningfully dependent on extra information conveyed in versus and preparation games. Even if the amount of private information conveyed through first-hand experience was not fully captured by utilization and the number of versus and preparation games, it would correlate strongly with these measures. Their lack of an apparent impact seems to rule out any secondary channel linking a player's relative performance and a team's private information that could explain the observed first-hand experience effects.[29]

5.6. Bad teams

Worse teams tend to allow better performances by opposing players. They also acquire new players at a higher rate. These two factors could generate a positive correlation between a player's relative performance against a team and the team's likelihood of acquiring the player. However, our controls for the mean performance of all players against the team and for the team's winning percentage help account for such possibilities. Furthermore, this hypothesis cannot explain the preparation effect since a player's relative performance in a game would not be correlated with the quality of the player's next opponent.[30]

5.7. Void-filling

Lastly, teams that lack players with a particular skill would tend to allow better performances by players who possess that skill while also acquiring such players at a higher rate. For example, a slow team may allow fast players to perform especially well, while also having a higher demand for fast players. However, this void-filling hypothesis also cannot explain the preparation effect because a player's relative performance does not plausibly depend on whether the next opponent lacks players of a similar skillset.

---

[27] A team only participates in its preparation game if they play the same opponent in two consecutive games. In our sample, teams play in less than 2 percent of their preparation games.

[28] Under standard assumptions, extra learning of this sort would increase the likelihood that a team's valuation of the player is extreme (in either direction) in comparison to other teams' valuations, thus increasing the likelihood of acquiring the player.

[29] The lack of a long-term effect of first-hand experience (see Appendix J) provides another reason to reject the private learning hypothesis, as it is implicitly predicated on the notion that a player's unreported attributes are serially correlated from year to year — otherwise, a team's private information concerning such attributes, as attained through first-hand experience in the preceding year, would be largely outdated and thus not indicative of the player's potential future value to the team. Thus, the private learning hypothesis would (incorrectly) imply substantial long-run persistence in the effect of first-hand experience.

[30] See Appendix Table A5, which shows that relevant characteristics (e.g. the team's winning percentage and new-player acquisition rate) of the player's current opponent — but not the next opponent — are related to the player's relative performance level.

6. Quantifying the magnitude of the bias

This section considers two exercises that will help us get a better sense of the magnitude of the apparent first-hand experience bias. First, in Section 6.1, we estimate the overattention parameter ω and use it to quantify the extent to which teams are disproportionately influenced by first-hand experience. Then, in Section 6.2, we estimate the proportion of players who are matched to the "wrong" team as a result of the bias. Also see Appendix G, where we estimate the dollar-value welfare loss associated with the bias, quantify the magnitude of the bias in terms of team wins, and compare the size of the first-hand experience effect to the sizes of the effects of other factors that affect team decision-making.

6.1. Measuring overattention to first-hand experience

The overattention parameter, ω, introduced in Eq. (1), captures the degree to which teams' player-acquisition decisions are swayed by first-hand experience relative to second-hand information. As previously noted, we can express ω in terms of γ and β, which capture the respective effects of first-hand experience and of second-hand information, as $\omega = \frac{\gamma}{\beta-\gamma}$. However, while γ was estimated in Section 4.1, we were not able to identify β because all of a player's potential new teams observe the same second-hand performance signal describing his overall performance in all games.

Therefore, to estimate ω, we first need to adapt our empirical framework in a manner that allows us to estimate β. To do this, we add an outside option representing the possibility that a player who leaves his original team might not be matched to a new team, while expanding our sample to include roughly 2,500 players who we observe leaving a team but remaining unmatched thereafter. The value of the outside option is normalized to zero, which captures the idea that a player will only be matched to a new team if at least one team has a positive valuation of the player.[31] With the no-match option, we estimate $\beta = .39$ and $\gamma = .02$ (with $p < .001$ for both effects; see Appendix K for details), which implies $\omega = \frac{\gamma}{\beta-\gamma} = .054$. This indicates that teams weight players' mean performance in games providing first-hand experience roughly 5.4 percent as heavily as they weight the second-hand signal conveying players' overall performance in all games. Recall, any $\omega > 0$ represents overattention in that players' performances in games providing first-hand experience are already embedded in the second-hand performance signal.

Next, we use ω to calculate the weighting of a player's performance in a game that is experienced first-hand relative to a game that is not.[32] Letting $v^+(g)$ denote the increase in a team's valuation of a player if the player's performance in game $g \in \mathcal{G}$ was one unit higher, we want to compute $\lambda \equiv \frac{v^+(g|g\notin\mathcal{G}^{\text{FH}})}{v^+(g|g\in\mathcal{G}^{\text{FH}})}$. Here, λ captures the effect of higher performance in a game that does *not* provide first-hand experience, expressed as a fraction of the effect in a game that does. Using Eq. (4), we can re-express λ as:

$$\lambda = \frac{|\mathcal{G}^{\text{FH}}|/|\mathcal{G}|}{\omega + |\mathcal{G}^{\text{FH}}|/|\mathcal{G}|}, \tag{7}$$

where $|\mathcal{G}^{\text{FH}}|/|\mathcal{G}|$ is the fraction of the player's games that the team experiences first-hand.

By calculating the distribution of λ across all applicable player-team observations, we find that the mean and the median are both $\lambda = .55$/ This indicates that a player's performance in a game that is experienced first-hand is typically weighted about $1/\lambda \approx 1.8$ as heavily as it would have been weighted if the game was not experienced first-hand.[33] It is also worth noting that $1-\lambda$ offers a natural analog to the 'inattention parameter' estimated in other contexts (here, reflecting the degree to which attention to second-hand information is diminished in relation to first-hand experience).[34] Following this interpretation, our estimates would imply a mean inattention parameter of $1-\lambda = .45$, which is within the range of estimates from other empirical settings.[35]

---

[31] A positive valuation may more accurately be understood as a positive net valuation, as the gross value of the player to the team would need to exceed his salary as well as the potential lost option value from filling the vacancy on the team's roster. Though we abstract from such complications, neither of these costs would be trivial in light of the NBA's rules mandating a minimum player salary and a maximum team roster size.

[32] The overattention parameter ω does not provide a direct measure of how a team's weighting of a player's performance in a single game depends on whether it was experienced first-hand. There are two reasons for this. First, part of the impact of a player's performance in a game the team experiences first-hand still arises through the second-hand signal conveying a player's overall performance in all games, yet this is not captured by ω. Second, ω represents the weight on the player's mean performance in games providing first-hand experience (i.e. $\bar{P}^{\text{FH}}$), relative to the weight on the player's mean performance in all games ($\bar{P}$). However, teams generally only attain first-hand experience in a small fraction of the player's games, which suggests the degree of overattention on a per-game basis would be substantially higher than ω.

[33] By this measure, the extra bias from first-hand experience is roughly 80% as strong as the influence of objective performance information in teams' player-acquisition decisions. Though not directly comparable, this is on par with the relative magnitude of NBA teams' previously-reported bias in favor of early draft picks, as Staw and Hoang (1995) estimate that the effect of being selected in the first round of the NBA draft instead of the second round on a player's career length is roughly 70% (3.3 years versus 4.6 years) as large as the effect of a one-standard-deviation increase in a "scoring index."

[34] In the simplest formulation, if the true value of an object is $a+b$, b (but not a) is an attribute drawing less than full attention, and θ is the inattention parameter, the object's perceived value is then $a+(1-\theta)b$ (see DellaVigna, 2009 for relevant background). With that said, since teams appear to place a nonzero weight on redundant performance information conveyed through first-hand experience, the discrepancy in the effective attention paid to games providing first-hand experience as compared to other games naturally reflects overattention to first-hand experiences as opposed to inattention to other information.

[35] Other estimates of the inattention parameter include .18 to .45 in Hossain and Morgan's (2006) study of shrouded shipping costs on eBay, .46 to .59 in DellaVigna and Pollet's (2009) study on investors' attention to earnings announcements occurring on Fridays, .75 in Chetty et al's. (2009) study on consumers' attention to sales taxes, and .31 in Lacetera et al.'s (2012) analysis of a left-digit bias in evaluating used car mileage.

Table 6
Summary of simulations.

		% of players with different matches		
		(A)	(B)	(C)
Estimated model with first-hand experience bias	(A)	0	2.3	35.2
'Optimal' model without first-hand experience bias	(B)		0	34.4
Pure random matching model	(C)			0

Simulations from the estimated model (A) use the coefficients from the conditional logit estimation described in Section 6.1 (see Appendix Table K1). Simulations from the 'optimal' model (B) use the coefficients from (A), except the coefficient on a player's performance deviation in versus and preparation games is set to zero. The random matching model (C) simply uses $v_{ijt} = \epsilon_{ijt}$. Each simulation uses the same type-I extreme value draw of ϵ_{ijt} for models (A), (B), and (C). Results are based on the average of 10,000 simulations.

6.2. Player-team mismatch

To get a sense of the proportion of player-team matches affected by the bias, Table 6 presents the simulated outcomes of three different models: (A) our estimated model from Section 6.1; (B) an 'optimal' model that lacks a first-hand experience bias (i.e. with γ set to zero), but is otherwise the same as model (A); and (C) a pure random matching model, in which each of a player's possible outcomes are equally likely. In each simulation, all models used the same draws of the unobserved ϵ_{ijt} error terms. Thus, whenever a player's simulated match differs between models (A) and (B), the discrepancy can be attributed to the first-hand experience bias. As seen, our simulations indicate that the bias causes roughly 2.3 percent of players in our sample to be mismatched in this sense.

7. Additional discussion

In this paper, we provided evidence from the NBA labor market that employers' hiring decisions are overly influenced by first-hand experiences with job candidates. In particular, we found that teams are biased in favor of acquiring players with better-than-usual performances in games against their own team (the versus effect) and in games immediately preceding such games (the preparation effect). In closing, we discuss whether teams may learn from their mistakes due to overattention (Section 7.1), whether the observed first-hand experience bias might generalize to hiring decisions in other industries (Section 7.2), and the potential implications of our findings for understanding first-hand experience effects in other decision-making contexts (Section 7.3).

7.1. Do teams learn?

Several of the empirical tests considered in this paper offer clues as to whether or not teams may learn from their mistakes due to overattention. As a whole, the evidence is mixed, and the answer may depend on what exactly is meant by learning. Accordingly, we will consider three separate aspects of learning and address the implications of our analysis for each of these aspects.

First, we consider whether a team that acquires a player based on overattention eventually learns that they overvalued that particular player. If teams do *not* learn in this sense, then a team's valuation of a newly-acquired player would continue to be positively associated with the player's relative performance in past games that the team experienced first-hand. However, a team's past first-hand experience with a newly-acquired player does not predict the team's utilization of that player (in terms of playing time) nor does it predict the length of the player's tenure (number of games played) as a member of that team (see Appendix Tables A9 and A10). Thus, to the extent that utilization and tenure length reflect a team's valuation of a player *after* acquisition, it appears that teams do eventually learn that players acquired due to overattention are less valuable than they previously believed. That said, if the first-hand experience bias is *fully* overcome in teams' valuations of acquired players, it may be reasonable to expect a *negative* relationship between past first-hand experience and tenure length. In that case, the lack of an observed relationship may reflect partial learning in the sense that a team's biased impressions are not completely overcome.[36] Furthermore, our finding that past first-hand experience is statistically and positively associated with the number of games played by that player in his *first* year on that team (again see Appendix Table A10) suggests that learning of this sort may be a gradual process.

As a second aspect of learning to consider, we might wonder whether a team that acquires a player based on overattention learns to avoid making the same mistake in future player-acquisition decisions involving *other* players. Presumably, if

[36] If we interpret the utilization result as evidence of complete learning and the tenure result as evidence of partial learning, the discrepancy could reflect heterogeneity in the extent of learning among team personnel. In particular, a player's utilization is generally at the discretion of the coach, and thus would primarily depend on the coach's valuation of the player, while tenure may, to a greater degree, reflect the valuations of other team personnel involved in the acquisition decision. In light of this, our findings could be interpreted as evidence of greater learning among coaches compared to other team personnel.

teams did learn to avoid repeating such mistakes, the effect of first-hand experience would decline over time. However, the lack of an apparent time trend (see Appendix Tables A15 and A16) casts doubt on this possibility. Thus, to the extent that teams learn that they overvalued players acquired due to overattention (as suggested by our utilization and tenure results), their apparent propensity to repeat the mistake with other players suggests that teams may not realize *why* they overvalued these players in the first place.

Lastly, we might wonder whether teams may ever learn to avoid the first-hand experience bias in player-acquisition decisions. While we cannot draw strong conclusions from our analysis (and despite the lack of an apparent time trend in the effect of first-hand experience over our full sample), our regime test results in Appendix B raise the possibility that teams became less susceptible to the bias in the final years of our sample. If so, such learning may have been spurred by the “analytics movement” and its far-reaching influence on NBA team decision-making in recent years (as opposed to teams learning as a direct result of their own mistakes). While the analytics movement did not suddenly give teams a newfound ability to recognize and overcome a first-hand experience bias in player-acquisition decisions (as the bias was always evident from traditional box score data), it may have allowed a widespread philosophical shift towards more objective and data-driven decision-making among NBA teams, leading to a reduction in the influence of first-hand experience. Note, this idea is still largely untested and future work would be needed to assess with confidence whether or not the effect of first-hand experience truly diminished in the final years of our sample, as our regime test results are also compatible with the possibility that the effect did not change.[37]

7.2. Generalizability to hiring decisions in other industries

Next, we consider whether a first-hand experience bias might generalize to hiring decisions in other industries. While this remains an open empirical question, some superlative characteristics of our setting may actually strengthen our confidence that the bias may apply elsewhere. With easy access to comprehensive performance data (and video footage of every game), the NBA provides unusually favorable conditions for employers to recognize and overcome a potential bias. Furthermore, NBA team owners tend to be successful executives in other industries, suggesting relatively high managerial competence and thus relatively low susceptibility to a potential hiring bias.[38]

With that said, most other jobs do not entail ‘contests’ involving direct competition with workers from other firms. Since the versus effect applies when teams are actively participating in such contests, its generalizability may, on its own, seem limited. However, the preparation effect suggests that the bias can also sway an impartial observer of a worker’s performance. Even so, both effects apply to a specific form of first-hand experience — on-the-job observation of a worker’s performance — attained in a contest-like setting.

As an example of another contest-like setting where similar effects may arise, suppose, after observing an opposing attorney at trial, a legal client (whether an individual or a business) is considering hiring that attorney for another case. If the client’s decision is excessively influenced by their prior observation of the attorney, this would be analogous to the versus effect in NBA teams’ decision-making since, in both cases, first-hand experience is attained from observing the worker perform on behalf of the employer’s opponent in the contest (whether a trial or a basketball game). A more natural analog to the preparation effect could then arise if, instead of being directly involved as a plaintiff or a defendant, the client previously observed the attorney while serving as a juror or as court stenographer.

Certainly, it may not be essential for a worker’s performance to be observed during a *contest* for a first-hand experience bias to apply. For instance, consider a decision to re-hire an independent contractor — such as an interior decorator, accountant, electrician, math tutor, or personal bodyguard — who the employer previously hired for another job. The key question is then whether the employer’s decision is excessively influenced by their first-hand observation of the contractor’s past job performance. Of note, the contractor’s previous job performance would have still directly affected the employer at that time, which is also true of the versus effect, though the preparation effect suggests that such ‘stakes’ are not necessary for a first-hand experience bias to exist.

To consider this idea in another non-contest setting, except without stakes in this sense, suppose an academic department is considering hiring a researcher from another institution. Incidentally, the chair of the hiring committee previously attended a conference seminar delivered by the researcher (who was not seeking a new job at the time). While the researcher’s previous seminar performance may not have directly affected the chair’s academic department at the time, it is conceivable that the performance could still be overweighted in the eventual hiring decision.

In all of the examples considered thus far, the employer’s first-hand experience is attained from observing the worker work on behalf of their original employer. With that said, if the chair of the academic department overweights the researcher’s conference seminar performance, then it is not much of a leap to think that such overweighting might also occur if the chair’s first-hand experience instead came from a recruiting seminar during an official campus visit (even though the researcher would no longer be observed working on behalf of their current employer). If that were the case, it would indicate that the bias may extend to first-hand experiences besides on-the-job observation of a worker’s performance —

[37] For additional details and discussion, see Appendix B.

[38] As of 2015, the average estimated net worth among NBA team owners was $3.3 billion (see http://www.businessinsider.com/sports-owners-net-worth-tenure-2015-10). With that said, owners do not necessarily have complete or direct influence in player-acquisition decisions (see Appendix C for details).

and could (returning to our example) perhaps even apply to the department chair's other first-hand experiences, such as a one-on-one meeting with the researcher during the campus visit or an initial screening interview with the department's full hiring committee.

Indeed, this notion that the bias could apply to other first-hand experiences (besides on-the-job observation of a worker's performance) fits with previously-discussed findings that employers tend to be too reliant on personal interviews (Highhouse, 2008; Dana et al., 2013). As noted earlier, however, this tendency could also reflect the personal motives of hiring managers. For instance, past research indicates that interview evaluations are often swayed by a candidate's physical attractiveness (Ruffle and Shtudiner, 2014), relationship status (Rivera, 2017), and whether they enjoy similar leisure activities (Rivera, 2012), which suggests that hiring managers may sometimes inflate their interview evaluations for candidates with whom they would personally like to affiliate.[39] In this light, the present work indicates that hiring decisions can be excessively influenced by first-hand experiences besides personal interviews, and without such discrepancies in the motivations of the hiring manager and the firm. In doing so, our findings lend support to the idea that a more pervasive and systematic decision-making bias may underlie (at least in part) previous findings regarding the overuse of personal interviews in hiring decisions.

A recent study by Leung (2017) finds that employers of online freelancers often over-extrapolate from past experiences (especially bad experiences) with employees of a given nationality when evaluating a current job candidate of the same nationality. Such findings raise the possibility that employers' hiring decisions may even be susceptible to a first-hand experience bias based on experiences with other individuals (besides the job candidate) with shared demographic traits. In this way, the bias may even contribute to ethnic and racial discrimination in hiring decisions (Bertrand and Mullainathan, 2004; Carlsson and Rooth, 2007), though such a link is only speculative at this point.

7.3. Relevance to other first-hand experience effects

While we have focused on employers' hiring decisions, the present work may also relate to evidence of first-hand experience effects in other decision-making contexts. For instance, investors tend to over-invest in assets (such as IPO subscriptions, 401k accounts, and stocks) that have previously brought high returns (Kaustia and Knüpfer, 2008; Choi et al., 2009; Chiang et al., 2011; Strahilevitz et al., 2011). First-hand experience may also distort consumers' expectations of future macroeconomic conditions, as those who have lived through better market conditions tend to be more bullish on future conditions, and those who have lived through higher-inflation periods tend to expect higher future inflation (Malmendier and Nagel, 2011, 2016).[40]

We note, however, that our findings are not explained by leading theories used to explain first-hand experience effects in these other domains. The consensus explanation for investors' over-reliance on past experience is reinforcement learning, which holds that good payoffs from a past action can bias future choices towards repeating that action.[41] In our setting, the relevant action (acquiring a particular player) is not generally repeated and payoffs from past actions are not the relevant source of first-hand experience. The broader premise that these effects are payoff-driven is also challenged by the preparation effect, as it suggests such effects can arise even when experiences are merely observational. In turn, Malmendier and Nagel's (2011, 2016) age-based learning hypothesis attributes differential weighting of past macroeconomic conditions to differences in age (and lifespan), and thus does not address a team's differential weighting of games occurring in the same timeframe.[42]

As discussed, our findings are instead compatible with a notion of overattention to first-hand experience, as modeled in Section 3. That is, while players' performances in all games are presumably conveyed (and objectively weighted) through second-hand summary statistics, teams can allocate extra attention to games they experience first-hand — while evidently assigning meaning to the redundant performance signals provided in such games. To the extent that an investor's returns from a past investment or the macroeconomic conditions during a consumer's lifetime are attentionally salient (despite

[39] As discussed in the introduction, a related line of research suggests that hiring decisions also tend to under-rely on algorithmic hiring aids (Kuncel et al., 2013; Hoffman et al., 2018). As with the apparent over-reliance on personal interviews, the under-reliance on algorithmic hiring aids may also be explained by hiring managers' personal motivations, as a hiring manager could perceive an algorithmic hiring aid as a threat to their autonomy, and thus resist its input.

[40] As other examples, Haselhuhn et al. (2012) find that a (non-informative) late-return fee on a video rental increases future compliance; Giuliano and Spilimbergo (2014) find that living through unfavorable macroeconomic conditions can affect beliefs regarding the degree to which luck determines success. Similarly, laboratory studies show that players in repeated games can be excessively swayed by prior experience (e.g. Erev and Roth, 1998; Camerer and Ho, 1999; Simonsohn et al., 2008), while psychology research shows that individuals, in choice situations ranging from gambling to medical diagnoses, tend to over-rely on personal experience in relation to relevant summary statistics (e.g. Weber et al., 1993; Hertwig et al., 2004).

[41] For theoretical background and laboratory evidence in multi-player games (in the non-sports sense), see Roth and Erev (1995), Erev and Roth (1998), as well as the generalizations by Camerer and Ho (1999) and Ho et al. (2007) permitting reinforcement of unchosen actions based on their hypothetical payoffs.

[42] Unlike reinforcement learning, age-based learning is compatible with a merely observational first-hand experience bias, though the necessity of potential) payoffs in driving such effects had remained an open question. As Malmendier and Nagel (2016) write, "it would be useful to further analyze the exact transmission channel of experience effects how do they depend on the prices of items personally consumed versus the CPI?" That said, the necessity of *realized* payoffs is inconsistent with Koudijs and Voth's (2016) finding of a first-hand experience effect among lenders exposed to a potential financial loss that was ultimately avoided.

the availability of more complete second-hand information sources), overattention may help explain first-hand experience effects observed in these domains too, though future work is needed to test such possibilities.

Declaration of Competing Interest

We declare that we have no relevant interests (financial or otherwise) to disclose.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2020.04.015.

References

Agness, S., 2016. Advance scout jimmy powell. Podcast Int. http://vigilantsports.com/ep-34-advance-scout-jimmy-powell/(accessed 4/18/2018).
Bertrand, M., Mullainathan, S., 2004. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. Am. Econ. Rev. 94, 991–1013.
Buckley, R., Norris, A., Wiese, D., 2000. A brief history of the selection interview: May the next 100 years be more fruitful. J. Manag. History 6, 113–126.
Camerer, C., Ho, T.-H., 1999. Experience-weighted attraction learning in normal-form games. Econometrica 67, 827–874.
Carlsson, M., Rooth, D.-O., 2007. Evidence of ethnic discrimination in the swedish labor market using experimental data. Labour Econ. 14, 716–729.
Chetty, R., Looney, A., Kroft, K., 2009. Salience and taxation: Theory and evidence. Am. Econ. Rev. 99, 1145–1177.
Chiang, Y.-M., Hirshleifer, D., Qian, Y., Sherman, A., 2011. Do investors learn from experience? evidence from frequent IPO investors. Rev. Financial Stud. 24, 1560–1589.
Choi, J., Laibson, D., Madrian, B., Metrick, A., 2009. Reinforcement learning and savings behavior. J. Finance 64, 2515–2534.
Dana, J., Dawes, R., Peterson, N., 2013. Belief in the unstructured interview: The persistence of an illusion. Judgment Decision Mak. 8, 512.
DellaVigna, S., 2009. Psychology and economics: Evidence from the field. J. Econ. Literat. 47, 315–372.
DellaVigna, S., Pollet, J., 2009. Investor inattention and friday earnings announcements. J. Finance 64, 709–749.
Dunleavy, M., & Eyen, J. (2009). Game preparation in. NBA Coaches Playbook: Techniques, Tactics, and Teaching Points, Ed. Giorgio Gandolfi, (pp. 287–298). Champaign, IL: Human Kinetics.
Erev, I., Roth, A., 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. Am. Econ. Rev. 88, 848–881.
Giuliano, P., Spilimbergo, A., 2014. Growing up in a recession. Rev. Econ. Stud. 81, 787–817.
Goldsmith, D., 2000. Soliciting advice: The role of sequential placement in mitigating face threat. Commun. Monogr. 67, 1–19.
Haselhuhn, M., Pope, D., Schweitzer, M., Fishman, P., 2012. The impact of personal experience on behavior: Evidence from video-rental fines. Manag. Sci. 58, 52–61.
Hertwig, R., Barron, G., Weber, E., Erev, I., 2004. Decisions from experience and the effect of rare events in risky choice. Psychol. Sci. 15, 534–539.
Highhouse, S., 2008. Stubborn reliance on intuition and subjectivity in employee selection. Ind. Organizat. Psychol. 1, 333–342.
Ho, T.-H., Camerer, C., Chong, J.-K., 2007. Self-tuning experience weighted attraction learning in games. J. Econ. Theory 133, 177–198.
Hoffman, M., Kahn, L., Li, D., 2018. Discretion in hiring. Q. J. Econ 133, 765–800.
Hossain, T., Morgan, J., 2006. plus shipping and handling: Revenue (non) equivalence in field experiments on ebay. BE J. Econ. Anal. Policy 5, 1–30.
Kaustia, M., Knüpfer, S., 2008. Do investors overweight personal experience? evidence from IPO subscriptions. J. Finance 63, 2679–2702.
Koudijs, P., Voth, H.-J., 2016. Leverage and beliefs: Personal experience and risk-taking in margin lending. Am. Econ. Rev. 106, 3367–3400.
Kuncel, N., Klieger, D., Connelly, B., Ones, D., 2013. Mechanical versus clinical data combination in selection and admissions decisions: A meta-analysis. J. Appl. Psychol. 98, 1060–1072.
Lacetera, N., Pope, D., Sydnor, J., 2012. Heuristic thinking and limited attention in the car market. Am. Econ. Rev. 102, 2206–2236.
Leung, M., 2017. Learning to hire? hiring as a dynamic experiential learning process in an online market for contract labor. Manag. Sci.. forthcoming.
Malmendier, U., Nagel, S., 2011. Depression babies: Do macroeconomic experiences affect risk taking? Q. J. Econ. 126, 373–416.
Malmendier, U., Nagel, S., 2016. Learning from inflation experiences. Q. J. Econ. 131, 53–87.
Rivera, L., 2012. Hiring as cultural matching: The case of elite professional service firms. Am. Sociol. Rev. 77, 999–1022.
Rivera, L., 2017. When two bodies are (not) a problem: Gender and relationship status discrimination in academic hiring. Am. Sociol. Rev. 82, 1111–1138.
Roth, A., Erev, I., 1995. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games Econ. Behav. 8, 164–212.
Ruffle, B., Shtudiner, Z., 2014. Are good-looking people more employable? Manag. Sci. 61, 1760–1776.
Simonsohn, U., Karlsson, N., Loewenstein, G., Ariely, D., 2008. The tree of experience in the forest of information: Overweighing experienced relative to observed information. Games Econ. Behav. 62, 263–286.
Staw, B., Hoang, H., 1995. Sunk costs in the NBA: Why draft order affects playing time and survival in professional basketball. Administrat. Sci. Q. 40, 474–494.
Strahilevitz, M.A., Odean, T., Barber, B., 2011. Once burned, twice shy: How naive learning, counterfactuals, and regret affect the repurchase of stocks previously sold. J. Market. Res. 48, S102–S120.
Weber, E., Bockenholt, U., Hilton, D., Wallace, B., 1993. Determinants of diagnostic hypothesis generation: Effects of information, base rates, and experience. J. Exper. Psychol. Learn. Memory Cogn. 19, 1151–1164.