The Causal Effects of Early-Career Playing Time on the Fourth-Year Performance of NBA Players *

Sean Fischer *The Annenberg School for Communication*

Recent shifts in professional basketball have led teams to place more urgency in drafting as well as possible. Draft picks must play out their initial years under team-friendly contracts that provide teams with increased salary cap flexibility. Yet, while this urgency has led to widespread discussion and research on improving teams' draft decisions, little attention has been given to identifying what teams can do to maximize their draft picks' performance and potential once they are added to their roster. However, learning and ecological psychology theories suggest that giving young players as much playing time as possible should lead to concrete improvements in their development and future performance. In this study, I test this causal theory by evaluating the relationship between the minutes a player receives in their first two seasons in the NBA and their fourth-year performance using a novel method of propensity score weighting that enables weighting for continuous treatment variables. I find that players who receive more minutes in their first two seasons have better fourth seasons and make larger jumps from their first two seasons to their fourth season, controlling for a broad set of potential confounders. These results have important implications for teams as they develop organizational strategies for the short- and medium-term.

Player development is a crucial aspect of achieving and maintaining success in all sporting domains. This process is made challenging in American professional basketball by relatively limited roster sizes and a small developmental league. Coupled with regulations limiting contract lengths to no more than five years, these limitations increase the risk associated with personnel decisions, especially those related to selecting new players through the NBA draft. As such, player evaluation has received substantial attention within NBA teams and among the sport analytic community. In particular, researchers and professional teams have devoted substantial attention to identifying player-level factors associated with future success to optimize the decisions made in the annual amateur draft.

The young players added to teams through the amateur draft process are their teams' best assets. After the draft, teams sign these young players to contracts with two guaran-

*Last updated January 22, 2021.

teed years and two team-option years. They are paid on a fixed scale, with little flexibility for negotiation, at a rate that is considered below market value (especially for those players drafted later in the first round and in the second round). If these players become regular contributors on the court, they give their teams additional salary cap flexibility, as they alleviate their teams' needs to sign as many higher-paid veterans.

Thus, as much as the individual players, teams are incentivized to do whatever they can to help their young players maximize their potential. However, teams only control a few elements of player development; they ostensibly control how their players practice and how they approach each game. In particular, teams control how much in-game experience their players get. This aspect is crucial as observers and participants commonly acknowledge that NBA basketball is played at a higher level than college or international basketball. Thus, if teams can help their players to learn how to play at the NBA level, they may reap future advantages.

So, what can teams do to increase the performance level of their young players on rookie-scale contracts once the players are already drafted?

Based on theories of learning, NBA-basketball should operate as a "kind" learning environment, in which players develop skills through repetition and feedback (Hogarth, 2001; Hogarth, Lejarraga, & Soyer, 2015). While teams vary in the specifics of their strategies, they share a general set of play styles and schemes that are not common in college basketball. As such, we should expect players who are given more opportunities to practice and play in actual-game or game-like situations to better adjust to playing at the higher pace, speed, and skill of NBA basketball, while also learning how best to integrate into their own team's playstyle and counter those of others. Players already regularly see early-career games in this light, describing them as regular learning opportunities and situations that will expose them to new play styles. This repeated exposure at game speeds allows them to develop the cognitive skills necessary to apply their already exceptional basketball skills faster. This learning process works best when players are given as many high-quality repetitions for which they can receive feedback beyond traditional coaching. Crucially, sensory feedback that assists in evaluating affordances and action possibilities is essential for future skill acquisition and performance improvement (Kiverstein & Rietveld, 2015), but is only accessible through game and highly game-like repetitions (Davids, Araújo, Vilar, Renshaw, & Pinder, 2013).

Therefore, we should expect those players that receive the most playing time early in their careers to improve the most by the end of their first contract. Yet, teams generally cannot afford to play their newest players the same amount as experienced veterans. Most teams are actively competing to win as many games as possible; playing their least experienced and least equipped players would lead to many avoidable losses. These team-side considerations ultimately lead to variation in the amount of playing time that young players new to the league receive in their first years.

In this study, I leverage this variation in playing time to test whether having more time to learn to play in the NBA affects players' future performance. I do so by estimating the causal effect of new players' playing time in their first two seasons affects their fourth-year performance. The fourth-year is an applicable benchmark because it is the last time every player is still playing under their rookie contracts before teams have to decide whether to extend each player a qualifying offer at the end of the season. I estimate this causal effect through a regression framework that leverages advances in propensity score weighting to minimize bias produced via confounding. I find that playing time in a player's first year as a significant and substantial effect on fourth-year performance.

Data

I restricted this study's scope to the analysis of players drafted between 2000 and 2013 who entered the NBA directly out of college. For each of these players, I collected their per-game and advanced stats from their last year of college play from Sports Reference. These data included their strength of schedule, minutes played, true-shooting percentage, effective-field-goal percentage, and win-shares-per-40. I also collected data on how many minutes these players played in their first and second seasons in the NBA and the length of their NBA careers using the nbastatR package in R (Bresler, 2021). nbastatR collects data from Basketball Reference and the NBA's official stats portal. All continuous measures were standardized in all analyses. Data collection was capped up through 2013 because players drafted after this year had not yet had the opportunity to begin playing through their rookie extensions.

For my outcome measure, I turned to Jacob Goldstein's Player Impact Plus-Minus (PIPM) (Goldstein, 2018). PIPM is an all-in-one metric intended to capture player contributions on offense and defense and summarize these contributions in a single number. Like other plus-minus metrics, it is a continuous measure that takes positive and negative values. PIPM improves on earlier all-in-one metrics like win shares and PER in that Goldstein designed the metric to accurately summarize past performance and accurately predict future performance. For each player in my sample, I collected their first-year, second-year, and fourth-year PIPM measures.

Notably, I standardized all of the continuous measures in my analyses. I did so following Gelman's general advice to standardize by two standard deviations (Gelman, 2008). Doing so increases the initial interpretability of results, as the continuous measures are placed on a scale that is closer to those of the included binary variables.

Method

The best way to assess the efficacy of any strategic decision is a randomized control trial. However, such a study is not plausible in this case, given that NBA teams are unwilling to randomly assign playing time during their seasons. In place of an experiment, I conduct an observational analysis using regression techniques.

To assess the relationship between early-career playing time and overall career performance, I utilized a regression-based framework for causal inference (Morgan & Winship,

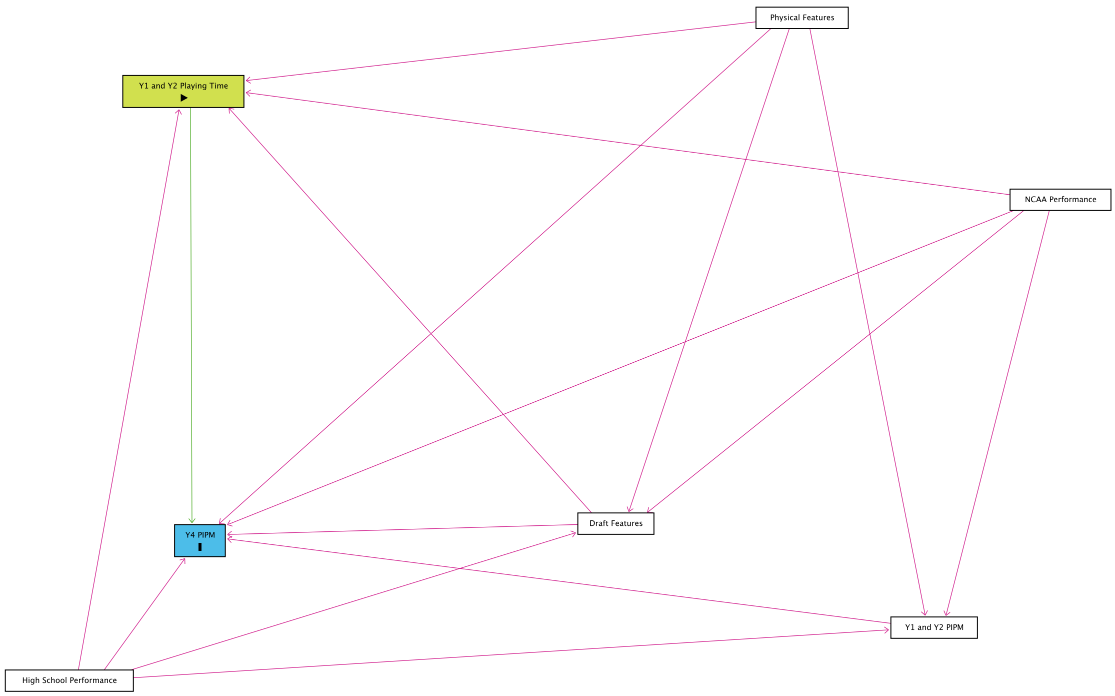

Figure 1: Directed acyclic graph (DAG) visualizing the hypothesized relationship between first year playing time, confounders, and career quality.

2015). For my model, I assume that a variety of observable features are related to fourth-year player performance. These include the treatment, playing time in a player's first two years; their prior skill level, operationalized by their placement on the RSCI ranking of high school players and their true-shooting percentage, effective field-goal percentage, win-shares-per-40, minutes played, and strength of schedule from their last NCAA season; their initial ability in the NBA, captured in their first- and second-year PIPM statistics; their age, operationalized as the number of years they spent in college; the features of their draft experience, in particular, the year they were drafted, the round they were drafted in, and the overall pick number they were selected at; and their physical attributes, proxied by their college body mass index.

Notably, all of the non-treatment measures in this model are associated with both the treatment of interest, early career playing time, and the outcome of interest, fourth-year performance. We can see in Figure 1 that these dual associations create multiple backdoor paths from early-career playing time to fourth-year performance. However, confounding raises concerns about potential bias in the estimated causal effect of treatment and needs to be addressed.

To resolve this concern over potential confounding, I would normally apply propensity score weighting or matching, but these methods generally require a binary treatment variable. The treatment here is a continuous measure without any well-defined cut points, making it difficult to justify any potential binarization.

However, Fong, Hazlett, Imai, et al. (2018) have introduced a method for estimating propensity scores for continuous treatment variables using the covariate balancing generalized propensity score (CBGPS). Applicable in both parametric and non-parametric cases, the CBGPS optimizes covariate balance while estimating each observation's propensity score. Here, I utilize the non-parametric CBGPS, as it allows for a reduction to zero in the correlations between treatment and observed covariates, even when the functional form of the propensity score model is misspecified (Fong et al., 2018).

Results

Early career Minutes Played

Do all rookies and second-year players play the same number of minutes? Intuitively, we know the answer is no. But, just how skewed is the allocation of playing time for young players?

We can see in the top panel of Figure 2 that there is a substantial amount of variation within and between years in regards to how many minutes teams give to first-year players. The mean number of minutes played for these draft classes ranges from a low of 1,868.27 minutes for the 2000 class to a high of 2,911.06 minutes for the 2008 class. Considering the yearly medians presents a similar story; the lowest median minutes played is 2000's 1,436.5 minutes and the highest median minutes played is 2005's 2,868 minutes. These averages tell us that, in most cases, NBA teams play their least experienced players for about 27% - 34% of the minutes available in their first two 82-game seasons. At the extreme, the most minutes played by any player in this dataset was 6,233 minutes, 79% of the minutes available in their first two seasons.

The lower panel of Figure 2 shows us the aggregate distribution of first-year playing time across all of the sampled draft classes. We can see here that the distribution is right-skewed. The mean value for the aggregate distribution is 2,369.53 minutes, slightly more than the median value of 2,206.5 minutes. These values represent about 30% and 28% of available minutes.

How Many Players Reach their Fourth Year?

Given that I am most interested in evaluating fourth-year performance when players are finishing their rookie contract, it is important to consider whether players regularly fail to reach this benchmark.

Of the 457 players captured in my original sample, only 50 failed to play in 4 full seasons. 31 of these 50 (62%) were second-round picks. The mean pick number used on these players was 32.56. There are only a few early first-round picks that fail to play in their fourth season. Of the 118 top-10 picks I include in my sample, only 3 (3%) failed to reach their fourth NBA season. These results should not come as a surprise, as it is well documented that NBA teams are likely to generate substantial sunk-cost effects when making personnel decisions (Staw & Hoang, 1995).

We should expect these patterns, as teams regularly have to weigh cutting some players. When teams cut players, they are still responsible for paying some or all of the money owed to a player in their contract. Second-round draft picks are particularly expendable in this case because they are not subject to the rookie pay scale and often sign contracts

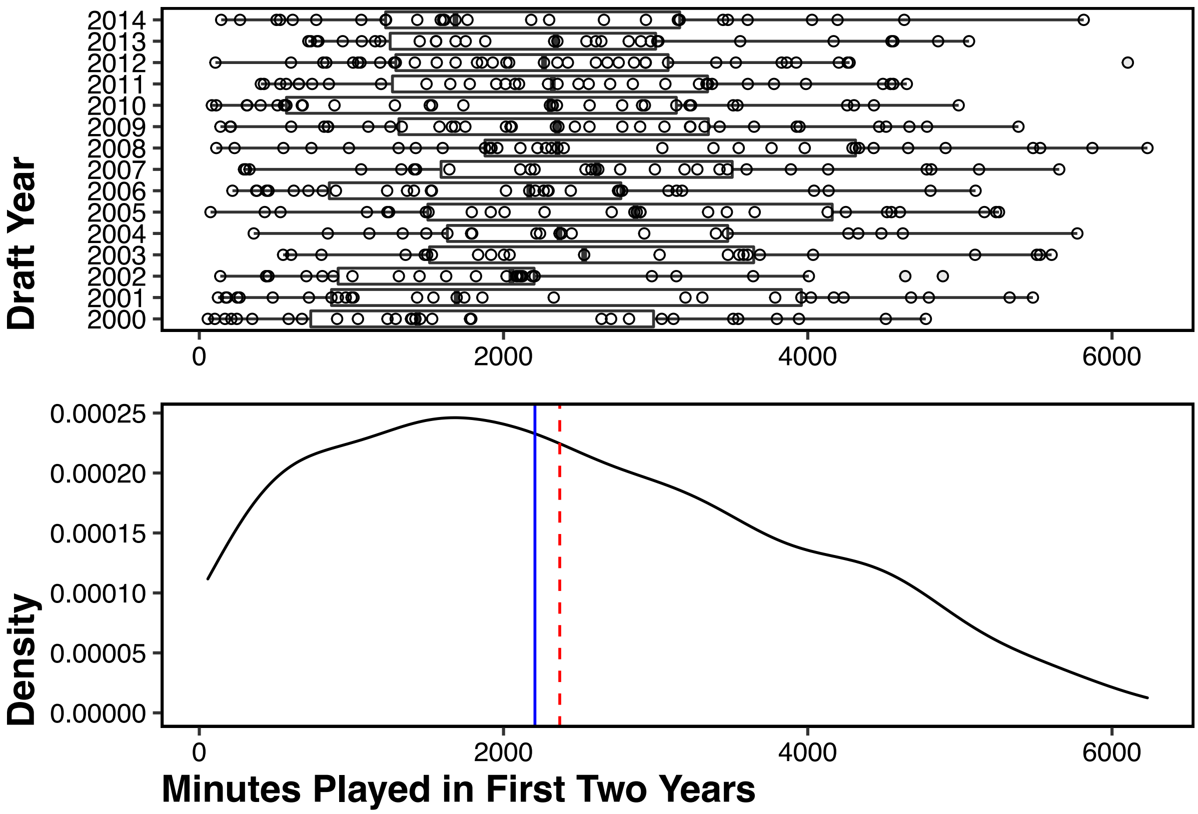

Figure 2: **Top** Distribution of minutes played by players from the 2000-2014 NBA Drafts in the first two years of their careers, split out by year. **Bottom** The distribution of minutes played by players from the 2000-2014 NBA Drafts in the first two years of their careers. The blue line marks the median value and the red dashed line marks the mean value.

that have fewer guaranteed years or no guaranteed years, reducing teams' salary cap penalties if they cut those players. Second-round draft picks are also often assigned to clubs' development teams to gain more experience. However, teams may ultimately decide never to bring a player in the development league to the NBA.

Given, then, that most players, and in particular a gross majority of first-round draft picks, who play in their first season also play in their fourth, I do not focus on the effects of playing time on players reaching their fourth professional season.

Propensity Score Weighting

While we can observe a basic relationship between increased playing time and career success, the causal relationship is undermined by potential confounding. This potential confounding manifests from the fact that both career success and first-year playing time may be associated with prior skill and success in college and high school, as well as features such as the draft investment the team made in the player, the amount of time they spent in college, and their physical attributes.

While randomized experiments alleviate such concerns, observational analyses need to employ alternative methods to resolve the issues raised by confounding. One such approach is propensity-score matching and weighting, which attempt to achieve balance on observed covariates for different treatment levels.

To generate propensity score weights for each observation, I used the CBPS package in R (Fong et al., 2019) and modeled first-year minutes played (centered and standardized continuous variable) on players' status as an RSCI-ranked high school player (binary); the number of years they spent in the NCAA (factor, five levels); their strength-of-schedule, minutes played, true-shooting percentage, effective-field-goal percentage, and win-shares-per-40-minutes from their last NCAA season (all centered and standardized continuous variables); the year they were drafted (factor, 15 levels); the round they were drafted in (factor, two levels); their pick number (centered and standardized continuous variable); and their college body mass index (BMI) (centered and standardized continuous variable).

To assess whether the propensity score model effectively addresses the potential confounding issues, we can evaluate the correlation between the treatment, first-year minutes played, and each covariate, before and after balancing. Panel A of Figure 3 displays the distribution of these correlations for both groups. As we can see, the balanced group is tighter to zero with a mean correlation of 0.0019, while the unbalanced group has a mean correlation of 0.096. Panel B of Figure 3 shows that the absolute value of the correlations

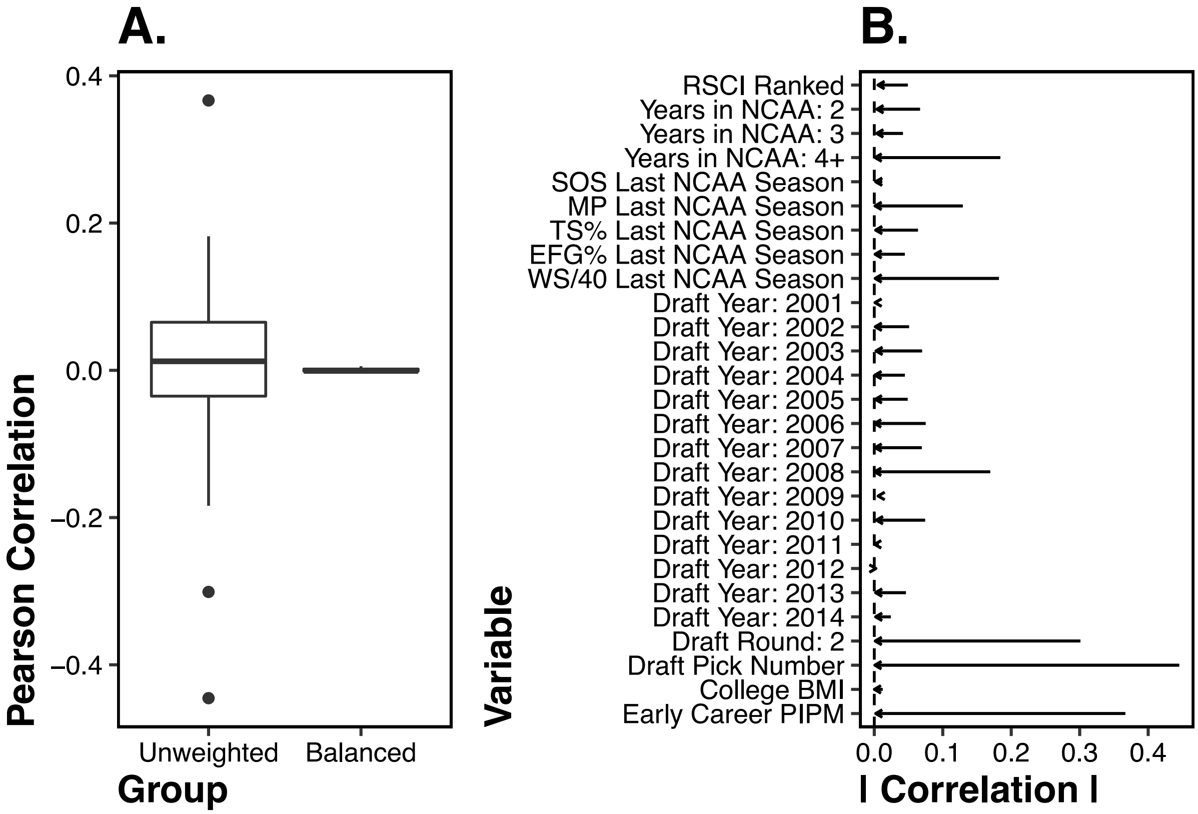

Figure 3: **A.** The distribution of correlations between first-year minutes played and potential confounders before and after balancing. **B.** The change in correlation between first-year minutes played and potential confounders when balancing applied.

for almost all covariates moved closer to zero through the balancing process. This move to zero signals that the covariate balancing process worked as intended.

By applying the weights generated through this process to standard OLS models regressing career success metrics on first-year minutes played and the additional confounders, I can estimate the causal effect of first-year minutes played on each outcome of interest.

Causal Effects of First-year Playing Time

To assess whether early-career playing time has a causal effect on fourth-year performance, I consider the effect of minutes played in players' first two years in the NBA on their fourth-year PIPM.

To estimate the causal effect of interest, I fit a model regressing players' fourth-year PIPM on their total minutes played in their first two years and the set of potential confounders using OLS. As we can see in Figure 4, early-career minutes positively and statistically significantly affect fourth-year PIPM. A two-standard-deviation increase in early-career minutes played produces a an increase in fourth-year PIPM approximately equal to one-fifth of two-standard deviations (β = 0.22, p = 0.000223, 95% CI = (0.10, 0.33)). In more practical terms, this means that increasing a player's total first- and second-year minutes by about 2,848.88 minutes, or about 17.37 minutes per game across two 82-game seasons, leads to a 0.86 gain in fourth-year PIPM. This 17.37 mpg increase is comparable to moving from the lower extreme to the higher extreme in our sample.

While this effect is statistically significant, is it substantive? These results indicate that moving a player's minutes from about what Marreese Speights received in his first two years to about what Deron Williams received in his first two years, would lead a player who otherwise would play as well as Jimmer Fredette in his fourth season to playing as well as Aaron Gordon did in his fourth season. Yet, teams may be put off by the notion of playing their least experienced players quite so much. An increase in playing time half as big would still push a player of Jimmer Fredette's level to Trevor Ariza's. These are not inconsequential changes when considering that teams would be paying these hypothetical players the same amount no matter what, given the structure of standard rookie contracts.

However, this observed effect may be sensitive to bias from unobserved confounding. To assess whether this result is sensitive to such unobserved confounding, I conducted a sensitivity analysis following the mode described by Cinelli and Hazlett (2020). In this case, for the null hypothesis that early-career minutes' treatment effect was equal to zero to not be rejected, an unobserved confounder would need to explain 9.66% of the residual variance in the treatment and the outcome. Alternatively, for the point estimate to go to

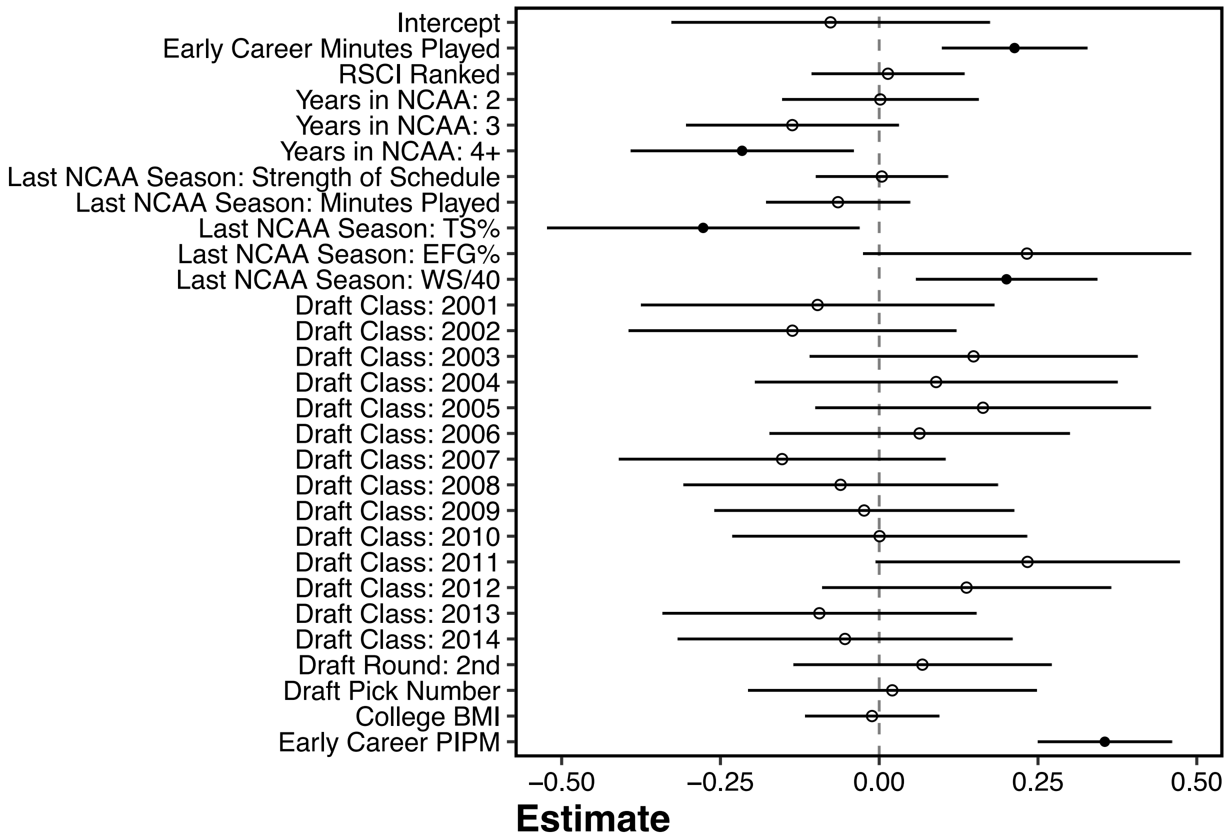

Figure 4: Estimated effect on fourth-year PIPM for early-career minutes played and covariates. Bars represent 95% confidence intervals. Estimates generated from weighted OLS.

zero, an unobserved confounder would need to explain 19.79% of the residual variance in the treatment and the outcome. With the propensity score weighting applied, an unobserved confounder is unlikely to explain this much variance. Using the same propensity score weights, an unobserved confounder five times as strong as the measure of early-career PIPM could only explain, at most, 0.01% of the treatment's residual variance.

Even so, we may wonder if some of this observed effect is the product of better players getting more early playing time and thus not indicative of an actual effect from receiving playing time. While the previous model did control for early-career PIPM, I can directly address this concern by estimating the causal effect of early-career playing time on the change in PIPM from a player's average in their first two years to their fourth year

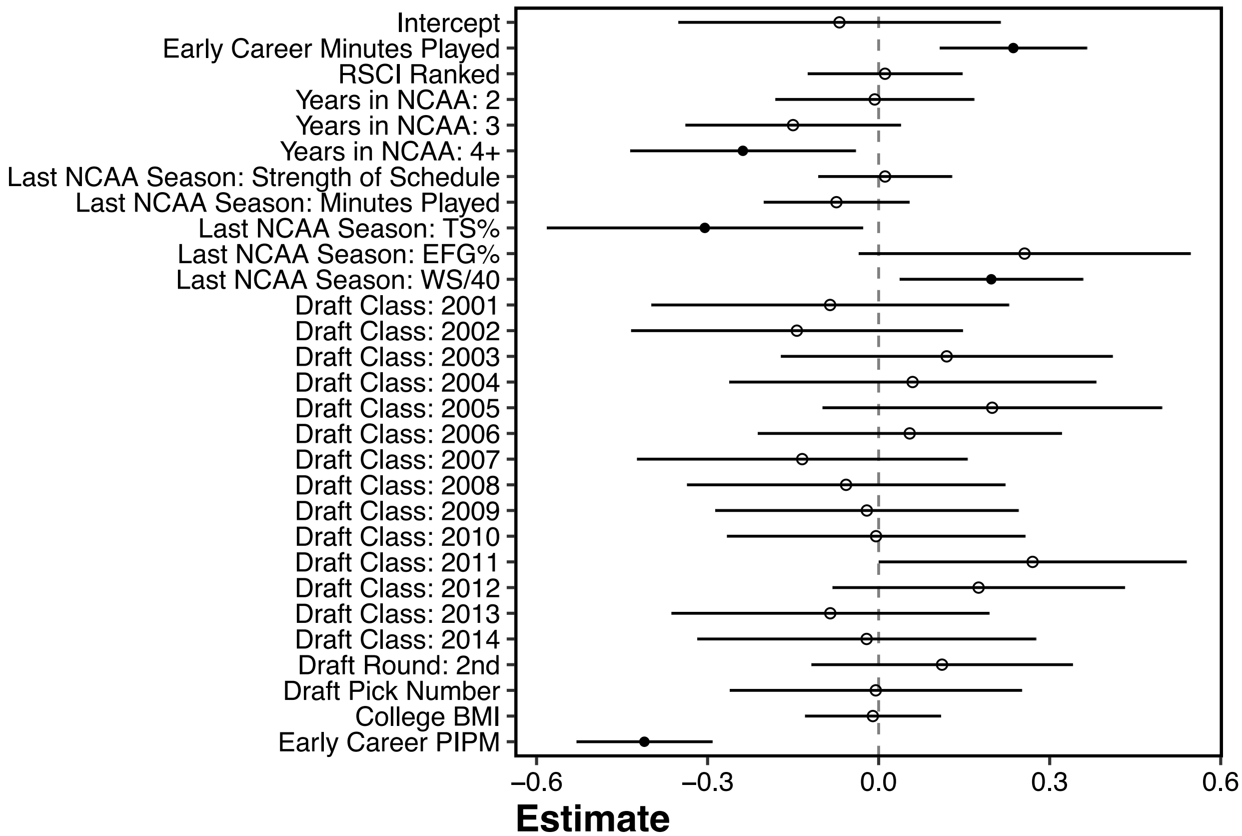

Figure 5: Estimated effect on the change in PIPM from an average of a player's first two years to their fourth year for early-career minutes played and covariates. Bars represent 95% confidence intervals. Estimates generated from weighted OLS.

($\Delta = PIPM_{y4} - \frac{PIPM_{y1}+PIPM_{y2}}{2}$). In estimating this model, I include the same controls and weights as used in the first model.

The results from this model fitting are visualized in Figure 5. We can see that, once again, early-career minutes have a significant positive effect, this time on performance improvement. The effect is similar in magnitude to that in the first model, where a two-standard-deviation change in early-career minutes played leads to about one-fourth of a two-standard-deviation increase in this change in PIPM (β = 0.24, p = 0.00039, 95% CI = (0.11, 0.37)). In practical terms, this means that the same 17.37 mpg increase over a player's first two years leads to a 0.84 PIPM increase in the jump in performance from their first two years to their fourth year.

This result shows that giving additional playing time to players in their first and second years increases the size of the jump they make going into their fourth year, holding all else constant. More directly, this evidence suggests that players are more likely to meet or exceed the expectations when given additional playing time, regardless of how good they already are.

Practically speaking, this estimate suggests that giving a player a 17.37 minutes-per-game increase would take a player's fourth-year jump from that of Corey Brewer (Δ_{PIPM} = 0.310), Jeremy Lamb (Δ_{PIPM} = 0.300), or Wilson Chandler (Δ_{PIPM} = 0.330) to that of Shelvin Mack (Δ_{PIPM} = 1.140), Jae Crowder (Δ_{PIPM} = 1.135), or Kevin Durant (Δ_{PIPM} = 1.125). Among players who would otherwise have a fourth-year jump similar to this latter group, the same increase in minutes played would make their fourth-year jump comparable to Chris Bosh (Δ_{PIPM} = 2.080) or Bradley Beal (Δ_{PIPM} = 2.105).

To assess whether this result would be sensitive to unobserved confounding, I repeated the first model's sensitivity analysis. In this case, for the null hypothesis that early-career minutes' treatment effect was equal to zero to not be rejected, an unobserved confounder would need to explain 9.33% of the residual variance in the treatment and the outcome. For the point estimate to go to zero, an unobserved confounder would need to explain 19.5% of the residual variance in the treatment and the outcome. With the propensity score weighting applied, an unobserved confounder is unlikely to explain this much variance; with weights applied, an unobserved confounder five-times as strong as the measure of early-career PIPM could only explain, at most, 0.01% of the residual variance of the treatment.

Notably, the relationships between early-career performance and the two outcomes of interest, fourth-year PIPM and the jump in performance between a player's first two years and their fourth year, are in opposite directions. This opposition makes sense, though, and provides important face validity to the model results. Players who are already successful in the NBA, as evidenced by higher PIPM metrics in their first two seasons, should also have high PIPM metrics in their fourth seasons since they are already good. However, because those players are already good, they often have a smaller performance gap to close to reach their peaks. As such, their jumps from their first two seasons to their fourth should be smaller because they are already near their top level of performance in their first two seasons. Take Chris Paul, for example. In his first two seasons, he generated PIPM values of 2.32 (12th best in the sample) and 5.71 (5th best in the sample), for an average of 4.015 (tied for 3rd best in the sample). His fourth-year PIPM was 3.75 (12th best in the sample), leading to a difference of -0.265 PIPM (-0.16 standard deviations away from the mean difference). Had Chris Paul struggled more early in his career before reaching a high level of success, his difference would have been greater. These patterns suggest that both models identify patterns we should expect to see in the data.

Discussion

NBA teams increasingly see their most recently drafted players as investment opportunities. Because league rules often force these players to accept less-than-market-value contracts, teams can use the amateur draft system to improve their rosters at a cost much less than they'd have to spend to make the same types of improvements through free agency. This realization has increased the overall perception of the value of draft picks across the league and raises the pressure placed on executives and scouts to draft the "correct" players. However, little attention has been paid to what teams can do to improve their early-career players' performance before they have to decide whether to extend a qualifying offer to a player at the end of the player's fourth year in the NBA. Here, I build on a psychological theory of learning environments (Hogarth, 2001; Hogarth et al., 2015) to consider whether teams can help their early-career players improve by increasing their in-game playing time in their first two seasons.

This theory posits that NBA basketball is a kind-learning environment, meaning that NBA games are structured so that patterns repeat and obvious feedback is regularly provided. As such, players who have more opportunities to receive that feedback by playing more early in their careers should be better by their fourth year than those who received less playing time. The hypothesized positive effect of playing time on performance is grounded in ecological theories of skill acquisition in sport, which argue that athletes need opportunities for feedback in-game or highly game-like contexts (Davids et al., 2013; Kiverstein & Rietveld, 2015).

However, testing this theory is complicated because playing time is not randomly assigned. Instead, it is based on various factors that are also associated with fourth-year performance, introducing concerns about confounding in any observational analysis. To account for these potential biases on any causal estimates, I employed covariate balancing generalized propensity score weighting (Fong et al., 2018) to break the association between the confounders and the continuous treatment measure, the sum total of minutes played in a player's first two seasons. The final result is an estimate of the causal effect of playing time in a player's first two seasons on their fourth-year performance, as well as the change in performance from their first two years to their fourth year.

My analyses show that players who play more early in their careers become better players later in their careers, as measured by PIPM (Goldstein, 2018), than those who play fewer minutes early in their careers. The effect is modest, but substantial; players who receive an increase of about 8.7 minutes-per-game in their first two years (assuming an 82-game season) should improve from the level of a player who only started seven games over six seasons in his NBA career (Jimmer Fredette) to one who started 731 games over sixteen seasons (Trevor Ariza). Furthermore, an increase in playing time in a player's first two years is associated with an increase in the size of the jump they make from their performance in their first two years to their performance in their fourth year.

This analysis is limited in a few crucial ways. First, the players considered were those drafted between 2000 and 2013 who entered the NBA after playing in the NCAA. This sampling strategy was necessary to establish a consistent pool of data, but omitted notable NBA players who entered out of high school (e.g. LeBron James and Dwight Howard), as well as those who entered the NBA after playing in an international league (e.g. Yao Ming and Pau Gasol). This sample also drops any players who never played in the NBA. Yet, my sample captures 457 of the 829 (55%) players drafted during this time period, even with these constraints. Beyond missing the set of players who did not play in the NCAA, this sampling and data collection strategy overlooks and disregards minutes played in the NBA's G-League (formerly the D-League), a developmental league for players not quite ready to contribute to NBA rosters. Minutes accrued in these settings may be valuable in skill acquisition and player development. However, that possibility seems implausible because the quality of competition in the G-League makes the games insufficient information spaces for advancing skills to an NBA level.

Additionally, the analysis relies on the assumption of no unobserved confounders. Any potential features that went unmeasured that are associated with both the treatment (early-career playing time) and outcome (fourth-year PIPM) could bias the results. Even so, sensitivity analysis indicates that an unobserved confounder would need to explain a little more than 9% of the residual variance of both the treatment and outcome for the null hypothesis of a true-effect equal to zero to not be rejected in both models and a little less than 20% of the residual variance of both the treatment and outcome for the point estimate to go to zero. As such, while these limitations are present, they are not likely to undermine the results.

Finally, estimating these effects via linear models is not an accurate representation of the true system. For various reasons, we should not expect all of the observed relationships to function linearly. For example, the relationship between early-career PIPM and the jump a player makes from their first two years into their fourth year is likely nonlinear, with some players (e.g. Chris Paul) being so good early that their space for improvement is smaller, while an opposite subset of players are so bad early that they lack the ability or opportunity to learn that they also have limited space for improvement. The strength of the relationship should rise and fall as one moves across this spectrum. That being said, results included in Supplementary Tables 1 and 2 show that generalized additive models also find significant linear effects of early-career minutes played on fourth-year PIPM and the change in PIPM from a player's first two seasons to their fourth.

Given that the limitations to this analysis are limited, it is reasonable to consider the implications of these results for organizational decision making. Teams heavily invested in short-term success likely cannot afford to allocate substantial playing time to new players who have not adjusted to the level of NBA competition. However, my evidence suggests that teams whose championships aspirations extend out at least four years should heavily consider giving modest additional playing time to young players over established veterans on short-term contracts. In doing so, those teams will enable themselves to make more informed decisions about to whom they should extend qualifying offers, while at the same time increasing the quality of their best assets. If they choose to keep those players, the players' performances should be better for the additional time, and if the teams, instead, choose to trade them, those players should command a higher return in a trade. Additionally, teams with at least medium-term championship aspirations that have reached a steady level of success and roster stability, preventing them from accruing high draft picks to use on the highest-rated prospects, likely would benefit from placing increased value on potential draft picks' perceptual attunement, their ability to isolate the most important sensory information for athletic skill acquisition (Jacobs & Michaels, 2007). Draft picks left to the mid-first-round and later in the draft who are high in perceptual attunement are most likely to benefit from team-controlled opportunities for skill growth and performance improvement.

At the same time, my results suggest that there is likely some optimal number of young players to roster at each position. Young players who do not receive substantial playing time eat up roster spots and force organizations to make tough decisions when choosing whether to pick up their team options or extend them a qualifying offer. Young players who cannot find playing time are likely to return less in a trade, both because their performance level is stagnating or because teams will lack sufficient information about their abilities to assess the risk and reward in acquiring them properly.

Going forward, more research is necessary to help contextualize and expand these results. First, teams and researchers should assess the reliability and replicability of these results. For example, we should ask does playing time matter in other sports? If not, is there something specific to basketball that would suggest that this result is believable? Second, future research should also attempt to contextualize the observed causal effects. Is all playing time equal? Or, does playing under particular conditions elicit a positive effect, while playing in other situations, such as at the end of blowouts when neither team is giving a full effort, does not? Similarly, does playing in consistent lineups facilitate skill acquisition? Or, can young players improve regardless of the stability of those with whom they are playing? Finally, future investigators should use player tracking data to identify cases in which young players face similar on-court scenarios and evaluate whether a positive outcome increases with the number of times the players encounter such a scenario.

References

Bresler, A. (2021). *nbastatr.* Retrieved from http://asbcllc.com/nbastatR/index.html

Cinelli, C., & Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *82*(1), 39–67.

Davids, K., Araújo, D., Vilar, L., Renshaw, I., & Pinder, R. (2013). An ecological dynamics approach to skill acquisition: Implications for development of talent in sport. *Talent Development and Excellence*, *5*(1), 21–34.

Fong, C., Hazlett, C., Imai, K., et al. (2018). Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *The Annals of Applied Statistics*, *12*(1), 156–177.

Fong, C., Ratkovic, M., Imai, K., Hazlett, C., Yang, X., & Peng, S. (2019, August). *Cbps: Covariate balancing propensity score.* Retrieved from https://cran.r-project.org/package=CBPS

Gelman, A. (2008). Scaling regression inputs by dividing by two standard deviations. *Statistics in medicine*, *27*(15), 2865–2873.

Goldstein, J. (2018, January). *Nylon calculus: Introducing player impact plus-minus.* Retrieved from https://fansided.com/2018/01/11/nylon-calculus-introducing-player-impact-plus-minus/

Hogarth, R. M. (2001). *Educating intuition*. University of Chicago Press.

Hogarth, R. M., Lejarraga, T., & Soyer, E. (2015). The two settings of kind and wicked learning environments. *Current Directions in Psychological Science*, *24*(5), 379–385.

Jacobs, D. M., & Michaels, C. F. (2007). Direct learning. *Ecological psychology*, *19*(4), 321–349.

Kiverstein, J., & Rietveld, E. (2015). The primacy of skilled intentionality: on hutto & satne's the natural origins of content. *Philosophia*, *43*(3), 701–721.

Morgan, S. L., & Winship, C. (2015). *Counterfactuals and causal inference*. Cambridge University Press.

Staw, B. M., & Hoang, H. (1995). Sunk costs in the nba: Why draft order affects playing time and survival in professional basketball. *Administrative Science Quarterly*, 474–494.

Supplementary Information: The Causal Effects of Early-Career Playing Time on the Fourth-Year Performance of NBA Players *

Sean Fischer *The Annenberg School for Communication*

Potential Non-Linear Effects

To assess the possibility of some nonlinear relationships between both outcomes, fourth-year PIPM and the change in PIPM from a player's first two years and their fourth year, and the continuous treatment and other covariates, I re-estimated my models using generalized additive models. Additive models provide the benefit of using smoothing functions to estimate nonlinear relationships between the independent variables and the outcome of interest. I use the same covariate balancing generalized propensity score weights as in the main text in both cases.

Table 1 shows the results from the generalized additive model regressing fourth-year PIPM on the set of covariates. Notably, we see that the expected degrees of freedom of the smoothed early-career minutes term are 1. By contrast, the early-career PIPM term has 4.715 expected degrees of freedom, indicating a highly nonlinear relationship.

Table 2 displays the model results for the generalized additive model regressing the change in PIPM from a player's first two seasons to their fourth on the set of predictors. Again, we see that the early-career-minutes term is significant with an expected degrees of freedom of 1.

*Last updated January 22, 2021.

A. parametric coefficients	Estimate	Std. Error	t-value	p-value
Intercept	0.0587	0.1180	0.4978	0.6191
RSCI Ranked	0.0236	0.0604	0.3916	0.6957
Years in NCAA: 2	-0.0007	0.0767	-0.0095	0.9924
Years in NCAA: 3	-0.1296	0.0846	-1.5321	0.1267
Years in NCAA: 4+	-0.2129	0.0879	-2.4219	0.0161
Draft Class: 2001	-0.1386	0.1373	-1.0095	0.3137
Draft Class: 2002	-0.2030	0.1304	-1.5568	0.1207
Draft Class: 2003	0.0845	0.1288	0.6564	0.5122
Draft Class: 2004	0.0927	0.1402	0.6615	0.5089
Draft Class: 2005	0.0637	0.1320	0.4827	0.6297
Draft Class: 2006	-0.0556	0.1209	-0.4596	0.6462
Draft Class: 2007	-0.1616	0.1288	-1.2551	0.2106
Draft Class: 2008	-0.1115	0.1224	-0.9104	0.3634
Draft Class: 2009	-0.1117	0.1193	-0.9369	0.3497
Draft Class: 2010	-0.0267	0.1154	-0.2312	0.8174
Draft Class: 2011	0.1849	0.1206	1.5326	0.1266
Draft Class: 2012	0.0276	0.1161	0.2376	0.8124
Draft Class: 2013	-0.1456	0.1247	-1.1676	0.2440
Draft Class: 2014	-0.1108	0.1300	-0.8526	0.3947
Draft Round: 2nd	0.0233	0.1008	0.2309	0.8176
B. smooth terms	edf	Ref.df	F-value	p-value
Early Career Minutes Played	1.0000	1.0000	12.3271	0.0005
Last NCAA Season: Strength of Schedule	1.0000	1.0000	0.1452	0.7035
Last NCAA Season: Minutes Played	1.0000	1.0000	1.4676	0.2268
Last NCAA Season: TS%	1.0000	1.0000	6.7908	0.0097
Last NCAA Season: EFG%	1.0000	1.0000	4.1580	0.0424
Last NCAA Season: WS/40	1.0000	1.0000	10.4904	0.0014
Draft Pick Number	1.0000	1.0000	0.5248	0.4694
College BMI	5.5799	6.7952	2.2416	0.0428
Early Career PIPM	4.7145	5.7581	10.5649	< 0.0001

Table 1: Results from generalized additive model regressing fourth-year PIPM on early-career minutes and other covariates. All continuous variables have smooth functions applied. Propensity score weights applied. All continuous measures centered and standardized.

A. parametric coefficients	Estimate	Std. Error	t-value	p-value
Intercept	0.0633	0.1272	0.4978	0.6191
RSCI Ranked	0.0255	0.0651	0.3916	0.6957
Years in NCAA: 2	-0.0008	0.0827	-0.0095	0.9924
Years in NCAA: 3	-0.1397	0.0912	-1.5321	0.1267
Years in NCAA: 4+	-0.2295	0.0948	-2.4219	0.0161
Draft Class: 2001	-0.1494	0.1480	-1.0095	0.3137
Draft Class: 2002	-0.2188	0.1406	-1.5568	0.1207
Draft Class: 2003	0.0911	0.1388	0.6564	0.5122
Draft Class: 2004	0.1000	0.1511	0.6615	0.5089
Draft Class: 2005	0.0687	0.1423	0.4827	0.6297
Draft Class: 2006	-0.0599	0.1303	-0.4596	0.6462
Draft Class: 2007	-0.1742	0.1388	-1.2551	0.2106
Draft Class: 2008	-0.1202	0.1320	-0.9104	0.3634
Draft Class: 2009	-0.1204	0.1286	-0.9369	0.3497
Draft Class: 2010	-0.0287	0.1244	-0.2312	0.8174
Draft Class: 2011	0.1993	0.1300	1.5326	0.1266
Draft Class: 2012	0.0297	0.1252	0.2376	0.8124
Draft Class: 2013	-0.1570	0.1344	-1.1676	0.2440
Draft Class: 2014	-0.1195	0.1401	-0.8526	0.3947
Draft Round: 2nd	0.0251	0.1086	0.2309	0.8176
B. smooth terms	edf	Ref.df	F-value	p-value
Early Career Minutes Played	1.0000	1.0000	12.3271	0.0005
Last NCAA Season: Strength of Schedule	1.0000	1.0000	0.1452	0.7035
Last NCAA Season: Minutes Played	1.0000	1.0000	1.4676	0.2268
Last NCAA Season: TS%	1.0000	1.0000	6.7908	0.0097
Last NCAA Season: EFG%	1.0000	1.0000	4.1580	0.0424
Last NCAA Season: WS/40	1.0000	1.0000	10.4904	0.0014
Draft Pick Number	1.0000	1.0000	0.5248	0.4694
College BMI	5.5799	6.7952	2.2416	0.0428
Early Career PIPM	4.7145	5.7581	13.8393	< 0.0001

Table 2: Results from generalized additive model regressing change in PIPM from first two years to fourth year on early-career minutes and other covariates. All continuous variables have smooth functions applied. Propensity score weights applied. All continuous measures centered and standardized.