

Productivity Spillovers in Team Production: Evidence from Professional Basketball*

Peter Arcidiacono
Duke University
& NBER

Josh Kinsler
University of Georgia

Joseph Price
Brigham Young University
& NBER

August 25, 2015

Abstract

We estimate a model where workers are heterogeneous both in their own productivity and in their ability to facilitate the productivity of others. We use data from professional basketball to measure the importance of peers in productivity because we have clear measures of output, and members of a worker's group change on a regular basis. Our empirical results highlight that productivity spillovers play an important role in team production. Despite this, we find that worker compensation is largely determined by own productivity with little weight given to productivity spillovers.

*We thank Patrick Coate, Fabian Lange, Lars Lefgren, Craig Palsson, Michael Ransom, and seminar participants at Iowa State, McGill, Boston College, Brigham Young University, and the University of Georgia for helpful comments.

1 Introduction

The classic economic model predicts that workers will be paid the value of their marginal product of labor. Estimating this marginal product may be complicated by team environments in which workers contribute to team production directly but also indirectly through their effect on the productivity of other team members. If firms are able to identify workers who boost peer productivity, they can leverage complementarities in team production through team and task assignments. Workers who bring out the best in others will likely be assigned to tasks essential for firm production.

Mas and Moretti (2009) provide an excellent example of spillovers in team production by looking at the placement of cashiers in a supermarket. Placing the most productive cashiers in full view of the other cashiers resulted in the other cashiers working faster. However, Mas and Moretti provide one of the few examples where actual productivity is observed. Other examples include Hamilton, Nickerson, and Owan (2003), who examine worker interactions in the garment industry, and a set of papers analyzing productivity in the academy, Azoulay, Zivin, and Wang (2010) and Waldinger (2010, 2012).¹

The assumption made in this literature—as well as the abundant literature on peer effects in education—is that the individuals who are most productive themselves are also the ones who will make others most productive. This assumption may not be true in many contexts. For example, there are professors who choose to focus exclusively on their own research, providing little in terms of public goods while other professors who are particularly adept at helping their colleagues in their research and may do so even at the expense of their own research. Similarly, a brilliant but introverted student may not be as helpful to the learning of the other students as the perhaps not-so-brilliant student who asks good questions in class.

We may expect workers to be compensated for both their productivities and their abilities to make others more productive. However, peer effects, particularly heterogeneous peer effects, are notoriously difficult to measure, in part because of the data requirements. But the advent of “Big Data”, as well as the accompanying means of estimating models with large state spaces, may result in measures of productivity spillovers becoming more readily available. For example, patent scientists and financial advisors are both occupations for which there is rich data on individual output and information about network structure within the firm. These are both settings where a firm could identify workers who improve the productivity of their peers through interaction and advice.

¹Field experiments have also been used to examine peer effects in the workplace. See the papers by Bandiera, Baranakay, and Rasul (2009, 2010) as well as Falk and Ichino (2006).

In this paper, we develop and estimate a model of team production where individuals are heterogeneous in their own productivity as well as in their ability to help others be productive. We then relate our measures to compensation, focusing in particular on the extent to which the ability to make others more productive is rewarded in the marketplace. We focus on an industry, professional basketball, where the ability to help others is clearly an important part of team production. Sports data provide an excellent opportunity to study team production because the members of a team can be clearly identified, there are frequent changes in the players that compose a particular team, and compensation data is available.

Two papers using sports data highlight the heterogeneity in how spillovers may operate. Gould and Winter (2009) use data on baseball players to analyze how batter performance is related to the performance of other batters on the team. This paper fits perfectly with the idea that the most productive players have the largest positive peer effects: batting in front of a high-performing player results in receiving better pitches because the pitcher will not want to risk a walk prior to facing the high-performing player. Guryan, Kroft, and Notowidigdo (2009) examine how the productivity of one's golf partners affects own performance, finding no significant effects from being paired with better golfers. Even though there is no team production in golf, it may be that individual spillovers are multidimensional in the sense that they work through multiple player attributes. Certain players may be very productive but are surly or disobey common golf etiquette, both of which may serve to distract their partners. The authors allow for this additional flexibility, but find no evidence of heterogeneous spillovers.

Using possession-level data from games played in the National Basketball Association (NBA), we demonstrate that productivity spillovers play an important role in team production. We find that a standard deviation increase in the spillover effect of one player improves team success by 63% as much as a standard deviation increase in the direct productivity of that player. Estimates of the model also allow us to form player rankings based on the overall contribution to team production. We compare these rankings to estimates of team production when spillovers are ignored. Players who are generally perceived by the public as selfish see their rankings fall once we account for productivity spillovers.

We also use our model to highlight how the value of a particular player can vary depending on the composition of his teammates. Most firms have various teams within their organization and have the ability to reassign workers across teams. Since individual productivity and productivity spillovers play a complementary role, the overall contribution of a player will depend on the composition of the

other players already on the team.² We find that the assignment that produces the greatest increase in team productivity is often an assignment that does not maximize the direct productivity of the player. This suggests a tension that firms need to balance between team and player productivity, especially in firms where individual productivity has a large effect on compensation.

Given the large role spillovers play in team production in this industry, we would expect significant returns in the labor market to the ability to help others. This is not the case. Returns to own productivity are substantially higher than returns to the ability to help others, well beyond their differences in their contribution to team production. Part of the reason for this may be the difficulty in measuring the ability to help others. As in the academy, direct productivity is easily observed in ways that facilitating the productivity of others is not. To the extent that own productivity and facilitating the productivity of others is endogenous, the lack of returns to the latter may result in inefficient effort allocations among workers.

2 Data

To estimate a model of player performance, we use publicly available NBA play-by-play data covering all games during the 2006-2009 regular seasons gathered from espn.com. For those readers unfamiliar with the basic rules of basketball we have included a brief description in Appendix A. The raw play-by-play data provides a detailed account of all the decisive actions in a game, such as shots, turnovers, fouls, rebounds, and substitutions. Plays are team specific, meaning that there is a separate log for the home team and the away team. Associated with each play are the player(s) involved, the time the play occurred, and the current score of the game. While our model of player productivity is estimated using only the play-by-play data, we augment it with additional biographical and statistical information about each player gathered from various websites which we discuss later in this section. As described in Appendix B, we took a number of steps to clean the data. These included establishing which players were on the court, acquiring the outcomes of possessions, and matching the names of the players to data on their observed characteristics such as position and experience.

Table 1 describes our estimation sample in further detail. We use data from 905,378 possessions and 656 unique players active in the NBA from 2006-2009. On average, each player is part of 13,801 possessions, split evenly between offense and defense. The average number of possessions for

²Similarly, Ichniowski, Shaw, and Prennushi (1997) find that the returns to innovative work practices (e.g. teams, incentive pay, etc.) are complementary in the steel finishing industry.

each player-team-season combination is 4,507. The corresponding 25th and 75th percentile values are 1,130 and 7,470. The final four rows of Table 1 describe the typical outcomes for a possession. Slightly more than 50% of the time the offensive team scores, and, conditional on scoring, the offense scores on average 2.1 points.

To supplement the play-by-play data, we merge in biographical and statistical information about each player. Our primary source for this information is basketball-reference.com. The website contains basic player information such as date of birth, height, position, and college attended and a full set of statistics for each season the player is active. We also gathered information on salaries and contract years from prosportstransactions.com and storytellerscontracts.com. We also obtained additional measures of player performance from basketballvalue.com and 82games.com.

3 Model and Estimation

In this section we present a model of team production, discuss identification, and describe our estimation strategy. The innovation of the model is that the ability of an individual to influence the productivity of others is not directly tied to own productivity. We tailor the model to the NBA context, though it would be simple to expand the framework to other types of production.³ The number of parameters to be estimated is quite large and would be computationally prohibitive using straight maximum likelihood. Consequently, we take an iterative approach as in Arcidiacono et al. (2012).⁴

3.1 Model setup

Our unit of analysis is an offensive possession during an NBA game. There are five offensive and five defensive players on the court during every possession. For a given possession n , denote the set of players on the court as P_n , where P_n includes the offensive players on the court O_n and the defensive players on the court D_n . For notational ease we abstract from the fact that possessions are typically observed within games, which themselves are observed within seasons. Additionally we abstract from the concept of team, even though the potential sets of offensive and defensive players will be determined by team rosters. A possession can end in one of six ways, no score or one of the

³For example, in Mas and Moretti (2009) checkout cashiers are assumed to influence other cashiers through their own productivity. However, it would be straightforward to allow for completely separate effects.

⁴See Burke and Sass (2013) for an application of this method in education and Cornelissen, Dustmann, and Schonberg (2013) for an application in the labor market.

five offensive players in O_n scoring at least one point.⁵ We assume that each player i on the court is fully characterized by three parameters: (i) their ability to score, o_i , (ii) their ability to help others score, s_i , and (iii) their ability to stop others from scoring d_i .

Assume for the moment that there is no heterogeneity in defensive skills. The likelihood that offensive player i scores to conclude a possession will depend on i 's own ability to score and his ability to help others score, as well as the similar skills of his teammates on the court. Denote $y_{in} = 1$ if the individual scores and $y_{in} = 0$ otherwise. We assume that the probability that player i scores at least one point during possession n is given by⁶

$$Pr(y_{in} = 1|P_n) = \frac{\exp\left(o_i(1 - \sum_{k \in O_n, k \neq i} s_k)\right)}{1 + \sum_{h \in O_n} \exp\left(o_h(1 - \sum_{k \in O_n, k \neq h} s_k)\right)}. \quad (1)$$

The probability that player i scores to end possession n is increasing in o_i , the offensive intercept of player i . An increase in the offensive spillover of player $k \neq i$ will have an ambiguous effect on the probability that player i scores, since an increase in s_k also benefits the other offensive players in O_n .⁷ The probability that possession n ends with no points scored is simply $1 - \sum_{i \in O_n} Pr(y_{in} = 1|P_n)$.

The above model is inadequate since defenders will vary in ability. Thus, the probability that one of the players in O_n scores will depend on the composition of the players in D_n . To account for defense, we alter the above framework such that the probability that player i scores at least one point during possession n is given by

$$Pr(y_{in} = 1|P_n) = \frac{\exp\left(o_i(1 - \sum_{k \in O_n, k \neq i} s_k)\right)}{\exp\left(\sum_{j \in D_n} d_j\right) + \sum_{h \in O_n} \exp\left(o_h(1 - \sum_{k \in O_n, k \neq h} s_k)\right)} \quad (2)$$

The difference between Equation (1) and Equation (2) is that the index associated with no points being scored now varies with the abilities of the defenders in D_n . The joint defensive prowess of the

⁵Note that this is different from many other labor market settings where production for each person may be observed. Individual productivity measures would be equivalent to having many players be able to score on one possession, a feature that would make estimation easier as we would have output measures on everyone at each point in time. This is discussed in greater detail at the end of this section.

⁶Separate identification of own productivity from spillover productivity is not reliant on the multinomial logit production function. We choose this specification since possession outcomes are binary in nature. In fact, a linear model would be significantly easier to estimate.

⁷Note also that spillovers are not specific to particular player combinations, with certain players working well together beyond what their individual abilities would suggest. We maintain this assumption as well for defensive parameters. Substitutions are then exogenous conditional on the fixed effects (offensive intercept, defensive intercept, and spillovers) of the players.

players in D_n is a linear function of the defensive intercepts of each player j .⁸

Possessions that yield positive points do not necessarily contribute equally to team success since there are a range of plausible point outcomes. To determine the expected number of points per possession for each player, we scale the probability of scoring positive points by the expected number of points conditional on scoring for each player,

$$E[Points_{in}|P_n] = E[Points_i|y_i = 1] \times Pr(y_{in} = 1|P_n). \quad (3)$$

The above formulation assumes that the number of points a player scores, conditional on scoring, is unrelated to the identities of the other players on the court. We check this assumption by regressing points scored on player fixed effects and teammate fixed effects for all possessions that yield positive points. In this simple model, each player has two parameters, one describing his own typical point production and the other describing how he influences the points scored by his teammates when he is on the court. For each season in our data, we fail to reject that all of the spillover related parameters in point production are jointly equal to zero. Thus, it appears that our assumption is a reasonable approximation of the data even if it is not strictly true. The advantage of making this simplifying assumption is that we can estimate $E[Points_i|y_i = 1]$ outside of the model. Without the independence assumption we would need to append points scored to our likelihood function and in the process triple the number of parameters. Each player would be described by an offensive intercept, offensive spillover, and defensive intercept both for the probability of scoring and points scored.⁹ This would significantly impede estimation.

Our model of team production is complex; the production function is non-linear, there are a large number of parameters to estimate, and many of these parameters interact. The large parameter space reflects our desire to estimate both own and spillover productivity. To estimate these parameters we use a non-linear production function since only one player can score during a single offensive possession. The fact that only one player can score per possession means that we don't observe the output, or contribution, of all the players all the time. We match this feature of the data by using a multinomial logit type production technology that generates discreteness in individual outcomes. If output were available for each individual in each time frame, a much simpler model of own and

⁸Note that defensive spillovers are not present. We experimented with including defensive spillovers and found that they were poorly identified.

⁹There may be ways to restrict the parameters for each player across the two outcomes and thus limit the number of additional parameters, but it is not obvious a priori what those restrictions should be. The simpler approach, and the one supported by the data, is to assume points conditional on scoring is independent of the other players on the court.

spillover productivity might be possible. For example, Mas and Moretti (2009) observe productivity for all cashiers working a particular shift. It is possible in their context to estimate own productivity and the ability to help others be more productive using a linear specification. Estimation of such a linear model would require iteration, but would be easy to implement since it would essentially require running repeated OLS regressions as in Arcidiacono et al. (2012). Linear specifications could also be utilized to explore spillovers in education (test scores are available each year), health (doctors and nurses treat patients each period), or entertainment (actors perform in many movies).

An alternative approach for estimating individual contributions to team performance, when output measures for all team members are lacking, is to model team outcomes directly. Because team outcomes are always available, it is possible to estimate player contributions using a linear specification. The Adjusted Plus/Minus (APM) statistical model of player performance follows this approach.¹⁰ The basic idea is to regress a team outcome, such as points scored, on a set of offensive and defensive player fixed effects. The drawback to this methodology is an inability to separately identify own and spillover productivities. The estimated offensive fixed effect for each player will be a mix of own productivity and the ability to help others be more productive. Replacing player fixed effects in the APM with an aggregate of traditional box score statistics, such as points, assists, and rebounds faces the same criticism. The coefficients on these statistical measures will capture a mix of own and spillover productivities. In contrast, our model exploits information about both the players on the court and which of those players scores to separately identify the two parameters. Separate identification of own and spillover productivities is critical for evaluating counterfactual player groupings when there are potential complementarities in production.

3.2 Normalizations and identification

The offensive intercept parameters can be readily identified if data for many possessions is available. However, to identify the offensive spillovers it is necessary to observe player i in multiple groups. Put differently, the spillover parameters can only be separately identified if we observe player i with different sets of teammates across possessions. If not we could simply redefine o_i such that $o_i^* = o_i(1 - \sum_{k \in O_n, k \neq i} s_k)$ for all i and estimate the o_i^* 's. Without switching there is no way to separate the offensive intercepts from the offensive spillovers.

Even when the set of teammates in O_n varies across possessions, it is still necessary to make a normalization on the spillover parameters. For any set of o_i and s_i , it is always possible to redefine

¹⁰Further details on the APM are available at <http://www.82games.com/barzilai2.htm>

the parameters such that the predicted probabilities are identical. For example, define

$$o_i^* = o_i(1 + 4\bar{s}) \tag{4}$$

where \bar{s} is the mean spillover calculated using i 's spillover and the spillover of all the players i is ever grouped with. Additionally, define

$$s_i^* = \frac{s_i - \bar{s}}{1 + 4\bar{s}} \tag{5}$$

It can be shown that the predicted probability that i scores is identical across $\{o, s\}$ and $\{o^*, s^*\}$. A natural normalization is simply to constrain $\bar{s} = 0$.

Adding defensive heterogeneity to the scoring probabilities does not alter the identification argument regarding the offensive parameters. Similar to the offensive spillover parameters, it is not possible to identify the baseline defensive productivities unless player i is observed on defense with different sets of teammates. This is because defensive outcomes are essentially group outcomes. Thus, switching teammates is critical for identifying the defensive intercepts. Note that the mean defensive intercepts must also be normalized since the mean is not separately identified from the scale of the baseline offensive parameters. The normalization we make is that the possession-weighted average of the defensive intercepts across players is zero in each year.

The normalization and identification issues discussed above are technical concerns in the sense that it is not possible to identify all the parameters of the model even if infinite amounts of data were available. However, in order to consistently estimate the player productivities that are identified, we must make an additional exogeneity assumption, namely that changes in the sets of players in O_n and D_n are exogenous conditional on o_i , s_i , and d_i . In the context of NBA basketball, there are many sources of plausibly exogenous player substitutions, such as fatigue, injury, or foul trouble. However, there is also a concern that some player substitutions are strategic responses to idiosyncratic shocks to player performance, game situations, or individual match-ups between players in O_n and D_n . Strategic substitutions could bias our estimates of player productivities in unpredictable ways. We investigate the robustness of our findings to alternative modeling assumptions in Section 4.5.

3.3 Iterative algorithm

There are three parameters to be estimated for each player we observe in the data. As previously discussed, it is not possible to estimate the offensive spillovers without an additional restriction. If all the players in the sample are connected, in other words, every player can be linked to every other player through their teammates, then it would only be necessary to restrict the overall mean of the

offensive spillovers to zero. In our data, all the players are in fact connected, since players switch teams both within and across seasons. However, jointly estimating all of the player parameters by maximum likelihood imposing this one restriction is not computationally feasible.

As a result, we pursue an estimation strategy that treats each team-season as an independent entity. Players that switch teams within a season are treated as completely unrelated. For each team-season, we normalize the possession-weighted spillover to zero. Once we have estimated all the player parameters imposing these restrictions, we adjust team and player spillovers to be consistent with observed changes in the same player’s performance across different teams and seasons. This approach has two advantages. First, it facilitates estimation since we can iteratively estimate the offensive intercepts, offensive spillovers, and defensive intercepts team by team. Second, it allows player productivity to vary across seasons as a result of random factors such as health and luck.

Once we have defined players as any unique combination of player-season-team, we estimate the offensive and defensive parameters using an iterative approach. The method has three broad steps that correspond to estimating the offensive intercepts, defensive intercepts, and the offensive spillovers, where each step increases the log likelihood.¹¹ The following steps outline the estimation procedure more precisely, where now player i ’s parameters are indexed by team-season pairs (t).

- Step 0: Make an initial guess of the parameters, denoted by $\{o_{it}^0, s_{it}^0, d_{it}^0\}$.
- Step 1: Estimate by maximum likelihood o_{it}^1 conditional on $\{s_{it}^0, d_{it}^0\}$.
- Step 2: Estimate by maximum likelihood d_{it}^1 conditional on $\{o_{it}^1, s_{it}^0\}$.
- Step 3: Estimate by maximum likelihood s_{it}^1 conditional on $\{o_{it}^1, d_{it}^1\}$.

Estimation proceeds by iterating on Steps 1-3 until convergence, where the estimates at the q th iteration are characterized as $\{o_{it}^q, s_{it}^q, d_{it}^q\}$. Note that within Steps 1-3, estimation proceeds separately for each team-season, implying that each maximization step is searching over only approximately 15 parameters—one for each player who played on a particular team in a particular season. Our approach relies on the fact that there are no across-team interactions to be concerned about when the remaining parameters are taken as fixed—when estimating the offensive parameters for one team, all the defensive parameters of the players associated with the other teams are taken as given. Note

¹¹Increasing the log likelihood at each step does not guarantee the algorithm will converge to the global maximum. We used two different sets of starting values, one of which used the estimates from the model where we did not allow for spillovers and then set the starting values for the spillover parameter to zero. In both cases the estimates converged to the same values.

that this approach would not be possible if the parameters for a particular player were constrained to be the same across seasons or across teams within a season.

3.4 Linking across years

Once the above process converges, we are left with a set of parameters describing player performance for each team-season. In order to compare these parameters across team and season combinations, assumptions must be made. For the defensive parameters, we assume that the possession-weighted average defensive ability is the same in each season.¹² Since our normalization on the defensive parameters was such that the average defensive ability was zero in each season, no adjustments need to be made to compare defensive parameters across seasons. Note that coaches are embedded in these defensive parameters: certain coaches may positively affect the defensive abilities of their teams. Hence we can project the time-varying defensive parameters on, for example, player fixed effects and coach fixed effects in order to ascertain which coaches are particularly effective at improving the defensive skills of their players.¹³

More difficult is comparing offensive intercepts and spillover parameters across teams and seasons. Recall that the average spillover was normalized to zero for each team-season. This rules out the possibility that specific teams are composed of players with higher-than-average spillovers. To compare players across teams and seasons, we make an observationally-equivalent assumption, but one that is perhaps more tenable. Namely, we assume that the variation in a particular player's spillover parameter over time does not depend on team-specific effects in that season. Hence, coaches or teams may be particularly effective in generating offense in a particular year beyond what would be expected given the skills of their players, but this effect operates solely through the offensive intercept. Similar to the discussion of the defensive parameters, we can project the offensive intercept on player fixed effects and coach effects, but our normalization effectively rules out coach effects on the spillover parameter.

Now denoting s_{it}^0 as the initial estimate of the spillover parameter when the average spillover on each team in each season is constrained to zero, the procedure to change the spillover normalization

¹²Note that the defensive abilities come from both the players skills and the coaching strategies. Hence, the assumption is that, on average, the combination of the defensive skills of the players in the league as well as how coaches complement those skills does not change over time.

¹³In fact, we can go further than this by interacting the coach fixed effects with the positions of the players. For example, some coaches' defensive schemes may depend more on the skills of particular positions than the schemes of other coaches. We do not pursue this because the panel is relatively short.

iterates on the following two steps:

- Step 1: Regress s_{it}^0 on player fixed effects and team-season fixed effects, weighting each observation by the observed number of possessions.¹⁴ This leads to the following decomposition of s_{it}^0 :

$$s_{it}^0 = s_i^0 - \delta_t^0 + \epsilon_{it}^0 \quad (6)$$

where s_i^0 and δ_t^0 give the contribution of the player and the team-season respectively. The team season effect enters negatively in the decomposition since a higher δ_0 implies a better spillover team and thus a lower normalized spillover estimate.

- Step 2: Create adjusted offensive spillovers using: $s_{it}^1 = (1 + 4\delta_t^0)s_{it}^0 + \delta_t^0$. Note that this is essentially equation (5) which gives the mapping from one normalization to another, where s_i^* is s_{it}^0 , \bar{s} is δ_t^0 , and s_i is s_{it}^1 .

Estimation proceeds by iterating on Steps 1-2 until convergence, where the estimates at the q th iteration are denoted by s_{it}^q . Convergence occurs when the regression of s_{it}^q on player fixed effects and team-season effects yields team-season effects that are all zero, $\delta_t^q = 0$, implying $s_{it}^q = s_i^q + \epsilon_{it}^q$.¹⁵ After convergence, the offensive intercepts are adjusted according to equation (4) using the new spillover parameters.

To get a more intuitive sense for how the algorithm works, consider a player who is on team A in the first half of the season and on team B in the second half of the season. Suppose team A has higher spillover players than team B . Since the mean spillover under the initial normalization is zero for both teams A and B , this player will have a lower estimated spillover parameter when he played for team A as he was pooled with stronger players. Hence the regression of the player's spillover parameters on his fixed effect and the negative of the team-season effect will yield a positive estimate for team A 's fixed effect. This positive effect gives us our first estimate of how much better the players on team A in that season were on the basis of the spillover parameter which is then used to adjust the spillover parameters for the players on team A upward.

¹⁴One team-season effect must be normalized to zero to pin down the scale of the offensive intercept parameters. It is possible to control for experience or age when estimating the spillover regression. Our results do not change if we allow for these additional explanatory variables.

¹⁵Only a few iterations are required in practice.

4 Results

The estimation procedure yields three parameters for each player-season-team combination. If these parameters are capturing something permanent about player skill, they should be reasonably stable over time and across teams. To investigate this, we estimate separate fixed effects regressions for each of our skill measures. The outcome variables are the player-season-team estimates and the only explanatory variables are player fixed effects. Each player-season-team skill observation is weighted by the number of observed possessions. With these simple regressions we are able to explain approximately 83%, 50%, and 57% of the variation in the offensive intercept, offensive spillover, and defensive parameters respectively across seasons and teams. Thus, it does appear that we are capturing something intrinsic about each player.¹⁶

For evaluating players and teams, however, the parameter estimates themselves are not particularly informative. In the next few sections we demonstrate how the various player skills contribute to team performance, how to rank players using our estimates, and finally whether players and teams make decisions that are consistent with the results of our model.

4.1 Importance of the three factors

As noted above, the scale of the offensive intercept, offensive slope, and defensive intercept are not meaningful on their own. To illustrate how important each of these components are for team success, we perform the following exercise using player skill estimates from the 2009-2010 season. We first identify the four most utilized players on each team in the 2009-2010 season based on total possessions. We then ask how each group of four players would perform when various types of players are added. Our measure of team performance is the predicted per possession point differential against an average team.¹⁷

The results of this exercise are illustrated in Table 2. The first row of results shows the dis-

¹⁶Since the spillover parameter is a key contribution of our empirical model, we provide two pieces of external evidence about the validity of this measure. First, we find that the spillover parameter is strongly correlated with the statistics in box-score data that it should be most closely related to, such as assists and turnovers. Second, John Hollinger wrote an article for ESPN in 2010 that listed the 15 worst ball-hogs in the NBA. We find that these players have a spillover parameter that is 0.4 of a standard deviation lower than the other players in the league.

¹⁷The five offensive intercepts for the average team are chosen to match the average intercepts across teams by player offensive rank. For example, the offensive intercept for the most productive scorer on the average team is the mean across all 30 teams of the best scorer's offensive intercept. Each player on the average team is assigned the overall average spillover and defensive parameter since these enter the production function linearly.

tribution of predicted point differentials when an average player is added to each team’s top four utilized players.¹⁸ Across 30 teams, the average differential is slightly above zero, with a standard deviation equal to 0.087. Using this as a baseline we then explore how each team’s per possession point differentials change when players with particular skills are added. We consider six different player types, altering the offensive intercept, offensive slope, and defensive intercept by one standard deviation in either direction from the average.¹⁹

We find that all three factors are important for team performance. Adding a one standard deviation better offensive intercept, offensive spillover, or defensive intercept player improves a team’s per possession point differential by 0.027, 0.017, and 0.021 points respectively. Compared to the baseline standard deviation of point differentials of 0.087, these numbers indicate that adding a one standard deviation more skilled player increases a teams per possession point differential by 20% to 30% of a standard deviation. The largest change in team performance stems from the addition of a better offensive intercept player. This is true despite the fact that adding a good scorer necessarily decreases the opportunities that the other players on the team have to score.

The results in Table 2 also document the variability across teams in the effect that different types of offensive players have on team performance.²⁰ Because of the complementarities in the probability of scoring, the benefit of adding a particular type of player will vary by team. For example, a high offensive intercept player may be valued more by teams that have fewer high spillover players since a productive scorer doesn’t rely as much on his teammates to score. In contrast, a team with a number of productive scorers may prefer to bring in a high spillover teammate to enhance the productive skills already present.

¹⁸The average player is constructed by taking the possession weighted average of each of the offensive intercept, offensive spillover, and defensive intercept.

¹⁹To calculate the standard deviation for each skill type, we first compute the variance of our parameter estimates by skill type, weighting by each player’s total number of possessions. This variance is an overestimate of the true variation in player skill since it will be inflated by sampling error in our estimates. We correct for this by subtracting off the average variance of the sampling error across the parameter estimates, again weighting by the total number of possessions. The sampling variance for each parameter is simply the square of the standard error associated with each parameter, which we obtain by inverting a numerical approximation to the Hessian. Further details on calculating standard errors are available in Appendix C.

²⁰In contrast to the offensive skills, there is little variability across teams in point differential changes associated with adding a one standard deviation better (or worse) defensive player. This is primarily a reflection of the fact that we assume that there are no complementarities in defensive production.

4.2 Position comparisons

High offensive intercept, spillover, or defensive players are often associated with particular positions on a standard NBA team.²¹ For example, point guards are generally viewed as facilitators, while centers are expected to protect the basket on defense. Table 3 shows how the various skills break down by position. For each position, we show the average skill measure for each of our estimated parameters and a measure of a player's overall effectiveness (a combination of our three measures which we discuss further in the next section). For comparison purposes we also include two common measures of player effectiveness, player efficiency rating (PER) and adjusted plus minus (APM).²² All measures are standardized to have a mean of zero and standard deviation equal to one across positions.

The estimates of our model match the basic intuition about the types of skills different position players bring to a team. Point guards are by far the best spillover players but tend to be below average scorers and very poor defenders.²³ In contrast, centers are 0.68 standard deviations better than the average defensive player and the huge defensive benefit of centers provides their most important contribution from an overall effectiveness standpoint, with centers being 0.33 standard deviations more effective than the average player. The PER's ordering of overall effectiveness by position is similar to our ranking, except in the case of point guards which are ranked higher under

²¹The most common lineup in professional basketball contains a point guard, shooting guard, small forward, power forward, and center. However, teams face no restrictions regarding which positions players are allowed to play at one time.

²²PER is a rating of a player's per-minute productivity that is generated using a complicated formula based on box-score statistics. The precise formula can be found at <http://www.basketball-reference.com/about/per.html>. PER does not consider who each player plays with or against and is viewed largely as a measure of offensive effectiveness. PER has been criticized as a measure of player effectiveness since it emphasizes shot taking (see <http://wagesofwins.com/2006/11/17/a-comment-on-the-player-efficiency-rating/>). APM ratings indicate how many additional points are contributed to a team's scoring margin by a given player in comparison to the league-average player over the span of a typical game. APM is constructed using a fixed effects regression where the dependent variable is the per-possession point differential for a given set of players on the court and the explanatory variables are player fixed effects. Further details are available at <http://www.82games.com/barzilai2.htm>. The advantage of APM relative to PER is that it explicitly accounts for who a player plays with and against. However, because the model is linear, it cannot capture complementarities in production.

²³We find statistically significant correlations between the different skill parameters. A player's ability to score and ability to help others score have a possession-weighted correlation equal to -0.21 that is statistically significant at a 5% level. This does suggest a tradeoff between own and spillover productivity on average. Ability to score and ability to stop others from scoring are not significantly correlated. However, the ability to help others score and to stop others from scoring have a possession-weighted correlation of 0.05 that is statistically significant.

PER. This is consistent with the criticism that the PER fails to account for defensive contributions. The APM, on the other hand, ranks centers quite differently than either PER or our measure. In the next section we discuss further how to rank players individually.

4.3 Ranking workers overall contribution

The previous sections suggest that attempts to rank players individually in a team sport are misguided since the value of each player necessarily depends on who his teammates are. However, NBA player rankings are ubiquitous and often fail to account for the team nature of the sport. We develop a player ranking that directly accounts for the complementarities present in a team setting by using estimates of each player's underlying skills from our spillover model. We first describe how we construct our rankings, and then compare them to other common rankings and rankings we generate when ignoring spillovers in team production.

There are a number of ways to measure the effectiveness of each player given our estimated parameters. We construct our preferred measure by first taking each player and pairing him with an average team. A player's measured effectiveness is then the per possession point differential when this team plays against an average opponent.²⁴ The per possession point differentials are then standardized so players can be compared in standard deviation units. We create two additional measures of player effectiveness. The first takes our preferred measure and adjusts for position, since as Table 3 indicates there are significant differences in player effectiveness by position. Because teams typically field lineups with one player at each position, players who excel at underperforming positions will be valued more. Finally, rather than take each player and put him on an average team, our third measure of player effectiveness replaces each player with an average player and asks how his team's performance changes.²⁵ To some extent, this measure accords with how valuable each player is to their team.

Table 4 lists the top ten players in the 2009-2010 season according to our three rankings along with the rankings according to PER and APM.²⁶ Our preferred rankings indicate that Dwight

²⁴The average teammates a player is assigned and the average opponent are constructed using the average offensive intercepts across teams by player offensive rank. For example, the offensive intercept for the most productive scorer on the average team is the mean across all 30 teams of the best scorer's offensive intercept. Each player on the average team is assigned the overall average spillover and defensive parameter since these enter the production function linearly.

²⁵For this measure we create rankings only for the top five most utilized players on each team. This allows for a straightforward determination of who the teammates will be when each player is replaced with an average player.

²⁶We only consider players who accumulated at least 2,000 total possessions for any team in 2009-2010. This restriction limits the rankings to those players observed often enough to accurately estimate their underlying skills.

Howard is the most effective player, over three standard deviations better than the average NBA player. The primary reason that Dwight Howard is so highly ranked is that he is the top ranked defensive player, almost a full standard deviation better than the next best defender. The rest of the top ten is full of names that are familiar to basketball fans, but are not necessarily the brightest stars in the game. For example, Al Horford and Chris Andersen are highly ranked because they are well above average both offensively and defensively. Horford is an above average offensive intercept, offensive spillover, and defensive player. Andersen is actually a below average offensive intercept player, but his presence on the court generates enough extra opportunities for his teammates that his offensive spillover measure is two and a half standard deviations above the mean. The rankings based on APM also pick-up Andersen's overall effectiveness.

LeBron James, widely regarded as the best player in the game, is ranked number six according to our preferred method. This "low" ranking is a reflection of the fact that in 2009-2010 James is a good, but not great defender, and only an average spillover player. Based on offensive intercept alone, James would be the highest ranked player, meaning that when added to an average team James would have the highest scoring probability relative to adding any other player.²⁷ The PER measure is often criticized for over-valuing shooting and scoring and not surprisingly James comes out ahead on this measure.

When the rankings are adjusted for position or team there are slight changes. Because centers are on average the most effective players, Dwight Howard is de-valued when ranked relative to other centers and drops to the 4th best overall player. Point guards, shooting guards, and small forwards move up the ranks, with Kevin Durant now identified as the most effective player. The player rankings changed very little when players are assessed based on how their team would perform without them. Finally, many of the names on our preferred ranking list appear in the PER and APM rankings. In fact, the possession weighted correlation between our preferred measure of standardized point differential and PER is 0.42. The correlation with the APM rankings is significantly higher, equal to 0.78.²⁸ This is not surprising since our measure is more similar to APM since it measures

337 players, or 67% of active players in 2009-2010, played in 2,000 or more possessions. Each team has approximately 11 players who play more than 2,000 possessions. Among those who played at least 2,000 possession the average number of possession was 6,436. The mean for those who played fewer than 2,000 possessions is 741. Rankings for 2006-2007 through 2008-2009 can be compiled in a similar manner.

²⁷Note that a ranking based strictly on each player's offensive intercept would yield a top-five of LeBron James, Dwyane Wade, Kevin Durant, Kobe Bryant, and Carmelo Anthony. These are five of the most recognizable and renowned scorers in the NBA.

²⁸Again, only players with more than 2,000 possessions are considered.

a player’s effectiveness controlling for the identity of the other players on the court.

4.4 Ignoring spillovers

In Table 5, we examine how our player rankings change when estimating a model that ignores spillovers, essentially ruling out any complementarity in offensive production. The first column shows the top ten players based on point differentials when playing with average teammates against an average opponent. Many of the names remain the same, such as Dwight Howard and LeBron James, but there are significant changes. In particular, players that tend to score often and also play above average defense tend to move up in the rankings. Examples include Tim Duncan, Kobe Bryant, and Chris Bosh.

The second through fourth columns look more closely at some of the changes in standardized point differentials across the spillover and no spillover models. The second column lists the changes for some notable players. For example, Carmelo Anthony, a high volume shooter, is 1.31 standard deviations better in the no spillover model than in the spillover model.²⁹ In contrast, Steve Nash, a player widely believed to be one of the best offensive facilitators in the NBA, is 1.30 standard deviations better in the spillover model. Columns three and four of Table 5 lists the ten players who have the largest positive and negative swings in point differentials between the spillover and no spillover models. The players that tend to improve greatly when complementarities are modeled are pass-first point guards, such as Jason Williams and Jamaal Tinsley, and players who tend not to score but generate opportunities for their teammates through offensive rebounds, screens, and passing, such as Theo Ratliff, Chris Andersen, and Anderson Varejao. The list of the ten largest negative changes is full of players who are well known to be not only bad passers, but shoot-first players, like Chris Kaman and Carmelo Anthony.

While the results presented thus far coincide with widely held perceptions about individual players, there is still a concern that our estimates could be biased by endogenous changes to team composition. In the next section, we investigate the robustness of the estimated player productivities to two main sources of endogenous substitutions.

4.5 Robustness Checks

As discussed in Section 3.2, endogenous changes to team composition could arise in response to idiosyncratic shocks to player performance or changing game strategy as a function of the competitive

²⁹Anthony was second in shots per 36 minutes during the 2009-2010 NBA season.

environment. If players play more in games when they receive positive productivity shocks, then estimated productivity will be biased upwards. Alternatively, if players are asked to allocate effort differentially on offense or defense as a function of game situations, this can also affect our estimates of player productivity. In this section, we conduct two robustness checks. First, we see if allowing players to play more possessions when they are performing well affects our findings. Second, we see how our estimates change when only the first three quarters are used, as the fourth quarter may entail more or less pressure depending on the score of the game.

4.5.1 Endogenous playing time due to variation in player productivity across games

To investigate the concern that playing time in any particular game is endogenous to contemporaneous performance, we estimate separate productivity measures for each player according to how many possessions they play in a particular game. We re-estimate the model outlined in Section 3.1, but index players by two sets of parameters, $\{o_i^H, s_i^H, d_i^H\}$ and $\{o_i^L, s_i^L, d_i^L\}$. The first (second) set of parameters describes player productivity in games when they play more (fewer) than their median number of possessions per game.³⁰ Essentially we are treating each player as two distinct players according to how often they play in any particular game.

All the parameters are jointly estimated, meaning that there will be players on the same team playing above and below median possession games at the same time. Once all the parameters are estimated for each team-season, we link players using the same strategy discussed in Section 3.4. In order to compare the results to our baseline estimates we need one set of parameters for each player. We combine the productivity parameters from high and low possession games by taking a simple average, such that player i 's productivity is given by $\{\frac{o_i^L+o_i^H}{2}, \frac{s_i^L+s_i^H}{2}, \frac{d_i^L+d_i^H}{2}\}$. These productivity measures can be interpreted as expected productivity. At the start of any game there is a 50-50 chance a player plays as a high-possession player or as a low-possession player.

The top panel of Table 6 illustrates how expected productivity compares with our original productivity. Across seasons and skill, the estimates are very highly correlated. In particular, the offensive intercepts generated through the approach described above are almost perfectly correlated with our baseline estimates. The second three rows in the top panel also indicate that our player rankings are largely unaffected by endogenous substitutions as a function of player performance. The correlation between our baseline ranking and the ranking based on expected productivity is approximately 0.90. While the baseline findings are largely unaffected, the results of the high/low

³⁰Only games in which a player plays positive possession are included.

possession exercise are consistent with the existence of endogenous substitution. On average, a player's high possession productivities are larger than the low possession productivities. However, it is not as if the best players do not play if they are having an off night. To see this, consider that the average within player standard deviation in total possessions across high and low possession games is 712. In contrast, the across player standard deviation of total possessions is 4,731. Overall, the game-to-game variability in playing time for each player is small relative to the variability in playing time across players.

4.5.2 Endogenous playing time due to game situations

A second source of endogenous player substitutions that could bias our estimates of player productivity are changing game situations. Players may be instructed to play differently as a function of the competitive environment. If a particular player only plays under a certain set of game circumstances, then that player's productivity estimates are not reflective of his general skills. To investigate this concern, we eliminate from our data any possession that occurs after the 3rd quarter and re-estimate the model. The idea behind this approach is that game strategy is likely to change most significantly in the 4th quarter. As the game nears its end, coaches are more likely to take drastic measures. Eliminating data from the 4th quarter eliminates some of these concerns. Estimation follows the approach outlined in Section 3, the only difference is that our sample is approximately 25% smaller.

The bottom panel of Table 6 compares the baseline productivity estimates with the estimates generated after dropping all 4th quarter observations. Similar to the previous exercise, the offensive intercepts are highly correlated across the two models in all seasons, on the order of 0.95. The offensive slopes and defensive parameters are correlated at a slightly lower level, around 0.80. The lower correlation for these parameters could be partly the result of game strategy, but is also likely a result of the fact that we have thrown out a lot of useful variation. Both of these parameters are identified through changes in team composition and are less precisely estimated than the offensive intercepts. Thus, when we eliminate a quarter of our sample it is not surprising that these parameters are more sensitive. The last three rows of the bottom panel show that the sensitivity of the spillover and defensive parameters trickles over to our player rankings. Across the 2007-2010 seasons, the correlation between the original player rankings and the rankings without the 4th quarter is approximately 0.80.

Overall the results of the robustness exercises are encouraging. The player productivity measures and rankings change very little when we treat high possession games differently than low possession

games. Eliminating all 4th quarter observations results in slightly larger changes, but this is not surprising as we have excluded a quarter of our sample. In the next section we move from our individualistic approach to players and consider more directly how players interact to generate team success.

4.6 Allocating workers to the optimal team

For the purposes of ranking individual players we considered how each player performs with an average team. However, when actual player personnel decisions are made, success will hinge on how the various components of a team work together. This is true not only in the context of professional basketball, but in other industries as well. When a firm considers hiring a new worker it likely considers how that worker complements the skills of the workers already at the firm. Moreover, individual workers will seek out firms where their talents can best be utilized.

To illustrate how our model captures these ideas, we first examine one of the most high profile personnel decisions in the history of the NBA. In the summer following the 2009-2010 season, LeBron James' contract with the Cleveland Cavaliers expired and he became a free agent. Using our model, we evaluate James' decision, examining both his own performance and the likelihood of team success. Table 7 presents the model predictions for the teams most interested in signing James: Cleveland, Miami, Chicago, and New York. Cleveland provided the greatest opportunity for individual output, while Miami offered the greatest chance for team success. James' predicted per possession probability of scoring declines from 0.187 with Cleveland to 0.165 with Miami, a drop of 11.8%.³¹ By joining Miami, James would increase Miami's predicted per possession point differential from 0.043 to 0.219. James ultimately signed with Miami, sacrificing individual output for team success. Workers and firms, more generally, can also face these types of tradeoffs when choosing where to work or how to allocate workers. The result here is interesting in that it suggests a tension that firms need to balance between team and individual productivity, especially for firms where individual productivity has a large effect on compensation.

A second example of a how a worker's value can be heterogenous across firms according to the set of incumbent workers can be seen in the free agency case of Amar'e Stoudemire. At the end of the 2009-2010 season, Stoudemire had played eight consecutive seasons with the Phoenix Suns and at the time was described as an offensively skilled center with injury concerns. Reportedly, Phoenix

³¹Interestingly, James' scoring average per 36 minutes in his first year in Miami declined by 8.5% relative to his last three years in Cleveland.

was only willing to give Stoudemire a four year contract at an undisclosed salary, while the New York Knicks were willing to sign Stoudemire to a five year contract for 100 million dollars. Stoudemire ultimately signed with New York, but should New York have been willing to give Stoudemire more than Phoenix? Table 8 shows the predicted performance for both New York and Phoenix with and without Stoudemire on the team. The final column of the table shows the change in each team's per possession point differential with Stoudemire instead of an average player. Given their projected lineups, Stoudemire was more valuable to the Knicks than to the Suns based on predicted team performance. Thus, our model is consistent with the Knicks' decision to offer Stoudemire a more lucrative deal. From Stoudemire's standpoint, the Knicks were also more attractive in terms of individual performance, as his points per possession is predicted to be 8% higher.³²

4.7 Returns to the three factors

The previous section highlights the usefulness of our model for evaluating potential personnel decisions, but teams need to decide not only which players to obtain but also how much to pay them. Table 3 indicates that the three player skills we have identified, offensive intercept, offensive spillover, and defensive intercept, are associated with improved team performance. In this section, we examine whether player compensation correlates with our measures of player skill.

Table 9 provides the results from a series of OLS regressions where the dependent variable is log annual earnings. The annual earnings data comes from prosportstransactions.com and storytellerscontracts.com and is pro-rated equally over the course of a multi-year contract.³³ The skill measures we use as regressors are based on performance in the previous year and vary across columns, allowing us to compare the predictive power of our skill measures and standard player measures. The unit of observation in these regressions is a player-season-team combination, where we observe each player for up to four seasons, from 2007 to 2010.³⁴

³²In his first year with the Knicks, Stoudemire's scoring average per 36 minutes actually increased by 3%.

³³Most NBA contracts are fully guaranteed and some contracts include an incentive provision. The data we use includes information on these provisions and indicates whether the incentive is likely or unlikely to occur based on information at the start of the contract. We include incentives that are likely to occur in the annual earnings measure based on the definition of a likely incentive provided by storytellerscontracts.com.

³⁴In the NBA, first-round draft choices are assigned salaries according to their draft position. Each contract is for two years, with a team option for the third and fourth seasons. These structured contracts may weaken the relationship between skills and earnings, since players are drafted based on potential, not performance. However, when we estimate the earnings regressions using only players who have been in the league for at least four years, the results are nearly the same.

The first column of results indicate that a one standard deviation increase in a player’s offensive intercept is associated with a statistically significant 41% increase in annual earnings. A one standard deviation increase in the defensive intercept is associated with a 15% increase in earnings, while there is essentially no monetary gain to being a better spillover player. The results are robust to controls for player position and experience. In column 4 we examine whether the returns to skills changed over the four seasons in our sample. We include an interaction between a linear time trend and each of the three player parameters. We re-centered the linear time trend at the first year in our data so the main effects of the three parameters are the relationship between those parameters and earnings in the first year of our data and the interaction terms indicate how these have changed over time. None of the interactions are statistically significant, indicating little change in the return to skills across our sample. The last two columns in Table 9 provide estimates for two commonly used measures of player performance, the player efficiency rating (PER) and the adjusted plus-minus (APM). Players with higher PER and APM tend to earn significantly more than other players. Our skill measures explain more of the variation in log earnings than APM, but slightly less than PER. This makes sense as PER is widely available while APM is not. The results across regressions suggest that teams tend to compensate players for easily measured statistics (high R^2 for PER), but fail to identify players that add to team performance in difficult to observe ways (no effect of spillover skill).³⁵

One potential reason for the apparent lack of return to the spillover factor is that this parameter is somewhat noisier than either the offensive or defensive intercept. So rather than estimate a log earnings regression using a single player-season-team skill measure, we examine how average earnings over the four seasons in our sample is related to a player’s possession weighted average skill measures. Table 10 shows the results of these regressions. Similar to the results from Table 9, players with higher offensive and defensive intercepts are rewarded with higher total earnings, with earnings that are 46% and 18% higher per standard deviation of performance respectively. However, the results now indicate that high spillover players also earn significantly more than low spillover players. A one-standard deviation increase in a player’s average offensive spillover parameter is associated with an increase in total earnings of approximately 10%. Again, the PER measure explains the greatest amount of variation in total earnings, followed by the three skill factors and APM.

One way to measure the degree to which productivity spillovers are undervalued by NBA teams

³⁵All the results in Table 9 are nearly identical if we restrict the sample to just those players who are in the first year of a next contract. This alternative approach reduces the sample size but ensures that the performance measures pre-date the determination of the player’s salary.

is to compare differences in the relationship between the different productivity parameters on wages and team success. Based on the results from Table 2, we find that a standard deviation increase in own productivity results in 1.6 times as much team success as a standard deviation increase in the spillover parameter, but results in 8.7 times as much income for the player (Table 10). Thus, productivity spillovers are undervalued by more than a factor of 5 relative to own productivity.

We can take the idea of spillover mispricing one step further by asking what would happen if a team trades away a player that had an offensive intercept parameter one standard deviation above the mean and used the money to purchase a player(s) with higher spillover parameters. The results in Table 10 suggest that obtaining an extra standard deviation of offensive intercept costs about 5 times as much as obtaining an extra standard deviation of the offensive spillover. Combining this with the coefficients from Table 2, we find that a team that gives up a standard deviation better offensive intercept player experiences a 0.027 drop in per-possession point differential. However, with the money saved, the team is able to purchase players that result in an additional 5 standard deviations of spillover productivity. This increase in spillover productivity would raise the per-possession point differential for the team by 0.084 (based on 5×0.0168). Thus, the net benefit of making this trade would be an increase in the per-possession point differential of the team of 0.057. In our sample from 2007-2010, the average team had 92.89 possessions per game. This would give the team a net increase of about 5.30 points per game relative to their opponent. Between 2007 and 2010, about 3.85% of NBA games fall within one point, 14.72% fall within three points, and 26.84% of games fall within five points. Thus, the proposed trade for skills would have a significant impact on expected total wins with no associated increase in salary costs.

An important note about the numbers above is that they are based on the average expected change in performance across teams using the four most utilized players on each team. The benefit of acquiring a high spillover player will depend on the other players already on the team. As such, there are likely to be teams that wouldn't benefit from acquiring a player with a higher spillover parameter. An attractive feature of our model is that we can use it to evaluate the expected results of specific trades, as in Tables 7 and 8.

Finally, our conclusion that spillover productivity is mispriced is based on the assumption that teams are trying to maximize point differentials. If teams have other objectives, such as ratings, then they may be willing to pay more for high offensive intercept players. It also may be the case that offense is preferred to defense, explaining why the returns to offensive intercepts are higher than defensive intercepts. But in this case we might expect the returns to spillovers to be larger than

defensive intercepts. The fact that the returns to spillovers are so much smaller than the returns to the other two factors suggest that the difficulty in quantifying the presence of the spillovers likely plays a role in its low return.

5 Conclusion

Worker skills are multidimensional. One of the skills that may be important to a variety of production processes is the ability to bring out the best in others. In this paper, we use data from the NBA to identify three measures for each player: their ability to score, their ability to defend, and their ability to help others score. It is this last factor that differentiates our work and also substantially complicates estimation. Using an iterative approach along the lines of Arcidiacono et al. (2012), we show that estimating models of this type can be accomplished in a straightforward manner.

We find that all three factors are important components to overall team productivity and probability of success. Ignoring spillovers has a substantial effect on assessing the overall contribution of specific players causing previous approaches to underestimate the contribution of “team” players. We also find that there are complementarities in production between direct forms of productivity and indirect forms that operate through productivity spillovers. As such, some teams will value particular players more than others based on the current composition of their team. Players who are particularly strong at scoring but are not good facilitators will be more valued by teams that are composed of players who are not very strong at scoring themselves.

We also find that players are primarily compensated based on their direct contributions to team production with little weight given to their ability to increase the productivity of their teammates. This misalignment of incentives might reduce the incentive for players to invest in or engage in actions that increase their positive effects on the productivity of their teammates, especially in cases where compensation is based on relative performance.

The approach we take here to measuring both one’s productivity and one’s ability to help others be more productive, as well as the corresponding implications for compensation, could be applied to a number of other labor markets. One example is in the production of patents. Researchers in a lab may both be working on their own patents as well as serving as a resources for other members of the lab. The measurable output is who is on the patent. But others in the lab may have facilitated the process and these facilitation skills may not map one-to-one in the skills needed to produce a patent. When facilitation skills become more easily measured and rewarded, then incentives to invest these skills also increases. Financial advisors provide another example. Compensation in this industry

tends to be based on one's clients, not how helpful one is to other financial advisors. This is likely due to the difficulty in measuring how other advisors affect both the number of the clients an advisor has and the quality of the advice the advisor is giving. But with data on clients and investment returns, as well as variation in the composition of the team, it would be possible to measure the value of these sorts of interactions.

References

- Arcidiacono, Peter, Jennifer Foster, Natalie Goodpaster, and Josh Kinsler. 2012. Estimating spillovers using panel data, with an application to the classroom. *Quantitative Economics* 3, no. 3:421-70.
- Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang. 2010. Superstar extinction. *The Quarterly Journal of Economics* 125, no. 2:549-89.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2009. Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica* 77, no. 4:1047-94.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2010. Social incentives in the workplace. *Review of Economic Studies* 77, no. 2:417-58.
- Burke, Mary A., and Tim R. Sass. 2013. Classroom peer effects and student achievement. *Journal of Labor Economics* 31, no. 1:51-82.
- Cornelissen, Thomas, Christian Dustmann, and Uta Schonberg. 2013. Peer effects in the workplace. IZA DP No. 7617.
- Falk, Armin, and Andrea Ichino. 2006. Clean evidence on peer effects. *Journal of Labor Economics* 24, no. 1:39-57.
- Gould, Eric D., and Eyal Winter. 2009. Interactions between workers and the technology of production: Evidence from professional baseball. *Review of Economics and Statistics* 91, no. 1:188-200.
- Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo. 2009. Peer effects in the workplace: Evidence from random groupings in professional golf tournaments. *American Economic Journal: Applied Economics* 1, no. 4:34-68.
- Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan. 2003. Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation. *Journal of Political Economy* 111, no. 3:465-97.
- Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi. 1997. The effects of human resource management practices on productivity: A study of steel finishing lines. *American Economic Review* 87, no. 3:291-313.

Mas, Alexandre and Enrico Moretti. 2009. Peers at work. *American Economic Review* 99, no. 1:112-45.

Price, Joseph and Justin Wolfers. 2010. Racial discrimination among NBA referees. *The Quarterly Journal of Economics* 125, no. 4:1859-87.

Waldinger, Fabian. 2010. Quality matters: The expulsion of professors and the consequences for PhD student outcomes in Nazi Germany. *Journal of Political Economy* 118, no. 4:787-831.

Waldinger, Fabian. 2012. Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany. *Review of Economic Studies* 79, no. 2:838-61.

A Basic Rules of NBA Basketball

An NBA basketball game consists of four 12 minute quarters. The court is ninety-four feet long and fifty feet wide. At each end of the court is a horizontal metal ring ten feet above the floor with a glass board behind each ring. Each team has five players on the court at a given time, though coaches can substitute players rather freely throughout the contest. A team's goal is to score more points than the other team. Scoring is achieved by throwing or shooting the ball through the team's metal hoop. A successful shot scores two or three points depending on the distance and provides the opposing team with possession of the basketball. Only one player on a team can score per possession, while the remaining players act to facilitate this activity. Teams must shoot the ball every 24 seconds or relinquish possession to the opposing team.

Using physical contact to gain an advantage is by definition a foul. When a referee concludes a foul was committed by a player, a whistle is blown causing play to stop. If the fouled player was in the act of shooting while being fouled, the player is awarded a number of free throws equivalent to the value of the shot being attempted. A free throw is a shot from 15 feet away from the basket in which the player cannot be defended by the opposing team. Each free throw is worth one point. If a player was not shooting, the fouled player's team is awarded the ball out of the bounds of the court and play continues by the team passing the ball in play to a teammate. If players accumulate six fouls, they are excluded from the remainder of the game.

B Data Appendix

There were a number of intermediate steps required to transform the raw play-by-play data that we gathered from espn.com into our final dataset. First, we had to determine which players were on the court at each point during the game. Since the play-by-play data does not provide a running list of who was on the court over the course of the game, we had to infer who was on the court based on the players we observed in the data.

Teams can freely substitute players at the start of each quarter and none of these substitutions appear in the data. However, for any substitutions that occur during the quarter, we observe both the players coming in and the players going out. We combine this with information on the names of the players that record some action in the data to construct the set of five players on each time for every play-level observation in the data.

The second step required us to transform our play-level data into a single observation for each

possession. An offensive possession begins anytime a team obtains the ball and switches from defense to offense. Possessions can end in many ways, such as a made basket, a missed shot, a turnover, or the end of a quarter. For possessions that resulted in positive points, we also captured which player on the offensive team scores and how many points they scored. For this study, any play by the offensive team that extends a possession (such as an offensive rebound) does not create new possession but just becomes the continuation of the possession already going. The one exception is when an offensive rebound occurs following foul shots, since it is very common for substitutions to occur during foul shots.

If in the middle of a possession there is a substitution, the player entering the game is the one considered on the court for that possession. The one exception is substitutions that occur during fouls shots in which case the players coming out are considered part of the possession that resulted in the foul and any points scored from the foul shots are credited to that possession.

From the defensive standpoint, the only relevant outcome is whether the offensive team scores positive points. Steals, blocks, and defensive rebounds will get reflected in an increased probability that the offensive team does not score.

At the end of this process we were left with 915,580 unique possessions. We dropped around 1% of these possessions either because we could not identify all of the players on the court or identify the player who shot the basket. A possession that either has too few players on the court or a player on the court more than once typically indicates a data entry error in the play-by-play data. Often this implies that active lineups for other possessions during that quarter are likely to be incorrect. As a result, any quarter that has a possession with either too few or too many players is dropped.

Finally, our empirical strategy requires us to estimate the model separately by season. If during a season a player never scores nor is ever part of a defensive unit that keeps the other team from scoring, then that player's offensive and defensive parameters are not identified. Thus, we identify who these players are, and then eliminate all possessions during which these players are on the court. Typically there are about five to ten players per season who fall into this category.

C Standard Error Calculations

In Section 3.3 we outlined an iterative estimation procedure that avoided jointly estimating all the parameters in the model. The downside to this approach is that it is not straightforward to compute standard errors. However, once the iterative procedure has converged, we have all the parameter estimates and can compute standard errors using the Outer Product of the Gradient. This is a

particularly useful method in this case since the logit probabilities yield simple closed form solutions for estimating the gradient.

Our joint log-likelihood function is given by,

$$\ln L(o, s, d) = \sum_{n=1}^N \sum_{i=0}^5 (d_{in} = 1) \ln Pr(y_{in} = 1|P_n)$$

where

$$Pr(y_{in} = 1, i > 0|P_n) = \frac{\exp\left(o_i(1 - \sum_{k \in O_n, k \neq i} s_k)\right)}{\exp\left(\sum_{j \in D_n} d_j\right) + \sum_{h \in O_n} \exp\left(o_h(1 - \sum_{k \in O_n, k \neq h} s_k)\right)}$$

and

$$Pr(y_{0n} = 1|P_n) = \frac{\exp\left(\sum_{j \in D_n} d_j\right)}{\exp\left(\sum_{j \in D_n} d_j\right) + \sum_{h \in O_n} \exp\left(o_h(1 - \sum_{k \in O_n, k \neq h} s_k)\right)}$$

and d_{in} indicates whether outcome i occurred for possession n . For ease of exposition I refer to $Pr(y_{in} = 1|P_n) = P_{in}$. Also, define $\hat{\theta} = \{\hat{o}, \hat{s}, \hat{d}\}$.

The first step is to construct the estimate of the gradient vector, $\hat{g}_n = \frac{\partial \ln P_{in}(\hat{\theta})}{\partial \hat{\theta}}$. For any possession n , there are at most 15 parameters involved, 5 each for o , s , and d . Thus, the gradient for each observation will only have 15 non-zero elements. Again, because of the closed form for the logit probability, the derivatives of the log probability are rather straightforward.³⁶ The estimated variance covariance matrix is then given by $\Sigma_{\hat{\theta}} = \left[\sum_{n=1}^N \hat{g}_n \hat{g}'_n\right]^{-1}$.

For estimation we restrict the spillover parameter for one player on each team-season to equal zero. Post-estimation we adjust the spillover estimates such that the possession-weighted mean spillover for each team season is equal to zero. This way we have a spillover estimate for each player. In doing this we also have to adjust the offensive intercept parameters to keep the scoring probabilities constant. Both of these changes require adjusting $\Sigma_{\hat{\theta}}$, which is a straightforward application of the delta method.

The final step of our estimation procedure is to link players across teams and seasons. This process yields a string of team-season fixed effects (δ_t) with which we update our parameters according to the process outlined in Section 3.4. We simply apply the delta method iteratively using the δ_t 's to update the standard errors for the team-season specific parameters.

³⁶Precise formulas are available upon request.

Table 1: Sample Statistics

Seasons Covered	2006-2007 through 2009-2010
Total Possessions (Involving 5 Offensive and Defensive Players)	915,580
Utilized Possessions	905,378
Fraction of Possessions Discarded	0.01
Unique Players	656
Average Possessions Per Player	13,801
SD Possessions Per Player	12,882
25th Percentile Of Possession Distribution	2,342
75th Percentile Of Possession Distribution	22,609
Average Possessions Per Player-Season	5,081
SD Possessions Per Player-Season	3,557
25th Percentile Of Possession Distribution	1,697
75th Percentile Of Possession Distribution	8,083
Average Possessions Per Player-Season-Team	4,507
SD Possessions Per Player-Season-Team	3,570
25th Percentile Of Possession Distribution	1,130
75th Percentile Of Possession Distribution	7,470
Proportion of Possessions with Positive Points	50.8
Avg. Points Per Possession	1.06
SD Points Per Possession	1.11
Avg. Points Per Possession Points>0	2.1

Table 2: Skills and Winning, 2009-2010

Take the 4 most utilized players on each team and...	Point Differential	SD Point Differential	Minimum	Maximum
Add average player	0.0048	0.0873	-0.1536	0.257
	Δ Point Differential	SD Δ Point Differential	Minimum	Maximum
Add 1 SD Better Intercept Player	0.0270	0.0022	0.0217	0.0307
Add 1 SD Worse Intercept Player	-0.0217	0.0018	-0.0252	-0.0177
Add 1 SD Better Spillover Player	0.0168	0.0012	0.0146	0.0202
Add 1 SD Worse Spillover Player	-0.0167	0.0012	-0.0200	-0.0146
Add 1 SD Better Defensive Player	0.0210	0.0002	0.0200	0.0211
Add 1 SD Worse Defensive Player	-0.0209	0.0002	-0.0211	-0.0202

Unit of observation is an NBA team in the 2009-2010 season.

Table 3: Average Skills by Position

	Point Guard	Shooting Guard	Small Forward	Power Foward	Center
	(All measures standardized at the population level)				
Offensive Intercept	-0.147	0.081	-0.013	0.069	0.006
Offensive Spillover	0.168	-0.091	-0.012	-0.073	0.012
Defensive Intercept	-0.511	-0.305	-0.097	0.315	0.683
Overall Rank	-0.298	-0.108	-0.016	0.134	0.327
PER	-0.039	-0.112	-0.107	0.065	0.218
APM	-0.101	0.039	0.057	0.015	-0.011
Observations	395	410	349	438	417

Unit of observation is a player-season-team combination. Means are constructed by weighting the total number of possessions for a player-season-team combination. Across all positions the mean of each measure is zero with a standard deviation of 1. Note that the number of observations for the APM measure is smaller since players with few possessions are not assigned a rating.

Table 4: Player Rankings based on Standardized Point Differentials, 2009-2010

Point Differential w/ Preferred	Average Team	Position Adjusted	Δ in Team Point Differential when Replaced with Average	PER	APM				
Dwight Howard	3.195	Dwyane Wade	3.017	Dwight Howard	3.266	LeBron James	3.866	Dwyane Wade	3.873
Kevin Durant	2.907	Deron Williams	2.889	Kevin Durant	2.807	Dwyane Wade	3.119	LeBron James	3.657
Dwyane Wade	2.629	Kevin Durant	2.772	Dwyane Wade	2.484	Kevin Durant	2.686	Dwight Howard	3.283
Al Horford	2.501	Dwight Howard	2.648	Al Horford	2.283	Chris Bosh	2.396	Steve Nash	2.839
Deron Williams	2.496	Andre Miller	2.267	Deron Williams	2.256	Tim Duncan	2.324	Chris Andersen	2.397
LeBron James	2.187	Dirk Nowitzki	2.074	Dirk Nowitzki	2.010	Dwight Howard	2.156	Ray Allen	2.291
Dirk Nowitzki	2.029	LeBron James	2.062	LeBron James	1.968	Chris Paul	2.083	Kobe Bryant	2.243
Marc Gasol	1.950	Al Horford	2.023	Marc Gasol	1.909	Dirk Nowitzki	1.891	Kevin Durant	2.136
Chris Andersen	1.915	Chris Andersen	1.958	Andre Miller	1.874	Pau Gasol	1.891	Deron Williams	2.045
Andre Miller	1.897	Stephen Jackson	1.728	Chuck Hayes	1.542	Amar'e Stoudemire	1.818	Nene Hilario	1.935

Table 5: Player Rankings Ignoring Spillovers, 2009-2010

Point Differential w/ Average Team	Spillover Differential - No Spillover Differential			
	Notable Δ 's	Largest Positive Δ 's	Largest Negative Δ 's	
Dwight Howard 3.126	Carmelo Anthony -1.287	Jason Williams 2.788	James Singleton -1.898	
Kevin Durant 2.828	Vince Carter -1.083	Jon Brockman 2.145	Andray Blatche -1.781	
Andrew Bogut 2.345	Allen Iverson -1.086	Jamaal Tinsley 2.128	Eric Maynor -1.518	
LeBron James 2.334	Tim Duncan -1.033	Mike Miller 2.062	Chris Kaman -1.499	
Tim Duncan 2.121	Kobe Bryant -0.963	Theo Ratliff 1.704	Mo Williams -1.481	
Vince Carter 2.092	Shane Battier 1.036	Chris Andersen 1.703	Luis Scola -1.329	
Kobe Bryant 2.071	Deron Williams 1.130	Anderson Varejao 1.617	Sam Young -1.306	
Dwyane Wade 2.035	Al Horford 1.148	DeShawn Stevenson 1.598	Carmelo Anthony -1.287	
Chris Bosh 1.882	Steve Nash 1.334	Shaun Livingston 1.480	Sergio Rodriguez -1.245	
Dirk Nowitzki 1.855	Chris Andersen 1.703	Earl Watson 1.375	Mario Chalmers -1.211	

Table 6: Robustness Exercises

	2010	2009	2008	2007
High/Low Possessions				
Corr(Offensive Intercepts)	0.98	0.98	0.98	0.98
Corr(Offensive Slopes)	0.90	0.92	0.90	0.91
Corr(Defensive Intercepts)	0.94	0.94	0.94	0.94
Corr(Standardized Per Poss. Point Differential)	0.91	0.92	0.91	0.92
Corr(Standardized Per Poss. Point Differential), >2000 Possessions	0.90	0.92	0.91	0.92
Corr(Player Rank), >2000 Possessions	0.88	0.92	0.92	0.92
No 4th quarter				
Corr(Offensive Intercepts)	0.95	0.95	0.95	0.95
Corr(Offensive Slopes)	0.79	0.79	0.77	0.79
Corr(Defensive Intercepts)	0.85	0.82	0.80	0.82
Corr(Standardized Per Poss. Point Differential)	0.77	0.86	0.78	0.86
Corr(Standardized Per Poss. Point Differential), >2000 Possessions	0.75	0.86	0.79	0.86
Corr(Player Rank), >2000 Possessions	0.72	0.84	0.81	0.83

Unit of observation is a player-season-team combination. Correlations are between the baseline estimates and the estimates generated by the particular robustness exercise. All correlations are possession weighted. Further details on the High/Low Possession exercise and the No 4th quarter exercise can be found in Section 4.5.

Table 7: LeBron James' Free Agency

If James joins ...	Projected Teammates	Probability James Scores	Team per Possession	
		per Possession	Points	Point Differential
Chicago Bulls	Boozer, Deng, Noah, Rose	0.175	1.204	0.136
Cleveland Cavaliers	Williams, Hickson, Parker, Varejao	0.187	1.142	0.043
Miami Heat	Chalmers, Bosh, Ilgauskas, Wade	0.165	1.244	0.219
New York Knicks	Chandler, Felton, Gallinari, Stoudemire	0.172	1.155	0.126

To generate predicted outcomes for each team we utilize player estimates based on all four years of data.

Table 8: Amar'e Stoudemire's Free Agency

	Lineup	Stoudemire Points Per Possession	Team per Possession Points Point Differential	Δ in Team Differential w/ Stoudemire
If Knicks sign Stoudemire ...	Stoudemire, Chandler, Felton, Gallinari, Jeffries	0.352	1.088 0.068	0.026
If Knicks don't sign Stoudemire ...	Average Player, Chandler, Felton, Gallinari, Jeffries		1.066 0.042	
If Suns sign Stoudemire ...	Stoudemire, Nash, Hill, Frye, Richardson	0.326	1.205 0.102	0.016
If Suns don't sign Stoudemire ...	Average Player, Nash, Hill, Frye, Richardson		1.193 0.086	

To generate the predicted outcomes for each team we utilize player estimates based on all four years of data.

Table 9: Lagged Skills and Wages

Dependent Variable: Log Earnings						
Offensive Intercept	0.414*	0.414*	0.356*	0.340*		
	(0.025)	(0.025)	(0.018)	(0.027)		
Offensive Slope	-0.018	-0.017	-0.033	-0.016		
	(0.026)	(0.026)	(0.019)	(0.030)		
Defensive Intercept	0.150*	0.140*	0.069*	0.038		
	(0.023)	(0.026)	(0.022)	(0.029)		
Player Efficiency Rating					0.356*	
					(0.031)	
APM						0.223*
						(0.022)
Experience			0.394*	0.396*	0.388*	0.400*
			(0.023)	(0.023)	(0.024)	(0.027)
Experience ²			-0.019*	-0.019*	-0.019*	-0.020*
			(0.002)	(0.002)	(0.002)	(0.002)
Position Effects	N	Y	Y	Y	Y	Y
Time Trends	N	N	N	Y	N	N
R ²	0.239	0.241	0.593	0.596	0.605	0.498
N	1273	1273	1273	1273	1273	1094

Unit of observation is a player-season-team combination. Robust standard errors in parentheses. * Indicates a coefficient that is statistically significant at a 5% level. Coefficients are estimated by OLS weighting each observation by the total number of possessions associated with that player-season-team combination. All skill measures are standardized to have a mean of zero and a variance of one.

Table 10: Skills and Wages over Career

Log Average Earnings 2007-2010					
Average Offensive Intercept	0.464*	0.461*	0.419*		
	(0.039)	(0.039)	(0.024)		
Average Offensive Slope	0.096*	0.091*	0.048*		
	(0.038)	(0.037)	(0.022)		
Average Defensive Intercept	0.182*	0.188*	0.084*		
	(0.033)	(0.041)	(0.026)		
Average Player Efficiency Rating				0.426*	
				(0.022)	
Average APM					0.266*
					(0.025)
Average Experience			0.338*	0.334*	0.340*
			(0.022)	(0.021)	(0.028)
Average Experience ²			-0.015*	-0.016*	-0.017*
			(0.002)	(0.002)	(0.002)
Position Effects	N	Y	Y	Y	Y
R ²	0.284	0.295	0.702	0.719	0.580
N	656	656	656	656	494

Unit of observation is a player. Robust standard errors in parentheses. * Indicates a coefficient that is statistically significant at a 5% level. Coefficients are estimated by OLS weighting each observation by the total number of possessions associated with that player over all four seasons. Average skill measures are also constructed as a possession weighted average of the player-season-team measures. The average skill measures are then standardized to have a mean of zero and a variance of one.