

Performance Variation in the NBA: Guaranteed Contracts and the Contract Year Phenomenon

Neal Jean

Professor Peter Arcidiacono, Faculty Advisor

*Honors Thesis submitted in partial fulfillment of the requirements for Graduation with
Distinction in Economics in Trinity College of Duke University.*

Duke University
Durham, North Carolina
2010

Acknowledgments

I would like to thank my faculty advisor, Professor Peter Arcidiacono, an avid follower of the NBA in his own right. Without his insights and ideas, much of the quality of my paper would not have been possible. I would also like to thank Professor Kent Kimbrough, who encouraged me to pursue a topic that I was truly interested in and helped to guide the formulation of my research goals and objectives. My gratitude goes out to Professor Michelle Connolly who, despite having no prior interest in basketball, spent hours reading and discussing my paper with me, and helped with everything from the development of theoretical hypotheses to the empirical analysis. Finally, I am indebted to my Economics 198S and 199S classmates for their patience and comments. In the spirit of this paper, all fouls are my own.

Abstract

Asymmetric information in the typical principal-agent model may create incentives for workers to participate in opportunistic or strategic behavior. This paper examines evidence for the existence of *ex ante* and *ex post* strategic behavior in the National Basketball Association (NBA), where multi-year guaranteed contracts give players incentive to vary their effort level over the course of the contract cycle. I hypothesize that players will increase effort in a contract year, the season prior to signing a new contract, and decrease effort in the following season, with the assumption that variations in effort are expressed through variations in performance. Player productivity is found to increase significantly in a contract year, while no evidence is found to support the presence of shirking behavior. The resulting model may serve to guide NBA teams in structuring contracts, as appropriate performance incentives built into individual player contracts could positively influence team success.

I. Introduction

In the principal-agent models present in most organizations, inefficiencies may arise from the creation of long-term contract agreements. Knowing that their future wages are guaranteed, workers have an incentive to participate in opportunistic or strategic behavior, generally referred to as *shirking behavior* (Alchian and Demsetz 1972; Holmstrom 1979). The moral hazard problem leading to the occurrence of shirking behavior results from asymmetrical information between the parties involved.¹ In general, opportunistic behavior may occur under the following conditions:²

1. The principal and the agent enter into a binding, long-term contract, through which the agent is compensated by the principal for expending effort at some cost to himself.
2. The principal's level of utility or profit is a function of the agent's level of effort expended.
3. The principal is unable to monitor the effort of the agent directly.

Under these conditions, the agent may decide that he has sufficient incentive to strategically vary his effort level. This paper will examine the question of whether this type of strategic behavior exists in the context of long-term guaranteed contracts in the National Basketball Association (NBA).

Maxcy et al. (2002) introduces the terminology *ex ante* strategic behavior and *ex post* shirking to describe potential worker behavior under circumstances of asymmetrical

¹ Regarding asymmetrical information, Holmstrom (1979) mentions that situations in which the private actions of individuals affect the outcome distribution have long been recognized for the risk of moral hazard. Since individual actions cannot be observed and contracted upon, contracts are written without complete information.

² As described by Krautmann and Donley (2009).

information. *Ex ante* strategic behavior refers to an increase in effort expended by the agent prior to the signing of a new contract. In attempting to increase his perceived value to the principal, the agent seeks to maximize the level of compensation guaranteed him by the impending contract. *Ex post* opportunistic behavior, or *shirking*, is defined as a decrease in effort expended by the agent once the contract is signed. Since compensation is then guaranteed, regardless of effort level, the agent has no incentive to continue performing with maximal effort.³ A key point here is that the employer is unable to observe employee effort level directly; if effort-monitoring were possible, the employer would be able to structure contracts accordingly and guarantee an efficient outcome.⁴

Strategic behavior as a result of long-term guaranteed contracts has been widely cited in academic economic literature and also carries significant real-world implications, arising in such diverse situations as manufacturing sales, political decisions, nuclear power plants, and the academic institution of tenure (Krautmann 1990).⁵⁶ Whenever agents are faced with the incentives created by long-term contracts, they are also presented with motivation to vary their effort to maximize their own utility. Consider the example of university professors and tenure: One can imagine that a professor seeking

³ A brief note on terminology: Strategic behavior and opportunistic behavior will be used interchangeably to describe any type of changing behavior as a result of information asymmetry during the contract cycle. The term shirking will be used to describe *ex post* strategic or opportunistic behavior.

⁴ If a principal were able to observe effort directly, he could simply compensate agents directly for the level of effort that they contribute toward the principal's production. In reality, the principal must make an estimate of the agent's future effort level and productivity that may be unmet by the agent's actual contribution.

⁵ Oyer (1998) examines the impact of incentives as they relate to the sales of manufacturing firms, and finds that sales are often boosted near the end of the fiscal year as salesmen try to hit their incentive targets and increase their own compensation.

⁶ Krautmann (1990) gives an interesting example illustrating the usefulness of being able to distinguish shirking behavior from the outcome of a stochastic process. He says to imagine a safety regulation scheme for nuclear power plants that induces the socially optimal maintenance and safety effort. Even with this optimized system, he argues that accidents can still occur as the result of random processes. In these cases, regulators would be better served to categorize these failures as random events rather than assigning fault to and punishing the individual plant, as this would cause the plant to respond by overspending on future maintenance and accident prevention, rather than simply spending the original, optimal amount.

tenure would be very productive, publishing original research on a regular basis, but upon achieving tenure, might lose motivation and kick back, enjoying the flexible lifestyle afforded him. We of course hope that most professors are more scrupulous and resist the temptation to take advantage of the opportunity to exert effort strategically, but this example makes it clear that such opportunities exist.⁷

Professional sports offer an ideal setting to study the incentive effects of long-term contracts, as players are carefully observed by teams, fans, and media, resulting in an expansive and thorough set of productivity data that is widely available for analysis. Already, a number of researchers have utilized the unique environment of professional sports to search for evidence supporting the shirking hypothesis, with varying results.⁸ Most of these studies have focused on Major League Baseball (MLB), as performance metrics that measure player productivity are more widely agreed-upon in baseball, although a few have also examined the NBA.⁹ Perhaps as a result of employing different empirical specifications, research concerning professional sports has drawn conflicting conclusions. Stiroh (2007) claimed that there is strong evidence that players engage in strategic behavior, while many others have refuted the existence of such behavior (Krautmann 1990; Maxcy et al. 2002; Berri and Krautmann 2006; Krautmann and

⁷ Stiroh (2007) points out that in other situations involving long-term contract cycles (e.g., professors working for tenure or politicians seeking election), the inefficiencies caused by incentivized strategic behavior could be even higher than in the NBA, due to greater information asymmetry and less monitoring ability.

⁸ Krautmann and Donley (2009) give a brief description of the typical study in professional sports. Usually the researchers examine the player's performance in proximity to contract negotiations. Conclusions about the presence of shirking behavior are drawn from the discrepancies between observed *ex ante* and *ex post* performance and expected performance.

⁹ Berri and Krautmann (2006) remark that baseball has a number of indices that have been designed specifically to capture the productivity of an individual player in a single number.

Donley 2009).¹⁰ The goal of this paper is to explore the *ex ante* and *ex post* phenomena of opportunistic behavior in a specific labor market, the National Basketball Association (NBA).

The NBA was chosen from among the three major professional sports leagues in the United States as the league to be studied for several reasons. From a practicality standpoint, the prevalence of long-term guaranteed contracts in the NBA provides a textbook setting in which to study shirking behavior—unlike, for example, the National Football League (NFL), in which large portions of a player’s contract are not guaranteed, and players can be cut or released at the team’s discretion. The relative scarcity of performance-based incentives in the NBA also contributes to an environment that would support the shirking hypothesis.

In terms of impact, the question of whether or not strategic behavior exists in the NBA is arguably more important, and certainly more complex, than in the NFL or MLB. Team sizes are much smaller in the NBA, making each athlete—and accordingly, each athlete’s contract—that much more important to the success of the team. Think back to the 2006-07 season, when LeBron James carried a team of largely mediocre players all the way to the NBA Finals, or the next season, when the Celtics’ acquisition of Kevin Garnett and Ray Allen led to a year-on-year improvement of 42 wins and culminated in the 2007-08 NBA Championship. Unlike football, in which the vast majority of players play either offense or defense but not both, or baseball, where most players spend only a

¹⁰ Stiroh (2007) found evidence of shirking behavior in the NBA, while studies refuting the shirking hypothesis have been conducted using both NBA and MLB data. Maxcy (2002) rejects strategic behavior, citing that the data suggests that mechanisms incorporated into the structure of MLB contracts intended to prevent players from strategically adjusting their effort seem to be fulfilling their role. Examples of such mechanisms include performance incentives, team success bonuses, and clauses that allow for performance-based contract renegotiation.

fraction of the game directly participating in the action, basketball is a game in which players are constantly involved.¹¹ While on the court, NBA players are continuously contributing to both the offensive and defensive efficiency of their team. Individual athletes make a greater impact in basketball than in any other major team sport popular in the United States.

While individual contributions are generally more important in basketball, an added layer of complexity arises when one attempts to quantify these contributions. Compared to baseball, the other prominent arena for studying shirking behavior, player effort in basketball is notoriously difficult to measure. To illustrate, imagine a baseball player who is up to bat. Presumably, there are players who care more than others about whether or not they get a hit, but it seems unlikely that this will affect their effort to hit the ball. It is debatable whether or not a player can hit a curveball more effectively simply by trying harder. In basketball, the effect of effort on player performance is intuitively apparent, though not always easy to discern through statistical analysis. It makes perfect sense that a basketball player can play better defense by exerting more effort, getting into a lower defensive stance, and rotating to help his teammates. A player can also help his team by diving for loose balls or boxing out on offense to earn extra possessions. Of course, since teams cannot monitor the effort of their players directly—one of the conditions under which strategic behavior can occur—they must determine effective methods to indirectly measure player effort.

¹¹ One could make the argument that starting pitchers in baseball are able to affect the outcome of an individual game to a greater extent than any player in a basketball game. Though it is difficult to quantify contributions across sports to allow for comparison, intuitively this claim is not without merit. However, the key is that starting pitchers can influence an *individual* game. Since most pitchers start once every four to five games during the regular season, they are limited to affecting the team's performance at most 25% of the time.

Generally, teams resort to using statistics of player performance as a proxy for effort. The paradox here, however, is that while an individual can make a greater impact in basketball than in baseball, basketball is simultaneously a more team-oriented sport. If a baseball player wishes to improve his performance, he can easily see where he should focus his efforts. Hitting and fielding, the primary components of the game for most position players, are largely individual capacities. In basketball, however, the answer is much less easily defined. Consider two average and equally dedicated NBA players, Player A and Player B. Suppose Player A chooses to focus on improving his scoring ability; any improvement in this domain will show up in the box score at the end of the game. Now suppose Player B concentrates his effort on playing effective help defense or on boxing out for offensive rebounds; even the most dramatic improvements in these areas might not be reflected in the statistics. We see here that two players who devote an equal amount of energy may not appear to benefit the team in equal measure. A successful general manager, then, must discern a player's true value to his team in the face of divergent statistical inputs. Furthermore, how players decide to allocate effort may change over time. Clearly, these complexities will be problematic in determining an effective measure of player productivity to be used as a proxy for effort.

The two pieces of literature most relevant to my current research are Stiroh (2007) and Berri and Krautmann (2006). Both of these studies look for evidence of opportunistic behavior in the NBA, with Stiroh (2007) concluding that the hypothesis was strongly supported, and Berri and Krautmann (2006) claiming that there was no evidence that NBA players engaged in shirking behavior. My research will attempt to synthesize and build on the methods of Stiroh (2007) and Berri and Krautmann (2006),

ultimately arriving at a more effective method of determining the presence of strategic behavior in the context of NBA player contracts. Two main objectives will be pursued through this study: First, I hope to evaluate several methods of measuring player performance and then identify the measure that I believe to be most appropriate for this study. Second, I will determine whether or not sufficient evidence exists to support the hypothesis that players engage in strategic behavior, answering the original research question that was the inspiration behind this study.

Using a playing-time adjusted measure of player performance, I find that evidence for the contract year phenomenon, or *ex ante* strategic behavior, does exist in a dataset that includes more recent observations than those used in the aforementioned studies. *Ex post* shirking, however, is not supported by the data, although shirking behavior may be masked by confounding factors. Additionally, in the course of this study, I develop a method for separating the interrelated effects of age and experience on player productivity by measuring experience in such a way that it is no longer perfectly correlated with the passage of time.

The rest of this paper will be organized as follows: Section II will review the existing body of literature and discuss its relevance to the current research. Section III will describe the economic theory of shirking and its implications for my empirical analysis. Section IV will provide an overview of my data set and discuss any limitations that it might impose. Section V will establish my empirical specification. This section will include my expectations for the results of the study, as well as the actual results obtained. Section VI will conclude the paper, and discuss possibilities for further improvement and research on the topic.

II. Literature Review

Long-term Guaranteed Contracts in the NBA

Shirking behavior has the potential to occur when an agent is signed to a long-term, guaranteed contract, in which he is compensated by the principal for the expenditure of costly effort (Krautmann and Donley 2009). The NBA provides a large sample of these types of situations, as well over half of all players are under contract for two or more seasons.¹² Since contracts in the NBA are largely multi-year and guaranteed, players could choose to raise their effort prior to signing a long-term contract, then reduce their effort level substantially after locking in a guaranteed salary for the ensuing several years. Since opportunistic behavior is associated with the presence of long-term guaranteed contracts, one may reasonably wonder why the NBA continues to utilize multi-year contracts. Given that the nature of the industry limits the accuracy of information that employers can obtain about their workers' behavior, why wouldn't they simply eliminate the types of contracts that promote strategic behavior altogether?

Stiroh (2007) offers a discussion about why long-term contracts have not been replaced with alternatives such as incentive-based contracts or short-term contracts. Instead of fixing future compensation at the moment that the contract is signed, an alternative to the guaranteed contract would be to offer players long-term, incentive-based contracts. Incentive-based contracts compensate players for reaching certain

¹² Unlike contracts in the NFL, NBA contracts are fully guaranteed (Berri and Krautmann 2006). Regardless of whether or not a player plays in games, or even if the player is with the team, the team must continue to pay the player under the original terms of the contract.

standards of performance.¹³ A reduction in the guaranteed portion of the contract would motivate players because their pay would then be dependent upon their performance, rather than remaining fixed, as in the guaranteed contracts that currently prevail in the NBA. Many NBA contracts already contain performance-based incentives, though these usually compose only a small fraction of the total value of the contract.¹⁴ Unlike other professional sports leagues that make extensive use of performance incentives, such as the NFL or MLB, bonuses in the NBA are generally limited to far less than 25% of the player's total salary (Heubeck 2003; Stiroh 2007).¹⁵ Consider the example of Peyton Manning, the star quarterback of the NFL's Indianapolis Colts. In 2004, Manning signed a 7-year contract worth \$99.2 million. Manning could also earn up to an additional \$19 million in contract incentives. Of this total, only \$34.5 million was guaranteed money in the form of a signing bonus. In contrast, contracts in the NBA are almost fully guaranteed.

Prendergast (1999), in his work on firms providing incentives to their workers, discusses the benefits and drawbacks of pure incentive contracts.¹⁶ By construction, a contract that is entirely composed of performance incentives is far riskier than a contract in which all future payments are guaranteed. It follows that a risk-averse worker would

¹³ For example, it might be specified in a player's contract that he would be paid a base salary of \$10 million for the season, and incentives of \$1.5 million for averaging 20 points per game (ppg), \$1 million for averaging 10 rebounds per game (rpg), and \$1.25 million for averaging 5 assists per game (apg). If the player averaged 19 ppg, 12 rpg, and 4 apg, he would be paid the \$10 million base salary plus the extra \$1 million for averaging over 10 rpg for a total of \$11 million in compensation for the season.

¹⁴ Incentives can be included for achievements that are not purely statistical. For example, a player may have bonuses written into his contract that would reward him for winning league awards such as the Most Valuable Player (MVP), or for being voted onto the conference All-Star team or the All-NBA team.

¹⁵ Salary structure in the NBA is regulated by the collective bargaining agreement (CBA) between the National Basketball Players' Association (NBPA), which is the players' union, and the NBA. The CBA distinguishes between "likely bonuses" and "unlikely bonuses" in the incentives clauses included in player contracts, and state that "unlikely bonuses" cannot exceed 25% of the player's salary.

¹⁶ This would be the limiting case, in which all compensation is contingent on the agent fulfilling the terms of the incentive clauses written into the contract and there is no guaranteed compensation.

require the expected value of the total compensation under a pure incentive contract to be higher than the amount guaranteed by a fixed-wage long-term contract. This is especially true in professional sports leagues such as the NBA, where a highly volatile environment limits the ability of the player to accurately forecast his own future performance (Stiroh 2007). For example, for most players, playing time is distributed at the discretion of the team's coach, meaning that there is no guarantee that an individual player will even have the opportunity to accumulate the necessary statistics to meet their contract incentives. Additionally, injury is a common and unpredictable limit on the productivity of professional athletes, and as such, also presents a risk to a player's compensation. Given such high levels of uncertainty, if NBA teams wanted to implement incentive-based compensation schemes, they would have to pay players a much higher average salary to convince them to accept the increase in volatility from riskless guaranteed contracts. We can infer that the corresponding productivity gains from increased player effort would not be enough to offset the increased financial expense of incentive-based contracts for NBA teams.

Stiroh (2007) also considers short-term contracts as an alternative to long-term contracts. If teams are worried that players might increase their effort in contract years and then decrease it after signing a long-term contract, why not make every season a contract year, and motivate players to play hard every year to earn a lucrative deal for the next season?¹⁷ However, this solution of year-by-year contracts is not as ideal as it sounds. There are many advantages to long-term contracts in the NBA, both for the teams and for the players. From the team's perspective, signing a player for the long-

¹⁷ A contract year is defined here as the last season before signing a new contract. It is in the contract year that *ex ante* strategic behavior is hypothesized to occur.

term can provide positive returns in excess of the player's on-court contributions—professional athletes also serve as commercial commodities. By signing popular players, teams can create fan interest and develop loyalty among supporters, which over time results in increased box office revenues, as well as profits from sales of jerseys or other team paraphernalia. In terms of winning games, keeping players for multiple seasons can also have positive effects on team performance. It is commonly accepted—and anyone who has ever played competitive team sports would agree—that players learn to play more effectively with specific teammates over time, a performance-enhancing consequence that results naturally when a core group of players stays with a team for multiple seasons.

From the player's perspective, long-term guaranteed contracts are preferable to short-term contracts for the same primary reason that they are preferable to incentive-based contracts. Simply put, long-term contracts reduce volatility in compensation by smoothing out the risks posed by injury concerns and other unpredictable circumstances. Suppose a player is given a choice between guaranteeing a relatively high salary for the next few years and playing for a new, more profitable contract each season—but at the same time exposing himself to the risk of not earning any new contracts beyond the current season, as in the event of a career-ending injury. It is clear that in most situations the player would choose the guaranteed money over multiple seasons. There are other, less significant concerns that would predispose a player to prefer longer-term contracts, such as not having to make new living arrangements if moving to a new city, building a relationship with the team and city that he is currently playing for, and so forth.

Opportunistic Behavior in the NBA

As mentioned before, the two previous studies on incentive effects of long-term guaranteed contracts in the NBA reached conflicting conclusions. Both provide significant insights on the topic, and the theoretical framework supporting this paper has drawn on components of each previous study. We will begin with a review of Stiroh (2007), in which the author found strong evidence of opportunistic behavior by using both individual statistics and a composite statistic as measures of player performance. We will then examine Berri and Krautmann (2006), in which the authors found weak evidence supporting the shirking hypothesis using the NBA's measure of player productivity, and evidence against shirking when using a measure of each player's marginal productivity.

The motivating hypothesis in Stiroh (2007) is that imperfect information and multi-year contracts create an implicit incentive for workers to strategically alter their effort over the contract cycle. Stiroh defines information asymmetry as the condition under which an employer cannot perfectly distinguish between effort and ability in his employees, but rather sees only an indicator that results from a combination of the two. Using this instrument, the principal must form perceptions about the quality of the agent, and then structure long-term compensation using these incomplete perceptions. From the other side of the strategic game, theory tells us that a rational worker should choose to exert an optimal level of effort in each year of the contract cycle. Returning to the specific case of the NBA, Stiroh continues to explain that players in their contract year should choose their effort level by balancing the gains from higher wages in future contracts with the disutility that arises from increasing effort. Similarly, in the season following the signing of a new contract, players should have incentives to decrease their

effort level. At this point, the cost of exerting effort is not balanced by any marginal benefit, since compensation levels have already been fixed by the terms of the contract. Applying these theoretical results, his goals are to test for two indicators that he believes would support the shirking hypothesis: whether players exert above-average effort in contract years and whether players exert less effort after the new contract is signed.

In his study, Stiroh (2007) uses contract data from the 2000-2001 NBA season and player statistics from 1988 through 2002.¹⁸ He estimates the following weighted OLS regression in order to test his hypotheses:

$$P_{i,t} = \beta_{PRE}PRE + \beta_{POST}POST + \beta_{AGE}NAGE + \alpha_i + \alpha_p + \alpha_t + \alpha_j + \varepsilon_{i,t}$$

P , the dependent variable, is a performance metric that depends on the dummy variables PRE and $POST$, which determine whether or not the performance occurs in a contract year or in the season immediately following. $NAGE$ is a standardized measure of the player's age relative to the league average, and each of the α 's corresponds to dummy variables to control for unobserved individual ability, player position, year, and team.^{19,20} The shirking hypothesis predicts that the coefficients on the dummy variables PRE and $POST$ will be positive and negative respectively, showing that performance increases in the contract year and decreases following the signing of the contract.

The performance metrics employed by Stiroh (2007) are quite elementary, as the regression is estimated several times using individual measures of performance, and once using a composite rating. The following individual statistics were used: points scored, total rebounds, assists, blocked shots, shots attempted, free throws attempted, and

¹⁸ The year 1988 corresponds to the 1988-99 NBA season.

¹⁹ The variable $NAGE$ is used to control for predictable age-related effects (Stiroh 2007).

²⁰ Player position refers to the position that the player primarily plays when he is on the court. Position is relevant because players who play different positions can be expected to accumulate different distributions of statistics. Players in this study were classified as guards, forwards, or centers.

minutes played. The reasoning behind including shots and free throws attempted as dependent variables is that attempts might be a better proxy for player effort, as the player is able to control these statistics more directly. The composite rating used was designed to summarize a wide range of performance statistics and provide a comprehensive measure of overall performance.²¹

Using the composite statistic, Stiroh finds statistically significant evidence of both performance increases in the contract year and performance decreases in the following year. Using the seven independent statistics, he discovers that the data strongly supports an increase in performance during the contract year, but fails to provide evidence for a decrease in production after the contract has been signed, instead showing that player performance regresses to the long-run mean. In other words, these results support the *ex ante* strategic behavior hypothesis, but find no evidence in the data that players also display shirking behavior after signing a new contract.

The Berri and Krautmann (2006) study is motivated by the same hypothesis as Stiroh (2007), namely that opportunistic behavior can occur when there is incomplete information in the labor contract between principal and agent. Berri and Krautmann approach the problem differently in several important ways, beginning with the objective of the research. Instead of searching for evidence of opportunistic behavior both before and after the signing of a contract, they only look for evidence that supports *ex post* shirking behavior.

²¹ The composite rating is calculated as $(\text{field goals made} * 1.4 + \text{blocked shots} * 1.4 + \text{free throws made} * 1.0 + \text{assists} * 1.0 + \text{steals} * 1.0 + \text{offensive rebounds} * 0.85 + \text{defensive rebounds} * 0.5 - \text{turnovers} * 0.8 - \text{field goals missed} * 0.6) / (\text{minutes played} / 48)$.

The model estimated by Berri and Krautmann (2006) is listed below:

$$\Delta PROD = \beta_0 + \beta_1 D2 + \beta_2 D12 + \beta_3 \Delta GP + \beta_4 CEXP + \beta_5 CWPCT + \beta_6 TMWINS + \beta_7 \Delta ROSTER + \theta_1 SHIRKING + \varepsilon$$

The chosen measure of productivity is the dependent variable, regressed on dummy variables $D2$ and $D12$ for the years of experience that the player has in the NBA, the change in games played from the previous season to the current season as a proxy for injury, the years of experience that the coach has, the lifetime winning percentage of the coach, the change in the number of team wins, a variable that accounts for the roster turnover, and the dummy variable that indicates that a player is in their first season after signing a new long-term contract.²² Berri and Krautmann (2006) note that from prior work, they have discovered that player productivity is nearly constant over the course of a player's career, with two notable exceptions. Productivity increases sharply in the first two seasons of a player's career, and then decreases steadily once the player reaches his 12th year of experience. Therefore, a dummy is included for players with more than 2 years of experience, and another for players with more than 12 years. The change in team wins variable is designed to account for the impact of teammate quality on player productivity. The assumption is that a player's productivity falls with increasing teammate quality, as his teammates contribute a greater portion of the team's overall production.

²² The $\Delta ROSTER$ variable was calculated by first determining which players were on the team in both the previous season and the current season, adding up the minutes played by those players in both seasons, and then dividing by the total number of player minutes available in both seasons. This measure was developed in Berri and Jewell (2004).

The term designated *SHIRKING* was tested by Berri and Krautmann using three different variables. In the first model, the simple *SIGNED* variable was used, a dummy equal to 1 when a player was in the season immediately following the signing of a new contract. Next, they tested the interaction with contract length ($SIGNED \times LENGTH$), and the interaction with salary ($SIGNED \times SALARY$). If the *ex post* shirking hypothesis is supported by the data, then we would expect to see a negative coefficient on whichever variable was being used for the *SHIRKING* variable in the test. As stated previously, Berri and Krautmann (2006) found weak evidence supporting the shirking hypothesis when using the NBA's measure of player productivity, and evidence refuting the shirking hypothesis when using a marginal productivity measure, with the additional result that the coefficients on *SIGNED* and ($SIGNED \times LENGTH$) were significant, while the coefficient on ($SIGNED \times SALARY$) was not. From their study, they concluded that signing a new long-term contract would decrease player productivity by about 2-4%, while increasing the size of the contract signed had no effect on productivity.

Using the marginal product method of measuring player efficiency, the evidence obtained by Berri and Krautmann (2006) refutes the shirking hypothesis. They conclude that in the argument over whether opportunistic behavior takes place in professional basketball, the answer is largely dependent on the measure of player productivity chosen. However, they hint that the conventional methods of measuring player efficiency, such as that employed by the NBA, may be lacking in their ability to fully capture a player's contributions to his team's success. From their implication that the marginal product metric better represents the productivity of an individual player, we can infer that Berri

and Krautmann believe that shirking behavior in the NBA is not supported by the available data.

III. Theoretical Framework

Testing for Strategic Behavior

The underlying foundation of this paper is the theoretical prediction that the principal-agent model present in the NBA provides incentives for players to strategically alter their effort and therefore their performance over the course of the contract cycle.²³ Testing for the presence of *ex ante* strategic behavior and *ex post* shirking, terms introduced by Maxcy et al. (2002), will be the main goal of this study. In theory, players will exert more effort in the season prior to signing a long-term contract; their incentive for doing so is the promise of higher compensation over the duration of their new contract as a result of their increased productivity. In the following season, we would expect performance to drop; since their salary is no longer tied to their contemporaneous production, players lose the incentive to play hard and give maximum effort. I will follow Stiroh (2007) in testing for the existence of both *ex ante* and *ex post* strategic behavior, unlike Berri and Krautmann (2006), which only tested for shirking behavior. Evidence from Stiroh (2007) hinted that the *ex ante* effect of increased performance before a new contract may be stronger than the *ex post* effect, suggesting that a statistically significant relationship may be easier to uncover if we test for both instances of performance variation.

Choosing a Model of Player Productivity

²³ The straightforward hypotheses of the principal-agent model in the presence of long-term contracts will only be reviewed briefly in this section, as they have already been discussed at length in previous sections.

The problem with picking a single model of player productivity is that it is unlikely that we will ever find a combination of performance metrics that perfectly measures the level of effort expended by an individual player.²⁴ Thus, we must choose some productivity measure that is a function of effort, among other variables, and then attempt to control for those other variables to the best of our ability. The remaining variation in productivity—provided we have chosen a valid and thorough set of controls—can be attributed to fluctuations in player effort throughout the contract cycle. Furthermore, increases and decreases in productivity in the *ex ante* and *ex post* seasons, if found, will serve as evidence supporting the existence of strategic behavior in the NBA.

Although we must control for additional variables regardless of our choice of productivity measure, I assume here that players will expend effort in a way that helps their team win. If this assumption is largely true, then by picking a measure of productivity that is related to team success, we hope to find a reliable proxy for effort. In this way, we can ensure the consistency and accuracy of the results without an overreliance on the selection of perfect controls.

As a starting point, I began by comparing the productivity measures used in Berri and Krautmann (2006). Using the NBA measure of player efficiency as the performance measure, Berri and Krautmann found statistical evidence supporting the shirking hypothesis at the 5% significance level. However, this measure of productivity may have several deep flaws. The NBA measure of player efficiency is calculated as follows:²⁵

$$PROD_{NBA} = (PTS + TREB + STL + BLK + AST) - (TO + FGMS + FTMS)$$

²⁴ Of course, if this were possible, the hypothetical perfect measure would directly indicate whether players vary their effort level over the contract cycle, and this study would be unnecessary.

²⁵ The variables in this equation are defined as follows: *PTS* = points scored, *TREB* = total rebounds, *STL* = steals, *BLK* = blocked shots, *AST* = assists, *TO* = turnovers, *FGMS* = field goals missed, and *FTMS* = free throws missed.

We immediately see that there are weaknesses in using this measure of performance. Because all of the statistics used to construct this overall measure of player efficiency are weighted equally, we cannot discern the relative value of any single statistic with respect to any of the others.²⁶ In an attempt to develop a more robust measure of player productivity, Berri and Krautmann turn to an economics-inspired model that relies on the player's marginal contribution to team winning percentage.

Berri's approach to measuring player productivity hinges on the idea that team performance can be connected to player statistics through marginal productivity. This idea is developed in Berri (1999), Berri and Jewell (2004), and several other sources, and then reevaluated in Lee (2008). The general idea is to express team winning percentage as a function of points per possession employed and points given up per possession acquired. By differentiating this equation with respect to the independent variables, it is possible to determine the marginal impact of each individual player statistic. For example, through this method, it is possible to calculate the contribution toward winning of a single offensive rebound by a specific player. Summing all the individual statistics that impact team performance, we recover a formula that is very similar to the productivity measure used by the NBA:²⁷

$$PROD_{MP} = (PTS + TREB + STL) - (TO + FGA + 0.44FTA)$$

The key difference is that this measure, derived from marginal productivity, punishes players for shooting inefficiently. By the NBA's measure, a player can boost his performance just by taking a large number of shots, but in this measure, each shot attempt

²⁶ For example, we would not expect a missed field goal to have the same effect as a missed free throw.

²⁷ For a more detailed explanation of this derivation, refer to the Appendix of Berri and Krautmann (2006).

deducts from the player's productivity rating, forcing him to make a higher percentage of his attempts.²⁸

Adhering to the assumption that players funnel effort into production that increases their team's likelihood of winning, the marginal productivity measure of player efficiency used by Berri and Krautmann (2006) provides the basis for the measure of productivity that I will use to estimate my regressions. In order to compare player productivity across different seasons, during which players may not always play the same number of minutes, I will adjust marginal productivity to account for playing time. The final measure of productivity used in my base regression will be the player's marginal productivity per 48 minutes (*MP48*), the duration of a regulation NBA game.²⁹

Choosing Independent Variables and Controls

To control for any factors that might influence a player's productivity other than the contract cycle, it is important to choose appropriate independent variables for the model. As before, I rely on the previous literature to guide my choices of independent variables. To begin with the most important, a player's age and experience will clearly have some kind of impact on his productivity. Surprisingly, both of the previous studies included either age or experience in their regressions, but neither included both.

Although no explicit explanation is given, I surmise that the authors assumed that age and experience would be highly correlated, and that including either one would largely

²⁸ Using the NBA's measure of player efficiency, in order for a player to improve his rating, he must attempt more field goals while making them at rates above 33% for 2-point FGs and 25% for 3-point FGs. Since most players in the NBA shoot above those percentages, a player need only take more shots to improve his productivity. Using the marginal product method, which punishes players for the shots that they take instead of just the shots that they miss, a player must make at least 50% of 2-point FGs and 33% of 3-point FGs to improve his rating.

²⁹ Since player statistics are recorded on a per-game basis, I will divide the calculated marginal productivity by the number of minutes played per game, and then multiply by 48.

capture the effects of the other as well.³⁰ However, in theory, age and experience should have quite different effects on player productivity.

Age is representative of a player's physical condition, which includes factors such as athleticism, accumulated injuries, daily wear and tear, and the natural aging process. As such, I expect the relationship between player age and performance to be generally parabolic, with performance increasing with age up until some critical age—the average age at which a player reaches his physical prime—and then decreasing as a player begins to lose athleticism and feel the effects of the figurative “miles” that he has put on his body over the course of his career.

On the other hand, I expect that experience falls under the category of “you can never have too much of a good thing.” If age is controlled for separately from experience, I would anticipate the effect of experience on productivity to be monotonically increasing, although probably with decreasing returns. Upon entering the NBA, it makes sense that a player would learn and improve quickly, with marginal improvement slowing over his career, as there remains less and less to be gained from additional experience. As with age, experience should have a nonlinear relationship with productivity, so both will be included in the model with first and second order terms.³¹

One variable from Berri and Krautmann (2006) that I chose to include in my model is the change in games played from one season to the next. This variable will serve as a proxy for injuries, which have a definite impact on player performance. I

³⁰ Stiroh (2007) chooses to include player age in his regression using the *NAGE* variable, while Berri and Krautmann (2006) include experience using the variables *D2* and *D12*. It is possible that these variables are designed to account for the effects of both experience and age together (more likely to be the case with *D2* and *D12*).

³¹ Age will be controlled for with variables *AGE* and *AGE*² while experience will be controlled for with variables *EXP* and *EXP*².

expect that if a player is healthy for a larger portion of one season than the previous season, his performance should also increase, as he generally will not have to deal with minor injuries that may not stop him from playing, but can still inhibit his productivity.

Finally, it is important to control for the player's team, the season in which the performance measure was recorded, and individual player characteristics. Different teams in the NBA can have widely varying styles of play, not to mention differences in coaching, so a player's current team during a given season must also be accounted for.³² Similarly, league-wide characteristics such as pace and style may also change from season to season, so it may be important to control for the season in which player performance is recorded. Individual players will of course have different characteristics; for example, talent level, commitment, career ambition, and let's admit it, greed, may all differ among players. We expect there to be variation in productivity between players who play different positions, so player position must be accounted for as well.³³

Confounding Variables / Potential Difficulties

Stiroh (2007) discusses at length two confounding variables that he believes will work against the incentive effects and perhaps diminish the strength of his results. First, there is a selection effect that determines the population of players who are signed to long-term contracts. It is in the best interest of each team to make a significant effort to determine which players are truly high-ability—and thus deserving of long-term contracts—and which are merely opportunistically increasing their effort level in their contract season. If teams are at all successful in distinguishing between ability and effort, then the players being signed to long-term contracts will disproportionately be those who

³² Players can be traded in the middle of a season, so in order to determine a player's team for a particular season, I simply chose the team for which the player played the majority of his games in that season.

³³ Players are categorized under one of five positions: C, F-C, F, G-F, and G.

have high levels of natural ability. This type of high-ability player will have more potential than a similarly productive player who is producing at that level because of increased effort, and this potential will likely be expressed through higher levels of future production, even if effort levels still fall after the signing of the contract. The idea is that teams try to sign high-ability players who have the potential to continue improving, a strategy that, if successful, will work against the hypothesized decline in performance in the year following the new contract.

Another confounding effect might stem from players' concerns about their careers, specifically about future compensation. Although compensation may be guaranteed for the duration of the current contract, all players, particularly younger players, must keep in mind that they will likely be signing new contracts within the next few years. Even though their recent effort and performance will not affect their current pay, it could have adverse affects on their reputation or perceived ability level, which would result in a lower-valued subsequent contract than they could otherwise earn. Stiroh points out that this effect leads to two additional testable hypotheses: variation in effort should be correlated with both the length of the contract signed and the age of the player. The longer the contract, the less the player has to worry about negatively impacting future contracts, and the older the player is, the smaller the proportion of his career he has left to worry about.

It is important to note that both of these confounding effects, the selection effect and future career concerns, are hypothesized to act in the opposite direction of *ex post* shirking. Thus, we theorize that evidence for performance increases during contract

seasons may be easier to uncover than evidence for shirking in the season following the signing of a new contract.

IV. Data

Both Stiroh (2007) and Berri (2006) take their contract data from the USA Today website, as do I. This source provides the annual salary, the total value of the contract, the length of the contract, and also the end date of the contract for each player in the NBA. The database contains data from the 2001-02 season through the 2008-09 season.³⁴ One shortcoming of these data is that they do not differentiate between the guaranteed portion of the contract and any incentives that might be included. However, incentives are usually a small portion of the total contract, if they are included at all, so the results should not be greatly affected.

I plan to use contract data from the 2008-09 season to determine the players who will be included in my sample. The USA Today database contains data for 463 players who were signed to an NBA contract by any of the 30 different teams at some point during the season; I will keep the ones under multi-year contracts for my empirical study. In this case, multi-year is defined as two or more seasons.³⁵ After eliminating players under contract for only one season, the dataset was narrowed to 397 players.

Furthermore, players who were currently in their first NBA contract will not be included

³⁴ If only one year is given to designate a season, it will be the year in which the season began. For example, the 2008-09 season would be referred to as the 2008 season.

³⁵ Players signed to shorter term contracts, specifically one-year deals, may be inherently different from those players who sign multi-year contracts. These differences may lead to complications in analysis, and difficulties in drawing conclusions. For example, many players signed to one-year contracts do not play in many games, or if they do, receive far fewer minutes. As such, they may have a disproportionately large impact on the results of regressions in which they are included.

in the dataset, as shirking behavior cannot be checked for without the existence of pre-contract, NBA statistics.³⁶ After cross-checking the list with data from various sources and eliminating players with incomplete or inconsistent records, I finally arrived at a list of 231 players with contract characteristics that are ideal for this study.³⁷

The top panel of Table 1 reports contract data summary statistics for the 231 players with multi-year contracts. The mean contract length is 4.409 years, with a standard deviation of 1.436 years. The longest contract in the dataset was for 7 years, and the highest total value was a 7-year, \$136 million deal for Kobe Bryant with the Los Angeles Lakers. The average annual salary for an NBA player was \$7.76 million, with standard deviation of approximately \$4.88 million.

Table 1: Summary statistics

	Observations	Mean	Standard Deviation	Minimum	Maximum
Contract data					
Ex ante season	231	2005.17	1.68	2001	2010
Ex post season	231	2006.17	1.68	2002	2011
Contract end season	231	2009.58	1.29	2007	2013
Contract length	231	4.41	1.44	2	7
Total value (millions)	231	\$37.00	\$27.56	\$2.00	\$136.40
Annual salary (millions)	231	\$7.76	\$4.88	\$1.00	\$20.67
Player data					
Age	1864	25.61	3.55	18	36
Experience	1864	5.09	3.17	1	17
ΔGP	1864	0.41	20.33	-77	80
MP48	1864	8.24	4.57	-48	49.15

Player performance statistics have all been acquired from www.basketball-reference.com, a Sports Reference LLC site that keeps an extensive database of NBA

³⁶ Generally this subset of players is composed of players still in their rookie contracts, international players, or players called up from the NBDL (National Basketball Developmental League). It is important that players were active in the NBA prior to their current contract, as statistics cannot be accurately compared across leagues (NCAA, European leagues, NBDL, etc.).

³⁷ I had also planned on estimating my model using players under contract for three or more years, in case a two-year contract cycle was not long enough to provide the necessary incentive for shirking behavior. After running both datasets and finding nearly identical results, I decided to use the full dataset of players with multi-year contracts exclusively.

statistics. Statistics for each individual player over the course of his career are organized by season, in a per-game format. Combining these with the USA Today data, I have a panel including the performance and contract data for every season of each of the 231 players' careers. The panel is sorted by player and season, so each individual observation will be referred to as a player-season, and will include performance statistics and contract cycle data for that player, in that season. The lower panel of Table 1 reports some of the relevant summary statistics for individual player characteristics and productivity. In total, there are 1864 player-season observations, so the average player in the dataset has been in the NBA for approximately 8 seasons.

V. Empirical Specification

My base regression is similar to the one described in Stiroh (2007) in that it has the same objective of testing for both *ex ante* and *ex post* strategic behavior in the NBA labor market. Testing for the possibility of increased effort in the contract year provides an advantage over the model used in Berri and Krautmann (2006), which only tested for *ex post* shirking behavior, since Stiroh found that *ex ante* strategic behavior had a larger effect on player productivity. The change in games played variable from the Berri and Krautmann (2006) model will be added as a proxy for injury, which was unaccounted for in the Stiroh (2007) model. Unfortunately, this proxy is unable to distinguish between games missed for injury and games missed for other reasons, such as league and team suspensions, personal matters, or coaching decisions. However, suspensions are rare in the NBA, and games played as a result of coaching decisions may largely be controlled for by the individual player fixed effect, meaning that the effects of changes in the

number of games played result primarily from injury patterns. I also plan to use several new variables unaddressed by either of the previous studies on this topic. As stated before, experience and age can mean quite different things in the context of the NBA, so separate controls are needed, and since I expect the effects of age and experience to be nonlinear, second order terms are included for each. The base form of the regression that I plan to estimate is the following fixed effects model:³⁸

$$P_{i,t} = \beta_0 + \beta_{ANTE} ANTE + \beta_{POST} POST + \beta_{AGE} AGE + \beta_{AGE2} AGE^2 + \beta_{EXP} EXP + \beta_{EXP2} EXP^2 + \beta_{GP} \Delta GP + \alpha_i + \alpha_t + \alpha_j + \varepsilon_{i,t}$$

The α 's correspond to fixed effects controlling for unobserved individual ability, and dummy variables for season and team respectively.³⁹ A fixed effects model was chosen to account for any unobserved differences in ability between individual players. For example, it could be the case that unobserved player ability is correlated with contract length, meaning that players of varying ability levels in our sample would exhibit systematic differences in their contract characteristics. Such systematic variation could interfere with our capacity to attribute changes in player effort level solely to varying incentives throughout the contract cycle.

The dependent variable used in my base regression, $MP48$, will be similar to the $PROD_{MP}$ developed by Berri, with the adjustments mentioned in Section III. The productivity measure will be adjusted to reflect a per-minute rate of productivity, as players often switch teams when they sign new contracts; changing teams also changes players' roles and their playing time, which could potentially interfere with the effects

³⁸ Here, the dependent variable P represents the measure of productivity chosen for each individual model.

³⁹ Fixed effects and random effects models were both estimated, with a Hausman test showing that the random effects model was not consistent. All following models will be estimated using fixed effects.

that we hope to measure. The per-minute productivity will then be multiplied by 48, reflecting the marginal product that each player would contribute to his team over the course of a full 48-minute game. Figure 1 shows a histogram of the distribution of *MP48* values for all player-seasons:⁴⁰




Figure 1: MP48 Distribution

Using this playing-time-adjusted measure of marginal productivity, I hope to provide a clearer answer to the question of whether or not long-term guaranteed contracts in the NBA provide incentives for players to vary their effort levels strategically over the course of the contract cycle.

A major advancement of this paper with respect to the existing literature is the inclusion of age and experience as separate inputs in estimating player productivity. As

⁴⁰ As the distribution of MP48 values looks roughly normal, I see no reason to take logs or otherwise scale or normalize the values before running the regression.

stated before, high correlation between the two presents a problem of multicollinearity.⁴¹ The breakthrough comes in devising a way to measure player experience that does not rely heavily on the passage of time. For each player-season observation, instead of using the number of seasons played in the NBA, I record the total number of minutes played up until that point in the player's career. By dividing first by the number of minutes in a game (48) and then by the number of games in a full season (82), I arrive at the equivalent number of full seasons of experience if each player had played every minute of every game.⁴² Using this full season equivalent definition of experience has several benefits. It not only breaks down the correlation between age and experience, but also measures the experience gained by a player actually *playing* in the NBA, not just the amount of time that the player has spent employed by an NBA team. It is more intuitive to assume that players benefit from varying amounts of experience as a function of the number of minutes that they play, rather than believing that all players accumulate similar gains each season, regardless of playing time.

Since I expect the overall shapes of the graphs of productivity versus age and experience to be downward-opening parabolas, the coefficients on *AGE* and *EXP* should be positive while the coefficients on the corresponding squared terms should be negative. For values within our dataset range of 1 to 17, I expect that performance should be monotonically increasing with respect to experience. For age, I expect to see a maximum in performance somewhere within our age range of 18 to 36, representing the age at

⁴¹ Most players in the NBA enter the league at around the same age, and experience is typically measured by simply counting the number of seasons that the player has been in the league. Clearly, age and experience by this measure increase hand in hand, leading to a correlation between the two of $\rho = 0.874$.

⁴² In a simple example, if Player X had played a total of 7872 minutes over a 6 season career, *EXP* would be recorded as 2, since $7872 = 2 * 82 * 48$. Although he has been in the NBA for a total of 6 seasons, by my measure of experience, Player X has played the equivalent of 2 full seasons.

which players tend to reach their physical peak. Finally, for the ΔGP variable, I expect the coefficient to be positive, as an increase in games played should correspond to greater health, a condition under which we expect the player to be more productive.⁴³

The driving hypotheses behind this paper are that the coefficients on the *ANTE* and *POST* variables will be positive and negative respectively. Such results would confirm that NBA players do indeed demonstrate increased productivity in the contract season and decreased productivity after signing a long-term contract. Before estimating any regressions, I ran a few simple summary statistics on the dependent variable (*MP48*) to determine if these hypotheses were reasonable. Table 2 below contains these statistics, sorted by player position:⁴⁴

Table 2: MP48 by Player Position

Position	Ex Ante MP48		Ex Post MP48		All Other Seasons	
	Mean	SD	Mean	SD	Mean	SD
C	12.877	3.326	12.252	4.174	12.216	4.508
F	10.337	2.909	9.725	3.079	9.389	3.247
F-C	12.872	2.980	12.039	3.090	11.431	5.435
G	5.353	2.330	4.863	2.436	4.692	2.711
G-F	6.985	2.110	5.920	3.872	6.349	2.716

As we can see, in all cases, the mean marginal productivity is higher in the *ex ante* season than in the *ex post* season, while both *ex ante* and *ex post* averages are higher than the averages in all other seasons, except in the case of G-F.⁴⁵ This leads us to believe that, after controlling for experience, age, and other independent variables, we may find evidence supporting the existence of strategic behavior in the NBA. The results in this table seem to strengthen my belief that contract year increases in performance may be

⁴³ If the change in games played is negative, then since the coefficient is positive, we will see a decrease in predicted productivity, which follows our intuition.

⁴⁴ SD columns list standard deviations.

⁴⁵ The category All Other Seasons refers to any season that is not an *ex post* or *ex ante* season.

more prevalent than shirking behavior. One potential concern is the obvious systematic differences in productivity between players who play different positions. Although player position will be controlled for in the fixed effects model, it is possible that such a large difference in productivities across positions is a sign that *MP48* does not accurately capture different players' contributions towards their team's success.

Table 3: Fixed Effects Models with Different Controls Using MP48 as Dependent Variable

Variables	(1) Age and experience	(2) Age only	(3) No season or team controls	(4) Season only	(5) Team only
Ante	0.684*** (0.202)	0.685*** (0.202)	0.666*** (0.195)	0.649** (0.201)	0.677*** (0.195)
Post	-0.049 (0.203)	-0.044 (0.203)	-0.071 (0.200)	-0.049 (0.202)	-0.076 (0.202)
Δ GP	0.020*** (0.003)	0.020*** (0.003)	0.019*** (0.003)	0.020*** (0.003)	0.019*** (0.003)
Age	0.787* (0.395)	1.089*** (0.303)	1.055** (0.348)	0.876* (0.383)	1.005** (0.359)
Age ²	-0.018** (0.006)	-0.022*** (0.005)	-0.018** (0.006)	-0.019** (0.006)	-0.018** (0.006)
Exp	0.303 (0.269)		0.288 (0.261)	0.292 (0.265)	0.282 (0.265)
Exp ²	-0.024 (0.020)		-0.024 (0.019)	-0.023 (0.019)	-0.024 (0.020)
Season controls	Yes	Yes	No	Yes	No
Team controls	Yes	Yes	No	No	Yes
Constant	-0.152 (7.144)	-4.386 (6.089)	-6.980 (4.910)	-1.338 (6.947)	-6.579 (5.112)
Observations	1864	1864	1864	1864	1864
Number of players	231	231	231	231	231
R-squared	0.107	0.106	0.070	0.077	0.099
Adj. R-squared	-0.051	-0.051	-0.066	-0.067	-0.051

Standard errors in parentheses
 *** p<0.001, ** p<0.01, * p<0.05

Table 3 above reports the results of my base regression, along with the results for several other regressions with different sets of controls. Starting with the base regression in column (1) of the table, the results seem to support the hypothesis that players increase performance in the *ex ante* season, although there is no evidence that they decrease

productivity in the following season. Examining the variables of primary interest, we can see that the coefficient on the *ANTE* variable has a positive coefficient and is statistically significant at the 0.1% level. This confirms our hypothesis that players hoping to sign a new contract will tend to outperform, exerting maximum effort in order to lock in the highest possible salary. The evidence for *ex post* shirking behavior is not nearly as strong: We do find a negative coefficient on the *POST* variable in each case, as expected, but the magnitude of the effect is not close to being statistically significant.

These findings agreed with those of Berri and Krautmann, and might be explained by the confounding effects described in the theoretical framework. In theory, the *ex post* shirking effects might be overcome by two confounding variables acting in the opposite direction. First, teams attempt to sign players that they believe have the potential to continue performing at a high level, or even to improve, which would lead to an increase in productivity after the new contract is signed. Secondly, each player also has an economic incentive to perform well after signing a new contract. Unless the player knows that the current contract will be the last of his career, he must remember that his present productivity can affect his reputation and future valuation. Both of these confounding effects would work against the shirking hypothesis, so perhaps the incentives to shirk are balanced out by effective player valuation by NBA teams and players' awareness that present production may dictate future compensation.

Looking at the other independent variables in the base regression, we see that the calculated ΔGP , AGE , and AGE^2 variables are also statistically significant. The coefficient on ΔGP is positive, as we expected; players are likely to be more productive when they spend less time injured and on the bench.

In order to understand the effects of age and experience on productivity, I start with the results of column (2), which are based on a model in which only age is used to account for the combined effects of age and experience.⁴⁶ By including AGE and AGE^2 , I expected the graph of productivity versus age to be an inverted parabola, reaching a maximum within the domain of the sample.⁴⁷ Qualitatively, this maximum should correspond to the age at which a combination of the player's athleticism, knowledge of the game, experience in the NBA, and overall physical condition allow him to reach his peak productivity. Figure 2 graphs productivity against age for this model, which does not include experience as a separate variable:




Figure 2: Productivity vs. Age⁴⁸

The effect of age on productivity peaks around 24.6 years, which is a reasonable estimate of when most NBA athletes reach their physical prime. After this point, the marginal decrease in productivity as a result of declining athletic abilities or wear and

⁴⁶ This model, with only age, is similar to both Stroh (2007) and Berri and Krautmann (2006) in that the effects of age and experience are explained by variation in only one of the two highly correlated variables. Unlike either of the previous studies, however, I allowed for the possibility of a nonlinear relationship between age and experience and productivity.

⁴⁷ Players in the sample ranged from ages 18 through 36.

⁴⁸ The gain in productivity with respect to age follows the graph $y = 1.089*x - 0.022*x^2$. The domain of the graph includes values of experience in the interval [15,40]. This graph represents the model which includes only age and not experience.

tear from the daily grind of practicing and playing a professional sport begins to outweigh the positive but decreasing marginal returns from the accumulation of experience.

Although this simpler model seems to capture the combined effects of age and experience fairly accurately, one of the goals that drove the specification of my base regression was to create a new model of player productivity in which the entangled effects of age and experience on productivity could be identified and separated. By measuring experience with the equivalent number of seasons played by each player rather than the total number of seasons the player has been in the NBA, I can include both age and experience in my base model, which allows for more specific insight into their individual effects on productivity. Figure 3 at the top of the following page graphs productivity against age in the base regression, which also includes experience as an input of productivity.

Again, the results closely match our theoretical expectations. According to the model, productivity as a function of age increases until it peaks around 22.21 years, and then decreases at an increasing rate as the player continues to age. Comparing this result to the previous model, which only includes age, it is quite believable that the average player actually reaches his maximum athleticism or physical prime around 22 rather than 25, as the effect of age on productivity is now isolated from productivity gains due to increases in experience.⁴⁹

⁴⁹ When age is used to account for the effects of both age and experience, the increases in experience that begin upon the player entering the NBA have a monotonically increasing effect on production, so it is not until the player is older that the negative effect of aging begins to overcome the positive effect of gaining experience. In the base model, age captures the pure physical effect of aging, so as soon as the player begins to lose physical ability, it is reflected through decreasing productivity with respect to age.




Figure 3: Productivity vs. Age⁵⁰




Figure 4: Productivity vs. Experience⁵¹

The effect of experience on productivity in the base regression is displayed in Figure 4 above. Since the players in the sample have equivalent levels of experience, ranging from 0 to 11 seasons, it is clear that productivity is not monotonically increasing in experience as postulated earlier. However, it may be that there are again two factors at

⁵⁰ The gain in productivity with respect to age follows the graph $y = 0.787*x - 0.018*x^2$. The domain of the graph includes values of age in the interval [15,40].

⁵¹ The gain in productivity with respect to experience follows the graph $y = 0.303*x - 0.024*x^2$. The domain of the graph includes values of experience in the interval [0,11].

work here. On the one hand, there is the positive effect that experience has on players both mentally and psychologically, which could in fact continuously increase with experience.⁵² At the same time, every additional minute of experience comes with a cost. Eventually, the player's body will start to feel the day-to-day attrition of playing professional basketball, as well as the accumulated effects of injuries sustained, both major and minor. This process is separate from that of the natural aging process, as it is directly related to the amount of time that the player has spent on the basketball court. At the quantity of equivalent seasons when the negative effects of physical deterioration begin to overcome the diminishing mental and psychological returns of increasing experience, productivity as a function of experience peaks. In the base model, this peak occurs when the player has played the equivalent of approximately 6.38 full seasons in the NBA, or about 523 complete games.

In models (3) through (5), I remove one or both of the team and season controls, finding that the removal of these controls from the model has no significant effect on the results. For all subsequent models, both season and team controls are included. Furthermore, note that the coefficients and standard errors on the *ANTE* and *POST* variables are almost exactly the same in the base regression and the regression using only age to account for the combined effects of age and experience. As neither model seems to have an advantage in consistency or efficiency, I use the base regression for the remainder of the paper, since it is supported by the theoretical framework and can offer more specific insight into the factors affecting productivity.

⁵² As players play more cumulative minutes in the NBA, one would assume that they gain a greater understanding of basketball and familiarity of the league, allowing them to make better decisions, handle additional pressure, and more capably evaluate their opponents, all of which would likely lead to increases in productivity.

**Table 4: Fixed Effects Models with Post Interactions
Using MP48 as Dependent Variable**

Variables	(1) Base	(2) Post x Length	(3) Post x Age	(4) Post x Length x Age
Ante	0.684*** (0.202)	0.683*** (0.202)	0.682*** (0.202)	0.684*** (0.202)
Post	-0.049 (0.203)	0.209 (0.664)	-0.521 (1.696)	0.256 (0.675)
ΔGP	0.020*** (0.003)	0.020*** (0.003)	0.020*** (0.003)	0.020*** (0.003)
Age	0.787* (0.395)	0.791* (0.395)	0.806* (0.401)	0.784* (0.395)
Age ²	-0.018** (0.006)	-0.018** (0.006)	-0.018** (0.006)	-0.018** (0.006)
Exp	0.303 (0.269)	0.312 (0.270)	0.302 (0.269)	0.314 (0.270)
Exp ²	-0.024 (0.020)	-0.024 (0.020)	-0.024 (0.020)	-0.024 (0.020)
Post x Length		-0.057 (0.140)		
Post x Age			0.018 (0.062)	
Post x Length x Age				-0.003 (0.005)
Season controls	Yes	Yes	Yes	Yes
Team controls	Yes	Yes	Yes	Yes
Constant	-0.152 (7.144)	-0.165 (7.146)	-0.377 (7.191)	-0.076 (7.148)
Observations	1864	1864	1864	1864
Number of player	231	231	231	231
R-squared	0.107	0.107	0.107	0.107
Adj. R-squared	-0.051	-0.052	-0.052	-0.052

Standard errors in parentheses
*** p<0.001, ** p<0.01, * p<0.05

In models (2) through (4) of Table 4 above, I add independent variables for the *ex post* variable interacted with contract length, player age, and both contract length and age, respectively. The underlying intuition is that for longer and longer contracts, a player in his *ex post* season has less and less to worry about in terms of signing a subsequent contract. This should provide further incentive for him to decrease his effort. Similarly,

for age, an older player would expect to be spending less time in the NBA in the future than a younger player, all else equal, so increasing age would also increase the incentive to shirk. In all three cases, the coefficients on the interaction terms were within the margin of error, so the results show no evidence that either contract length or player age has an effect on the propensity of a player to shirk in the *ex post* season.

After running my base regression using marginal productivity per 48 minutes as my dependent variable, I decided to conduct a sensitivity test by substituting several different measures as my dependent variable. Using the independent variables of my base regression, I tried to predict successively marginal productivity, the NBA measure of efficiency, points, total rebounds, assists, field goals attempted, and free throws attempted.⁵³ Table 5 on the following page reports the results.

In models (2) and (3), where comprehensive productivity measures of marginal productivity per game and NBA efficiency per 48 minutes are used, the coefficients on *ANTE* are positive and significant at the 1% level.⁵⁴ However, unlike the base regression with marginal productivity per 48 minutes as the dependent variable, the coefficient on *POST* is positive as well, though not significantly different from zero. Models (4) through (8), which used individual statistics as the dependent variables, found varying results. Evidence for contract year strategic behavior was found to be significant at the 1% level for *PTS48* and at the 10% level for *TRB48*, *FGA48*, and *FTA48*.⁵⁵⁵⁶ The coefficient on *ANTE* is slightly negative for *AST48*, though not statistically significant.

⁵³ This was the marginal productivity unadjusted for minutes played, as MP48 was already used in the base regression. Each of the other variables, NBA efficiency, points, total rebounds, FGA, and FTA, was adjusted to a per-48 minutes basis.

⁵⁴ MP = marginal productivity per game; NBA48 = NBA efficiency per 48 minutes.

⁵⁵ PTS48 = points per 48 minutes. The other individual statistics are denoted similarly.

⁵⁶ FGA and FTA were hypothesized to be good proxies for player effort, as a player would seemingly have more direct control over how many shots he takes than over how many rebounds he is able to get.

There is no statistical evidence of *ex post* shirking behavior in any of the estimated regressions. Since similar results are obtained using several different measures of player productivity, this sensitivity test increases my confidence in the robustness of my model.

**Table 5: Fixed Effects Models of the Base Regression
Using Different Productivity Measures as Dependent Variable**

Variables	(1) MP48	(2) MP	(3) NBA48	(4) PTS48	(5) TRB48	(6) AST48	(7) FGA48	(8) FTA48
Ante	0.68*** (0.20)	0.46*** (0.11)	1.07*** (0.28)	0.77*** (0.22)	0.20* (0.10)	-0.01 (0.07)	0.31* (0.16)	0.19* (0.10)
Post	-0.05 (0.20)	0.09 (0.11)	0.10 (0.28)	0.07 (0.22)	-0.02 (0.10)	0.08 (0.07)	0.09 (0.16)	0.07 (0.10)
ΔGP	0.02*** (0.00)	0.01*** (0.00)	0.03*** (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	-0.01*** (0.00)	-0.00 (0.00)
Age	0.79** (0.39)	0.59*** (0.22)	1.79*** (0.55)	1.45*** (0.43)	0.17 (0.20)	0.19 (0.14)	0.88*** (0.31)	0.38* (0.20)
Age ²	-0.02*** (0.01)	-0.02*** (0.00)	-0.05*** (0.01)	-0.05*** (0.01)	-0.01** (0.00)	-0.01*** (0.00)	-0.03*** (0.00)	-0.02*** (0.00)
Exp	0.30 (0.27)	1.62*** (0.15)	2.39*** (0.37)	3.79*** (0.29)	-0.25* (0.13)	0.59*** (0.10)	2.24*** (0.21)	1.41*** (0.14)
Exp ²	-0.02 (0.02)	-0.11*** (0.01)	-0.13*** (0.03)	-0.21*** (0.02)	0.02* (0.01)	-0.03*** (0.01)	-0.13*** (0.02)	-0.06*** (0.01)
Constant	-0.15 (7.14)	2.26 (3.95)	5.28 (9.88)	3.84 (7.77)	9.38*** (3.56)	2.37 (2.62)	6.14 (5.68)	2.40 (3.68)
Observations	1864	1864	1864	1864	1864	1864	1864	1864
Number of player	231	231	231	231	231	231	231	231
R-squared	0.11	0.32	0.18	0.29	0.07	0.13	0.25	0.22
Adj. R-squared	-0.05	0.20	0.03	0.16	-0.10	-0.02	0.12	0.08

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

An early concern that I had was with the weak predictive power offered by the fixed effects model. Comparing my low R-squared and adjusted R-squared values with adjusted R-squared values in the 0.7-0.8 range using the weighted OLS method of Stiroh (2007), I was initially wary of my results and concerned with both my methodology and my panel data. I conducted several checks to make sure that each was valid. Estimating the model used by Stiroh on my dataset, after adding dummies to capture individual effects, I was also able to reproduce similar adjusted R-squared values, convincing me that my dataset was not invalid or incomplete in any obvious way. Checking my results

with Berri and Krautmann (2006), I found that they had reported similar R-squared values. Econometrically, it makes sense that a weighted OLS model would better predict individual player productivities than a fixed effects model, as each player will be associated with a unique constant value that sets an appropriate base level of production for that individual. However, with a set of panel data, I choose to use the fixed effects model, rather than random effects or a pooled OLS model, as it is consistent and efficient under the condition that time-invariant individual effects are correlated with the independent variables.

VI. Conclusions

The existing literature on strategic behavior in the context of the National Basketball Association showed inconsistent results using different empirical approaches. By drawing upon and building on work done by Stiroh (2007) and Berri and Krautmann (2006), my research provides greater insight on the theory behind models of player productivity. Additionally, my results also present further evidence to substantiate the claim that players do indeed vary their effort levels in response to changing incentives over the course of their contract cycle, much as economic theory predicts.

Specifically, my model—which held up for a variety of measures of player productivity—provides relatively robust evidence that *ex ante* strategic behavior, or the contract year phenomenon, is quite real. Following the model’s underlying theory, players currently in a contract year have an incentive to increase their productivity because they know that the value of their next long-term contract will be determined in part by their performance in the current season. Thus, they will expend more effort now,

in order to lock in a higher guaranteed salary in the future. On the other hand, no evidence was found to support the shirking hypothesis, corroborating the findings of Berri and Krautmann (2006). *Ex post* performance declines either do not exist or are masked by a variety of confounding factors. For example, there may be a selection effect working in the opposite direction as shirking behavior if teams are succeeding in their goal of identifying players who have the most potential for continued improvement.

A possible extension of this study involves examining the players on individual NBA teams as separate subgroups, and attempting to draw a link between the success of a franchise and the types of players that it signs, relative to the norm. For example, if players on the Knicks show greater variation in their performance level over the contract cycle relative to the league average, are the Knicks more likely or less likely to have a winning team? If a team has many players who are much more productive in their contract year than in other seasons of their contract cycle, it would probably have a negative impact on team success, as these players tend to exert maximum effort only when faced with the proper incentives.

My findings have the potential to benefit NBA teams in their player evaluation process, as it allows them to more accurately determine contract cycle effects on player performance. For example, when trying to decide whether to sign a player to a new contract who is currently in his contract year, my model of player productivity could offer insight into how much of the player's productivity is explained by his being in a contract year, and how much is a function of his individual characteristics. The inclusion of both age and experience as inputs in estimating production should allow for greater accuracy in predicting future productivity, especially of players who do not follow the

typical path to the NBA. These players might not exhibit the usual correlation between age and experience, so modeling the two as separate effects will be much more accurate in these cases. With greater success in player evaluation, teams would be able to save money that they previously would have spent on unworthy players, and parity and competitiveness in the NBA would likely increase. Additionally, teams might apply this research on a player-by-player basis to determine which players would benefit most from the motivation of performance incentives. Guided by the theories of player productivity variation formulated in this paper and working under the inevitable financial constraints imposed by the league salary cap, an NBA team could maximize its aggregate productivity, and thus its overall success, by limiting the contractual incentives for players to engage in strategic behavior.

Works Cited

- Alchian, Armen A., and Harold Demsetz. "Production , Information Costs, and Economic Organization." *American Economic Review* 62, no. 5 (1972): 777-95.
- Basketball-Reference.com. Sports Reference LLC. Available online at <http://www.basketball-reference.com>.
- Berri, David J. "Who Is 'Most Valuable'? Measuring the Player's Production of Wins in the National Basketball Association." *Managerial and Decision Economics* 20, no. 8 (1999): 411-27.
- Berri, David J., Stacey L. Brook, and Martin B. Schmidt. "Does One Simply Need to Score to Score?" *International Journal of Sport Finance* 2, no. 4 (2007): 190-205.
- Berri, David J., and R. Todd Jewell. "Wage Inequality and Firm Performance: Professional Basketball's Natural Experiment." *Atlantic Economic Journal* 32, no. 2 (2004): 130-39.
- Berri, David J., and Anthony C. Krautmann. "Shirking on the Court: Testing for the Incentive Effects of Guaranteed Pay." *Economic Inquiry* 44, no. 3 (2006): 536-46.
- Heubeck, Tina, and Jochen Scheuer. "Incentive Clauses in Players' Contracts in Team Sports - Theory and Practice." *German Working Papers in Law and Economics* 2003 (2003): 1-30.
- Holmstrom, Bengt. "Moral Hazard and Observability." *Bell Journal of Economics* 10, no. 1 (1979): 74-91.
- Kahn, Lawrence M., and Malav Shah. "Race, Compensation and Contract Length in the Nba: 2001-2002." *Industrial Relations* 44, no. 3 (2005): 444-62.

Krautmann, Anthony C. "Shirking or Stochastic Productivity in Major League Baseball?"

Southern Economic Journal 56, no. 4 (1990): 961-68.

Krautmann, Anthony C., and Thomas D. Donley. "Shirking in Major League Baseball:

Revisited." *Journal of Sports Economics* 10, no. 3 (2009): 292-304.

Lee, Young Hoon, and David Berri. "A Re-Examination of Production Functions and

Efficiency Estimates for the National Basketball Association." *Scottish Journal of*

Political Economy 55, no. 1 (2008): 51-66.

Maxcy, Joel G., Rodney D. Fort, and Anthony C. Krautmann. "The Effectiveness of

Incentive Mechanisms in Major League Baseball." *Journal of Sports Economics* 3,

no. 3 (2002): 246-55.

Oyer, Paul. "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business

Seasonality." *Quarterly Journal of Economics* 113, no. 1 (1998): 149-85.

Prendergast, Canice. "Intrinsic Motivation and Incentives." *American Economic Review*

98, no. 2 (2008): 201-05.

———. "The Provision of Incentives in Firms." *Journal of Economic Literature* 37, no. 1

(1999): 7-63.

Stiroh, Kevin J. "Playing for Keeps: Pay and Performance in the Nba." *Economic Inquiry*

45, no. 1 (2007): 145-61.

USA Today.com. "Basketball Salaries Database." Available online at

www.usatoday.com/sports/basketball/nba/salaries/default.aspx.