Moral hazard in long-term guaranteed contracts: theory and evidence from the NBA

J. Bradford Rice
Department of Economics
Boston University

Arup Sen
Department of Economics
Boston University

March 25, 2008

Abstract: Conventional wisdom suggests that offering short contracts to workers will elicit maximum effort, which is why the long term-contracts with guaranteed wages that abound in sports markets seem to lack economic justification. Performance and contract data from an unbalanced panel of NBA players, suggest that effort-related productivity increases monotonically and nonlinearly, by as much as 30% over the duration of a fixed-period contract. We develop a game theoretic model of optimal contract length to show that the insurance aspect of long-term contracts and workers' career concerns can nonetheless overcome the adverse incentive effects..

JEL classification: D86, J41

Contact information:

Brad Rice
Boston University
270 Bay State Road
Boston MA 02215 USA
jbrice@bu.edu

Arup Sen
Boston University
270 Bay State Road
Boston MA 02215 USA
arupsen@bu.edu

Acknowledgment: Thank you Kevin Lang, Barton Lipman and Andrew Newman, for valuable input to this paper. Thank you, also, to our colleagues and Empirical Microeconomics and Microeconomic Theory seminar participants at Boston University, and participants at several other seminars for helpful comments.

1 Introduction

When complete contracts are written (i.e., when contracts specify a menu of payments contingent on documentable levels of effort or performance), longer contracts must Pareto dominate shorter contracts. When markets must rely on incomplete contracts, however, the length of the contract has crucial implications for agents' performance during the contract. When contracts are incomplete and either cannot or do not rely on performance bonuses, firms must rely on the inherent incentives that workers have to perform well in order to obtain the best possible terms in subsequent contracts. In the case of incomplete contracting, long-term contracts will undermine that incentive since the disutility of effort is borne by the worker immediately, while any benefits only begin to accrue at the time of recontracting.

The presence of moral hazard in incomplete contracts suggests that employers will prefer short-term contracts in settings where the length of the contract is part of the negotiation. Labor markets, however, are rife with examples of firms and workers agreeing to incomplete contracts that stipulate a guaranteed wage for a specified multi-year duration. For instance, contracts for junior academic faculty often include a guaranteed salary throughout a standard length of service. In several European countries, the only lawful alternative to non-permanent hiring is a "fixed-term" contract that specifies guaranteed payment over a negotiated duration. This type of fixed-term contract is similar to those signed by many professional team-sport athletes. At the end of these contracts, individuals are evaluated for reappointment and a new contract is negotiated.

In this paper, we develop a simple model of optimal contract length that draws on the inherent trade-off between the incentive effects of short-term contracts and risk aversion. The model shows that while career concerns moderate moral hazard, long-term contracts do contain adverse effort incentives. When firms and workers can bargain over wages, however, risk aversion ensures that workers gain by having long-term contracts. Because workers are willing to make wage concessions, firms can benefit since long-term contracts provide a cheaper way to hire workers. Our results hold without the inclusion of any explicit notion of recontracting costs. The model weds the most salient features of the literature on career concerns (e.g., Holmström 1982, 1999; and Dewatripont, Jewitt, and Tirole 1999) with the relevant literature from long-term contract theory (e.g., Holmström 1983; Cantor 1988; Fudenberg, Holmström, and Milgrom 1990; Fudenberg and Tirole 1990) to find conditions under which long-term contracts Pareto-dominate short-term contracts. The optimal contract length does not always cover multiple periods. Adjusting the parameters in the model helps explain the intraindustry heterogeneity in contract length which oftentimes occurs.

The focus of the empirical section of the paper is to quantify the change in effort incentives within a long-term deal, controlling for confounding factors. To do so, we focus on the agency relationship between National Basketball Association (NBA) players and team ownership. We have compiled a unique dataset containing information on 654 NBA players, their contractual terms, annual performance across several dimensions, information on team performance, and physical characteristics.

Historically, testing for moral hazard in labor markets has proved difficult in practice. As a result, our paper contributes to a literature with limited empirical evidence. The primary reason for this scarcity of evidence is that the data requirements for such tests are restrictive. At a minimum, the researcher must observe micro-level worker performance and contractual terms. Because of the public nature of the data, a relatively large proportion of the work testing the effort implications of contract structures analyzes the agency relationship in sports.[1] The bulk of this literature, however, studies Major League Baseball, whose contracts tend to include more individual incentives and bonuses than professional basketball contracts (see Lehn [1990] for an example or Kahn [2000] for a review of this literature). In addition, because contractual terms are not distributed randomly across workers, and inherent ability is imperfectly measured, it is difficult to disentangle the output effects due to ability and the output effects due to effort. The important feature of our data is that we observe each player for an average of 2 contracts, varying from 1 to 12 years in length. Because over 90% of the individuals in our data are observed in multiple years we are able to control for unobserved individual-specific heterogeneity. The data set we utilize is over four times larger, includes more individuals, more observations per individual, and more data on individual characteristics than any other comparable study. In addition, this paper appears to be the first of its kind to quantify the continuous, within-contract path of effort.

We find strong evidence that the effort of NBA players increases monotonically as their contracts near completion—a pattern that can be explained by our model of contracting. Fixed-effects estimates indicate that effort-related performance (as measured by the NBA efficiency score, which combines several performance statistics into a single index) in the final year of a multi-year deal is approximately 8% higher than in the year prior. Moreover, our point estimates indicate that effort is a non-linear function of the number of years until contract renewal, with performance in the third-to-last year of a contract falling only an additional 4%, and with no statistically significant differences in effort when a player has four or more years remaining in his contract. Estimating the model via OLS without player fixed-effects fails to find any adverse effort incentives associated with long-term contracts, implying that unobserved heterogeneity is an important source of bias.

---

[1] There is a literature on testing for moral hazard in labor markets outside of sports, however. For instance, Lazear (2000) and Foster and Rosenzweig (1994) perform a test for moral hazard in the labor market under different payment schemes. In addition, Gibbons and Murphy (1992) test the empirical implications of the career concerns literature applied to chief executives. While they do not address the moral hazard effects of contracts specifically, they test for differences in contractual structures designed to combat moral hazard.

2 Theory

We start our theoretical framework by setting up a three-period principal-agent model. The model derives three central predictions,

1. effort increases from one period to the next within a multi-period guaranteed contract,
2. cumulative effort generated within a multi-period contract falls short of the effort levels elicited by a series of single-period contracts, and
3. in spite of this 'shirking,' the multi-period guaranteed contract may be Pareto efficient.

The potential efficiency of multi-period contracts is driven by the risk aversion of the agents. Since output in our model is partially determined by a random component, risk aversion motivates the agent to offer wage concessions in exchange for a pre-specified wage, guaranteed over a longer period of time. Multi-period contracts enable the worker to face less income risk.

Overall, career concerns ensure that the agent does put in a positive level of effort (except in the last contract) and that problems of shirking are mitigated. The optimal contract in this framework, however, must weigh the risk smoothing advantages of multi-period contracts with their adverse incentives, much like in Zeckhauser's (1970) classic paper on medical insurance.

The market consists of two players: one principal and one agent. Output produced in period t, y_t, is a stochastic function of the agent's exogenously determined ability, a, and effort, $e_t \in \mathbb{R}^+$, according to the production function $y_t = a + e_t + z_t$. Ability is initially unknown to the principal and the agent, but both share a common prior belief that it is drawn from a normal distribution with mean μ, variance σ_a^2, and a precision parameter defined as $h_a = \frac{1}{\sigma_a^2}$. The stochastic production element, z_t, is independent of a, normally distributed, and i.i.d. with zero mean, variance σ_z^2, and precision $h_z = \frac{1}{\sigma_z^2}$. The cost to the agent of exerting effort is described by a convex cost function, $c(e_t) = e_t^2$.

The agent's total utility in any period t as a function of wage, w_t, and level of effort exerted, e_t, is given by $U(w_t) - e_t^2$. As opposed to most career concerns models, the agent is assumed averse to unexpected income shocks so that the utility that he derives from a monetary reward follows a CARA functional form, $U_t = -\frac{1}{\alpha} \exp -\alpha w_t$. The parameter α is the agent's coefficient of constant absolute risk aversion. There is no borrowing or lending, i.e. the agent must consume w_t in period t. The principal's payoff is determined by the profit function $y_t - w_t$.

We assume that for any length of contract, τ, there is an associated exogenous profit sharing rule, $\kappa_\tau \in 0, 1$ such that $w_t = \kappa_\tau E\bar{y}$, with $E\bar{y}$ representing the average (discounted) expected output within a contract. We interpret this profit share as the result of a bargaining process between the principal and agent. When the principal is part of a perfectly competitive market, $\kappa_\tau = 1$. It is worth emphasizing again that wage in any period must be determined before the start of the period, and can only be based on expected levels of output conditional upon all publicly available information up to that point.

Though the true ability of the agent is unobservable, the market is able to obtain progressively more precise estimates by observing a series of $y_t - e_t^*$, in which e_t^* is the the amount of effort that the principal expects the agent to exert in period t. In a Perfect Bayesian Equilibrium, the principal correctly anticipates the chosen level of effort, hence there is never any information asymmetry. A series of realizations $\{y_s - e_s^*\}_{s=1}^{t-1} = \{a + z_t\}_{s=1}^{t-1}$ gives rise to a normal posterior distribution of a with the following means and precisions:

$$\begin{aligned} \mu_1 &= \frac{\mu h_a + h_z(y_1 - e_1^*)}{h_a + h_z} \\ h_{a_1} &= h_a + h_z \end{aligned}$$

after period one, and

$$\begin{aligned} \mu_2 &= \frac{\mu h_a + h_z(y_1 - e_1^* + y_2 - e_2^*)}{h_a + 2h_z} \\ h_{a_2} &= h_a + 2h_z \end{aligned}$$

after period 2 has transpired. The market uses this updated belief of ability as its basis to predict future output and wages.

With this setup, we compare the results of two possible contract structures. In the first case, we consider an agent who signs a two-period contract followed by a one-period contract—a multi-period contract case. We then consider a scenario in which the agent signs three one-period contracts.

The aim of this exercise is to generate three theoretical predictions that can be tested empirically, as mentioned above: effort within a multi-period contract is not constant, aggregate effort in a multi-period contract is less than that elicited from multiple single-period contracts, and to prove the existence of an optimal multi-period contract. We are not attempting to make a qualitative distinction as to which type of multi-period contact is optimal. For instance, we omit from this analysis the case of the principal and agent agreeing to a lifetime contract (τ=3). Of course, while such a contract may be Pareto preferred if the agent is sufficiently risk averse, the contract would elicit zero effort in every period. Similarly, we forgo analyzing the case of a one-period contract followed by a two-period contract. While the principal and agent may agree to a two-period contract at the start of period 2, this multi-period contract would again have the implication of zero effort in periods 2 and 3. The particular multi-period contract analyzed allows for meaningful comparisons of the optimal effort decisions with the single-period contractual scenario.

Since effort in period t may differ according to changes in current and past contractual terms, let superscripts on effort, e_t^i, and wages, w_t^i denote effort and wage in period t. The superscript $i = \{M, S\}$ denotes the multi-period versus single-period contract case, respectively (e.g., e_2^S denotes second-period effort in the scenario of three single-period contracts). In addition, to call attention to the information on which the agent bases his effort choice, we denote E_X as taking expectations over the random variable(s) X.

2.1 Case 1: Multi-period contract

Given the contract structure, the agent chooses effort to maximize lifetime utility. Since first and second-period wages are fixed, the utility from these wages is known at the time of signing the contract. Because effort in period t will be an optimal response to the entire history of the game, second-period effort will depend on first-period outcomes. Likewise, third-period effort will depend on the outcomes of both periods 1 and 2 though, in this case, it is equal to zero for all y_1 and y_2. The agent's maximization problem is

$$\max_{e_1^M, e_2^M(\cdot), e_3^M(\cdot)} U(w_1^M) - \left(e_1^M\right)^2 + RE_{y_1}\left(U(w_2^M) - \left(e_2^M(y_1)\right)^2\right) + \\ R^2 E_{y_1,y_2}\left(U_3(w_3^M(y_1,y_2)) - \left(e_3^M(y_1,y_2)\right)^2\right)$$

in which $R \in (0,1)$ is the discount factor common to both the principal and agent.

The first-order conditions

$$\begin{aligned} e_1^M &= \frac{R^2}{2}\frac{\partial E_{y_1,y_2}U_3}{\partial e_1^1} + \frac{R}{2}\frac{\partial E_{y_1}e_2^M(y_1)}{\partial e_1^1} + \frac{R^2}{2}\frac{\partial E_{y_1,y_2}e_3^M(y_1)}{\partial e_1^1} \\ e_2^M(y_1) &= \frac{R}{2}\frac{\partial E_{y_2}U_3}{\partial e_2^M} + \frac{R}{2}\frac{\partial E_{y_2}e_3^M(y_1,y_2)}{\partial e_2^M} \\ e_3^M(y_1,y_2) &= 0 \end{aligned} \tag{1}$$

indicate that the agent is behaving optimally in setting the marginal benefit equal to the marginal cost of effort. Even though effort will not affect second-period wages, output (and thereby effort) will affect wages in period 3, since the principal and agent will renegotiate a third-period contract taking into account the observed performance in periods 1 and 2.

The agent chooses effort to maximize utility taking the principal's beliefs on effort as given. Any marginal increase (decrease) in effort beyond those beliefs will lead the principal to think that the agent is of relatively high (low) ability when the time comes to renegotiate a new contract. The principal does not observe the agent's effort directly, but he is able to guess it by solving the same maximization problem above. The Perfect Bayesian Equilibrium requirement of beliefs being correct in equilibrium implies the following additional equilibrium conditions: $e_1^M = e_1^{M*}, e_2^M = e_2^{M*}$,and $e_3^M = e_3^{M*}$.

Essentially, the principal observes output net of effort every period, $\bar{z}_t \equiv y_t - e_t^* = a + z_t$. Using the first-order conditions (1) and the equilibrium conditions on effort beliefs, we find that optimal first and second-period effort are

$$e_1^{M*} = \frac{R^2}{2} E_{y_1,y_2}\left(\frac{\kappa_1 h_z}{h_a + 2h_z} \exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1 + \bar{z}_2)}{h_a + 2h_z}\right)\right]\right) \tag{2}$$

and

$$e_2^{M*}(y_1) = \frac{R}{2} E_{y_2}\left(\frac{\kappa_1 h_z}{h_a + 2h_z} \exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1 + \bar{z}_2)}{h_a + 2h_z}\right)\right]\right). \tag{3}$$

For a more complete derivation of these equations, see Appendix A.

Note that optimal second-period effort (e_2^{M*} in equation (3)) is a function of realized first-period output ($\bar{z}_1 = a + z_1$). For high output realizations, the agent exerts less effort than when observed first-period output is low. Even though actual second-period effort obeys this function, the principal and the agent must make an *ex ante* decision on the type of contract (single or multi-period). This *ex ante* decision therefore must be based on the *expected* level of second-period effort, conditional upon period 1 information. Equations (2) and (3) lead to our first proposition.

Proposition 1 *(Expected) within-contract effort is increasing,* $e_1^{M*} < e_2^{M*}$.

Proof. *See Appendix A.* ■

The intuition is simple: the agent realizes that the benefits from exerting any effort during the contract will accrue to him when the current contract is over. The costs of exerting effort, however, are realized contemporaneously. In addition, the principal weighs first and second-period outcomes equally when updating beliefs about the ability. In each period the agent chooses an effort level that sets the marginal costs of exerting effort equal to the (future) marginal benefits. Since the returns to effort in period 1 are two periods away, he chooses a lower level of effort in period one than in period two, when the present-value of effort exertion benefits are higher.

2.2 Case 2: Single-period contracts

As in Case 1, the agent will choose effort to maximize expected utility,

$$\max_{e_1^S, e_2^S(\cdot), e_3^S(\cdot)} U(w_1^S) - \left(e_1^S\right)^2 + RE_{y_1}\left(U\left(w_2^S(y_1)\right) - \left(e_2^S(y_1)\right)^2\right) + \\ R^2 E_{y_1,y_2}\left(U_3(w_3^S(y_1,y_2)) - \left(e_3^S(y_1,y_2)\right)^2\right),$$

which results in the following first-order conditions,

$$\begin{aligned} e_1^S &= \frac{R}{2}\frac{\partial E_{y_1,y_2}U_2}{\partial e_1^2} + \frac{R}{2}\frac{\partial E_{y_1}e_2^S(y_1)}{\partial e_1^2} + \frac{R^2}{2}\frac{\partial E_{y_1,y_2}U_3}{\partial e_1^2} + \frac{R^2}{2}\frac{\partial E_{y_1,y_2}e_3^S(y_1,y_2)}{\partial e_1^2} \\ e_2^S(y_1) &= \frac{R}{2}\frac{\partial E_{y_2}U_3}{\partial e_2^2} + \frac{R}{2}\frac{\partial E_{y_2}e_3^S(y_1,y_2)}{\partial e_1^2} \\ e_3^S(y_1,y_2) &= 0. \end{aligned} \quad (4)$$

In addition, the PBE condition that beliefs are correct in equilibrium applies: $e_1^S = e_1^{S*}$, $e_2^S = e_2^{S*}$, and $e_3^S = e_3^{S*}$.

In this case, the second-period wage is negotiated incorporating the information conveyed from the agent's first-period output. Furthermore, just as in Case 1, the agent's optimal effort choice at any given time is a response to past observations.

The equilibrium conditions lead to two propositions, which together show that aggregate effort in a multi-period contract is less than the effort exerted when the agent is covered by multiple single-period contracts. The effort incentives in the multi-period contract, seen in isolation, make it less desirable to the principal.

Proposition 2 *First-period effort under a single-period contract is greater than first-period effort under a multi-period contract;* $e_1^{S*} \geq e_1^{M*}$.

Proof. *See Appendix A.* ■

At first glance this proposition might seem at odds with a model that talks of netting out effort levels and information structures that are independent of effort. The principal correctly anticipates the effort level by imputing the *highest* incentive compatible level of effort for the agent. This result is driven by the fact that the agent in our model is finitely lived. The agent in a single-period deal can affect a longer stream of revenues through is current actions than can the agent in the first year of a multi-period contract. This is because the payoff next period for the agent in a multi-period contract is fixed, whereas the agent in a single-period contract will sign a new contract at the end of the current period, incorporating any new information that arises.

Proposition 3 *Effort in the second period of the agent's life is the same regardless of whether the agent is in a one-period contract or in the second period of a two-period contract;* $e_2^{M*} = e_2^{S*}$.

Proof. *See Appendix A.* ■

The only unknown parameter with a fixed value throughout the game is the agent's true ability. Both the agent and the principal, therefore, update their beliefs on ability using realizations of past output. In equilibrium, the market correctly predicts what the agent's level of effort will be, consequently information updating is independent of effort. Recall that updating is done in a Bayesian fashion (described above) with output from the first and second period affecting beliefs equally. Thus, an agent in the second period is choosing effort to affect third-period wages, accounting for the fact that the principal has observed one-period of production. The implication is that the effort incentives in period 2 are unconditional on the current or past contractual structures, and instead only depend on the effect that second-period effort will have on third-period wage. In essence, the last year of a multi-period deal is tantamount to a one-period contract, in which a new contract has to be negotiated in the following period.

Viewed jointly, propositions 1, 2 and 3 unambiguously show that the total amount of effort elicited from the agent in a multi-period contract is less than in a series of single-period contracts. Despite these adverse effort incentives, the insurance benefits of a long-term contract can cause it to be Pareto optimal, which we discuss in the next section.

2.3 Optimal contract choice

At the start of the game the agent and the principal have a choice of signing a contract that specifies a fixed wage for 1 or 2 periods. Propositions 2 and 3 jointly imply that, all else equal, the principal is worse off under the multi-period contract scenario. Given that the agent is risk averse, however, there are potential gains from insuring the agent against possible negative output shocks.

As in Holmström (1983), firms may prefer a multi-period contract, since the per-period wage concessions given by the agent result in long-term contracts being the less expensive mode of hiring labor. More formally, if there exists a 2-period profit sharing rule, κ_2, which when compared to the corresponding 1-period rule, κ_1, yields the following outcomes,

$$E_{y_1,y_2}((1-\kappa_2)\left[y_1(e_1^{M*}) + Ry_2(e_2^{M*})\right] > (1-\kappa_1)y_1(e_1^{S*}) + Ry_2(e_2^{S*})) \tag{5}$$

and

$$E_{y_1,y_2}(U(\kappa_2 \cdot y_1(e_1^{M*})) - \left(e_1^{M*}\right)^2 + R\Big[U(\kappa_2 \cdot y_2(e_2^{M*})) - \left(e_2^{M*}\right)^2\Big] > \\ U(\kappa_1 \cdot y_1(e_1^{S*})) - \left(e_1^{S*}\right)^2 + RE_{y_1}\Big[U(\kappa_1 \cdot y_2(e_2^{S*})) - \left(e_2^{S*}\right)^2\Big]\,), \tag{6}$$

then both the principal and the agent benefit from the multi-period deal. The left-hand side of equation (5) shows the profit that the principal receives in the multi-period case. While the principal receives a higher profit share than in the 1-period case, the adverse effort incentives simultaneously ensure that $e_1^{M*} < e_1^{S*}$. For equation (5) to hold, the increased share of profits the principal receives must be enough to at least offset the reduced output in a multi-period contract.

The guaranteed wage paid over two periods provides the agent with increased certainty (note the difference in second-period utility in equation (6) above). Since the increased share of profits associated with a higher κ_2 directly lowers the wage paid to the agent, κ_2 cannot be set arbitrarily low. The choice of contract in equilibrium depends on whether such a multi-period profit sharing rule exists.

2.4 Numerical simulation

We provide a numerical simulation of the model to (i) get a sense of how much effort fluctuates within and across alternative contracts, and (ii) to demonstrate a scenario in which a multi-period contract Pareto-dominates a series of single-period contracts (thus showing existence). The CARA utility functional form makes it relatively easy to find closed-form solutions to our first-order conditions on effort above (see Appendix B for a detailed outline of how we compute effort in each case). To begin, we set the model's underlying parameters to the following values: $\alpha = 2$, $\mu = 0.5$, $h_a = 1$, $h_z = 1$, $R = 0.9$, $\kappa_1 = 0.5$.

With closed-form solutions for effort, and the initial parameter values, we can calculate the differences in effort that arise from the two contractual scenarios considered. Figure 1 shows the resultant equilibrium effort paths. The solid line shows the market's *ex ante* expectation of effort given by the agent in Case 2, i.e. under 3 one-period contracts. The single period contract case is exactly analogous to the type of contract analyzed in the career concerns literature (e.g., Holmstrom 1999), with the same resulting optimal effort portfolio. The dashed line shows the expected effort portfolio under Case 1, i.e. a two-period contract followed by a single-period contract.

As expected from Proposition 2, effort in period 1 is lower for a two-period deal than for a one-period deal. In particular, our parameter values imply that first-period effort is 60% lower under a multi-period contract. Nevertheless, our simulation allows us to find a profit sharing rule such that a multiperiod contract is Pareto preferred. If we set the share of profits going to the agent at 46.5% ($\kappa_2 = 0.465$), instead of 50% in a 1-period contract, the present value of the principal's expected lifetime profits increases by 1.5% in the multi-period case relative to the single-period case. Meanwhile, the agent benefits enough from the certainty of a multi-period deal that his present value of expected lifetime utility increases by 2.1%.[2]

By adjusting the underlying parameters of the model, we can mimic various real-world situa-

[2] We also calculate the scenario of the wage structure needed for the principal and agent to agree to a single lifetime contract. In this case the agent will exert zero effort in every period. The principal will require in return a $\kappa_3 = 0.42$ in order to break even against the single-period contract case. The wage in this case is so low that the agent prefers a single-period contract, which itself is Pareto dominated by the multi-period contract scenario.

tions. The model allows for comparisons, for instance, of optimal contract length for 2 different workers, e.g. an experienced worker (with many past output realizations and precise h_a) with a less experienced worker (low h_a) having the same expected ability. Specifically, if we increase the degree of precision on ability so that $h_a = 5$, the principal will require a multi-period profit share of $\kappa_2 = 0.49$ to be indifferent between the multi-period and single-period contract cases. With such a profit share, however, the agent's expected lifetime utility from a series of single-period contracts is larger than the multi-period case: with such a precise estimate of ability, the insurance benefits of a long-term contract are small.

3 Empirical application

We analyze the agency implications of our model using 15 years of data from the National Basketball Association. We view this sector as a suitable setting for this kind of analysis for several reasons. First, the extensive (and oftentimes public) information on individual and firm characteristics makes sports an ideal setting for testing many labor market phenomena (Kahn 2000). As opposed to other professional sports, which include many incentive-based, non-guaranteed contracts, basketball has simple, fixed-wage contractual structures and rarely uses incentive clauses based on achievement.[3] Second, basketball players of all positions are evaluated roughly uniformly across a variety of productivity categories. This is distinct from corporate employees and other professional athletes, whose job performance and productivity can be measured differently (or cannot be compared at all) depending on their position, job title, or the organization to which they belong. Finally, the bargaining between team owners and players (via their representatives) is well-known and we believe matches the profit sharing in the theoretical construct.

3.1 Estimation strategy

Consistent with our theoretical framework, we model log performance (Y)—measured by the NBA efficiency score (described below)—as a linear function of player ability (A), effort (E), and an idiosyncratic component ($\tilde{\zeta}$),

$$\ln(Y_{it}) = \tilde{\alpha} + \tilde{\delta}E_{it} + \beta A_{it} + \tilde{\zeta}_{it}. \tag{7}$$

Here i indexes individual players, and t indexes time.

Because the model above predicts that effort is directly related to the time until the contract's expiration, we model effort as a possibly nonlinear function of years remaining in the contract. In

[3] We cannot identify the players whose contracts include incentive clauses. All players have a team-based incentive, which is that players receive a bonus based on "pool" of playoff revenues if their team makes the playoffs.

the simplest (linear) case, we have

$$E_{it} = \gamma_0 + \gamma_1 r_{it} + u_{it} \tag{8}$$

with r denoting the years remaining in the contract. Combining equations (7) and (8), we have

$$\begin{aligned} \ln(Y_{it}) &= (\tilde{\alpha} + \tilde{\delta}\gamma_0) + \beta A_{it} + \tilde{\delta}\gamma_1 r_{it} + \tilde{\zeta}_{it} + \tilde{\delta} u_{it} \\ &= \alpha + \beta A_{it} + \delta r_{it} + \zeta_{it} \end{aligned} \tag{9}$$

for $\alpha = (\tilde{\alpha} + \tilde{\delta}\gamma_0)$, $\delta = \tilde{\delta}\gamma_1$ and $\zeta_{it} = \tilde{\zeta}_{it} + \tilde{\delta} u_{it}$. The coefficient of interest is δ. The identifying assumption in the empirical estimation is that the number of years remaining in a current contract at time t is a predetermined variable (i.e., it is decided at some time $t - s$).

Least squares estimation of δ is consistent if the number of years remaining in a player's contract is exogenous to the error term, conditional on other ability measures, $Cov \zeta_{it}, r_{it} | A_{it} = 0$. One concern is that the observable player characteristics will only imperfectly proxy for a player's true ability. That is, ability should be modelled as $A_{it} = \beta X_{it} + \eta_i$, with X denoting the vector of observable player characteristics including the player's current team, experience (number of years in the league), age, height, college attended, race, when the player was selected in the NBA draft (first selection, second selection, ..., N^{th} selection), and superstar status. The unobserved ability component, η_i, is problematic because contract length is not randomly distributed: longer contracts tend to be given to players who have exhibited greater talent. As long as talent is autocorrelated, having many years remaining in a contract will be associated with high (unobserved) ability, which leads to greater output, i.e. the estimation has an omitted variable which exerts upward bias on δ.

Our data allows us to estimate a fixed-effects (within-groups) model in which we specify the random component, $\zeta_{it} = \theta_t + \eta_i + \varepsilon_{it}$, to include year-specific shocks distributed to all players (θ_t), as well as individual-specific (η_i) and idiosyncratic components (ε_{it}). In addition, we include in vector A the duration and total value of the current contract (within a fixed contract duration, more able players are paid more).[4] Conditional on the duration of the current contract, δ represents the percentage change in output when a player is one year further from his contract termination date. In the presence of moral hazard, we expect $\delta < 0$.

3.2 Data

The data were collected from a variety of sources. First, we compiled player characteristics and performance statistics from Sports-Reference Inc., which collects information on professional athletes. We obtained data on NBA player characteristics including height, weight, position played, team, where the player was selected in the NBA draft, the first year in the league, college attended,

[4] While higher value contracts are given to more able players, there may be an income effect which dampens the career concerns. Which effect dominates is a quantitative question.

and birth date. Contract information from 2000 to 2006 was primarily derived from the *USA Today* NBA contract database, which gives the annual salary, and contract start and end dates for every player in the league. Using this information together with additional publicly available sources,[5] we assigned contract information prior to 2000 for each player for all years possible (1991 was the earliest). We also were able to identify and correct several mistakes and omissions from the *USA Today* database.

In the NBA, player performance is multidimensional and includes, among other variables, points scored, rebounds, shooting percentage, assists, and turnovers. To create an objective measure to judge performance, the NBA has created an "efficiency rating," which is a function that maps a player's multidimensional output from $\mathbb{R}^N$ to $\mathbb{R}$. The efficiency rating is claimed to be used by coaches and scouts to evaluate a player's performance. While certain beneficial components of the efficiency score can only increase with the number of minutes played (e.g., rebounds and assists), other damaging components (e.g., turnovers and missed shots) may also increase. The net effect is that the efficiency score is used to compare the performance of players of all positions. The efficiency score is calculated for each game played. We use average annual efficiency score (across all games played that year) as the dependent variable in our estimation.[6]

Table 1 reports the summary statistics for our sample. In all, we have 2,260 player-year observations which are derived from an unbalanced panel of 654 players. The average player in our sample is 6 feet 7 inches tall, was the 25th selection in the NBA draft, is observed for 3.5 years (max 13) throughout at least part of 2 contracts (max 5). Players enter the sample at mean age of 25 with almost 4 years of professional experience. The dominant racial category of our sample is American-born black; however, about a quarter of the sample is either American-born white (12.0%), or of "other" racial composition (10.7%), which overwhelmingly includes foreign-born players. Because "superstars" have a large positive, external effect on league notoriety and revenue aside from their performance on the court (Hausman and Leonard 1997) and often receive a large portion of income from endorsements outside of their NBA contract, we created a superstar indicator variable. Nearly two percent of players (4.4% of observations) are superstars.[7]

---

[5] These other sources include sports consulting companies, newspaper articles, and independent gurus and hobbyists.

[6] Information on the NBA efficiency score, including the formula can be found at http://www.nba.com/statistics/efficiency.html. There are several competing measures of player performance including the "win score" and "player efficiency rating." The correlation coefficient between each of these measures is approximately 0.9. If we substitute either of the alternative measures as the dependent variable in our empirical analysis, the results of shirking are, if anything, more pronounced. As an example of the efficiency index consider two players. Player A plays 43 minutes, makes 5 of 22 shots, 7 of 9 free throw attempts, gets 8 rebounds, 6 assists, 3 steals, commits 4 turnovers and scores 17 total points. Player B plays 29 minutes, making 5 of his 8 shots, 3 of 4 free throw attempts, gets 4 rebounds, 7 assists, commits 2 turnovers, and scores 15 total points. According to the NBA efficiency rating we find that Player A scored an efficiency of +11, while player B received +20 and had the better game. Kevin Garnett was the league leader by this measure of efficiency for the 2006-2007 NBA season.

[7] Players we identify as superstars during our sample period are Ray Allen, Kobe Bryant, Vince Carter, Tim Duncan, Kevin Garnett, Allen Iverson, Michael Jordan, Jason Kidd, Tracy McGrady, Alonzo Mourning, Dirk Nowitzki,

Panel B of Table 1 shows summary statistics for the contractual terms. The average contract in our sample lasts 3 years—80% of players sign a contract lasting 5 or fewer years—however, we observe contracts lasting every value between 1 and 12 years inclusive. The average player receives $2.3 million per year, with a large standard deviation.[8]

Panel C of Table 1 presents the descriptive performance statistics for our sample, as measured by the NBA efficiency score. Players achieve an average per game NBA efficiency rating of 10.0, though there is quite a lot of variability in this measure—the players we identify as superstars have a mean efficiency score of 23.5 (max=36.9). The table also shows that the mean efficiency score is increasing throughout the final 3 years of a player's current contract.[9]

Figure 2 presents a histogram of the distribution of years remaining in the current contract for the 2,260 player-year observations. This figure emphasizes that while we observe players who have a variety of contractual terms, we also observe players at many different stages in their contract.

Players leave our data once they are fired ("waive"), retire, or are otherwise not hired by any team following an expired contract. In the NBA, teams may hire a maximum of 15 players. If a team decides to fire a player, it must still pay him according to the contract terms, but the vacant roster spot may be filled with another player. A small minority of players sign extremely short-term contracts (often 10-day contracts). These contracts are usually given to relatively low-ability players signed to fill a suddenly available roster spot, usually stemming from injury. We include only players who are signed to contracts lasting one or more years. In addition, because the NBA efficiency score is measured on a per game basis, players receive missing values in years when they have zero games played due to injury.

3.3 Empirical Results

Table 2 presents the coefficients from four variations of our estimation equation (equation (9)). In column (1) we estimate the simplest version of a fixed-effects regression model under the assumption that the number of years until contract renewal has a linear relationship with output. Consistent with the equilibria in the theoretical model above, the results show that, *ceteris paribus*, the more years remaining in a player's current contract, the lower his performance. The estimated coefficient implies an average of a 2.2% reduction in output for each additional year remaining until contract renewal. This effect is statistically significant. Salary and total contract duration serve as (possibly time-variant) ability proxies, and performance increases with each. Productivity decreases

Shaquille O'Neal, and Paul Pierce. None of the results that we present below are sensitive to marginal changes to the set of players classified as superstars.

[8] All dollar values have been converted to 1995$ to account for inflation.

[9] These statistics were calculated using only the subset of players with contracts of duration 3 years or longer.

with experience, reflecting that aging in professional sports tends to dominate any positive effects of learning (Fair 2007). Because height, superstar status, and draft position are constant within players, they are dropped from the fixed-effect analysis.

In column (2) we allow for the number of years remaining in the contract to have a non-linear effect on output. Indeed, one plausible scenario is that the number of years remaining in the contract is negatively associated with performance, but that the marginal decrease in effort is a diminishing function of the number of years remaining in the contract. In this example, the linear prediction estimated in column (1) would underestimate the marginal moral hazard effect for initial increases in contract duration, and overestimate the marginal moral hazard effect in the early years of a long contract.

A second estimation issue with column (1) is that, while the model controls for gains in *league*-specific experience, it does not account for gains in *team*-specific learning or experience. Indeed, 43% of all instances of players switching teams occur at the end of an expiring contract. To the extent that players accrue team-specific knowledge, and that this learning tends to occur towards the beginning of contracts, the omission of team-specific tenure would bias the results towards finding increased within-contract productivity. In column (2), we present the results including controls for changes in team-specific experience, measured by the number of years on the current team.

The coefficients imply a nonlinear productivity effect, with the largest decreases in effort associated with initial movements away from the termination date. The first two coefficients reported in column (2) suggest that, all else equal, a player's effort in the penultimate year of his contract is 6% lower than effort in the final year of the contract. Similarly, effort in the player's third-to-last year of the contract is 4% lower than in his penultimate contract year, and effort in his fourth-to-last contract year is an additional 2% lower than in his third-to-last year. Although the coefficients in this model imply that output tends to rise for sufficiently large numbers of years remaining, the effort effects beyond year 4 are not statistically significant. In addition, we find that players do appear to acquire beneficial team-specific training, with productivity increasing approximately 6% per year. There are no significant changes to the other coefficients, with respect to those in column (1).

Column (3) reports the coefficients from a regression allowing for both a non-linear effect of years remaining and for an interaction between a player's experience and number of years remaining in his contract. The coefficient on the interaction between years remaining and experience is positive and statistically significant, implying that the adverse effort effect of a multi-period contract is reduced for more experienced players.

The reported coefficients allow for a comparison of the effect on output of a long-term contract for a player different stages of his career. For example, the coefficients suggest that the performance of a player with two years of experience and in the penultimate year of the contract will be 15%

less than the same player in the final year of his current contract. This output differential is approximately twice as large as would be the case for the same player with 6 years of experience. In addition, the coefficients imply that the within contract effort fluctuations approach zero for players with over 10 years of experience. This finding conforms with our theoretical results. Because the market has a history of outcomes on which to base its estimate of ability, the precision of this estimate will be relatively high. The outcome of any singular period will therefore have a relatively small impact on the beliefs about ability. Second, the performance of older players affects a shorter stream of future wages and therefore the present discounted benefit of effort exertion is relatively low. For these reasons, the maximum incentive compatible level of effort is reduced.[10]

The F statistic testing the joint significance of the player-specific effects is 5.47, 5.44, and 5.66 for the models in columns (1) through (3) respectively (the 99 percent critical value on the F table is 1.16). In addition, the estimated parameter ρ in the fixed-effects models, which represents the proportion of the total variance in log performance that can be attributed to the player fixed-effects, η_i, is larger than zero. Together, the large value of the F statistic and the ρ suggest that the panel-level estimates are important and are likely to be different than pooled OLS estimates. We test this in column (4), which presents the results of our estimation equation specifying the error term as $\zeta_{it} = \theta_t + \xi_{it}$, excluding player fixed effects. Estimation via pooled OLS allows the inclusion of many variables that are constant across time. As in the fixed-effects analysis, we find that salary, contract duration, and team-specific experience are positively associated with performance. Performance among players that we classify as superstars is, on average, 21% higher than their peers. Likewise, top selections in the NBA draft perform better than those selected later on. Finally, whites tend to have lower output than blacks, and age is negatively associated with output.

The OLS estimates, however, fail to find any within-contract performance effects, suggesting that, even when controlling for all observable player characteristics, including proxies for ability such as superstar status and draft position (first selection, second selection, ..., N^{th} selection), unobserved individual heterogeneity appears to be an important source of bias.

3.4 Illustrations and Extensions

To get a better sense of the magnitude of the incentive effects of contracts, we estimate a version of equation (9) identical to the fixed-effects model reported in column (2) of Table 2, with the exception that we group players on the basis of years remaining in their contract, rather than estimating years remaining as a continuous variable. In this estimation, we combine all players

[10] Although we control for the effect of experience on output, there may be a concern that the effects of this model are being driven by differentials in gains to experience between younger and older players. There is, however, effectively no difference between the models presented here and versions which additionally allow for a non-linear impact of experience on output, for instance by additionally controlling for the square of experience.

with eight or more years (less than 1% of the sample) remaining into one category. Figure 3 presents the combined coefficients from the years-remaining dummy variables, and corresponding years remaining × experience interaction terms. The reported coefficients are all for a hypothetical player with three years of experience and should be interpreted as being relative to output in the final year of the contract. Standard errors are shown in italics.

These results re-emphasize that player productivity is negatively associated with time until contract renewal. Players in the penultimate contractual year produce, on average, 7.2% fewer NBA efficiency points, every game. There is a steady drop in performance for each additional year remaining. The negative effect on output triples for players with 3 years remaining in the current contract (23% lower output than in the final year). In absolute terms, this difference amounts to a little more than 1/3 of a standard deviation in the NBA efficiency score for an average player, and is approximately equivalent to needing two additional shot attempts per game to get the same number of points, or having one standard deviation more assists per game, holding all other performance statistics constant. As in Table 2, the adverse effort effect of multi-period contracts level off quite quickly. Output for players with four or more years remaining is nearly identical, with performance leveling off at about 30% less than output in the terminal year. The magnitude of the within-contract changes in player performance are quantitatively similar to those in Lazear (2000), who finds an approximately 44-percent increase in output per worker associated with moving from a guaranteed salary to a piece rate compensation schedule.

In Table 3, we present the results of the performance effect of being in the first or last year of the contract. Again, all reported coefficients are derived from fixed-effects (within-group) estimation which includes all covariates reported in column (2) of Table 2. In particular, the results reported in Panel A are derived from a version of estimation equation (9), substituting dummy variables denoting a player in the first year or last year in a contract, for the continuous measure of years remaining in the contract. We find that players score 7 percent fewer efficiency points per game in the first year of their current contract, and 6 percent more efficiency points in the last year of their contract relative to all other years. These effects are statistically significant, and are similar to the results found in Table 2 and Figure 3, above. The results are larger in magnitude than findings in previous studies such as Berri and Krautman (2006), who use a smaller panel of players and a less detailed estimation method to find only a 1 percent decrease in output associated with being in the first year of a contract.

In Panel B, we select the subsample of players in the last year of their contract, for whom we also have data on the first year of the subsequent contract (714 player-year observations). Using only these two observations, we estimate a fixed-effects model with a dummy variable which takes the value of 1 if the player is in the first year of a new contract. We find that relative to the year prior (i.e., the final year of the previous contract), player productivity falls by 16%, lending further support to the hypothesis of decreased effort subsequent to signing a new contract.[11]

4 Concluding Remarks

Our results lend theoretical and empirical support to the notion that workers will be less industrious the more distant the expiration of the current contract is, even in the presence of career concerns. Results of this, and related work, have crucial implications for the contemporary labor market, in which contractual arrangements frequently specify payment over a fixed period of time in return for services rendered. The model lends a sound economic justification for the existence of multi-period contracts in spite of the effort effects: with uncertain output, these contracts act as insurance for the agent and, in return for bearing some of the risk, firms benefit by paying a lower wage. While the results we derive *per se* depend on the functional form assumptions, we are confident that changes in these assumptions will yield qualitatively similar results. In addition, while it is beyond the scope of this paper, we conjecture that our results carry out qualitatively to a model with N periods.

To the extent that multi-period contracts give reduced effort incentives, one might wonder if there is any economic justification as to why NBA contracts, and those in other labor markets, are written in this fixed-wage manner. Indeed, an optimal contract should include all (free) contingencies which provide information regarding the agent's (hidden) actions. The use of incentive-laden contracts, for instance, which offer performance bonuses contingent on output each year, could serve the dual purpose of allowing for a more precise proxy of the agent's underlying ability, and to offset the effect of shirking in the early years of the contract. The use of these incentive contracts in the NBA, however, is minimal.

Holmström and Milgrom (1987, 1991) suggest that one reason why real-world incentive schemes are relatively rare may be that when an agent's "output" is multidimensional, there may exist efficiency reasons for paying fixed-wages because any attempt to specifically engineer incentives to motivate hard work in every period may conflict with other team or league-wide profitability goals. Thus, including an array of bonuses is not "free" since their presence may give rise to new sources of inefficiencies. Including additional year-by-year player bonuses in the NBA for surpassing a prespecified points-scored threshold, for instance, may motivate the player to shoot at every opportunity which, in-turn, conflicts with sound teamwork that is valued by fans and is critical for overall franchise success. In the Holmström-Milgrom model of multitasking, the ability to separate job tasks among workers is essential for such incentives to be efficient.

---

[11] There is potential selection bias in panel B of Table 3. Namely, if the only players who are able to continue in the league are players who had large positive random output shocks in the final year of their contract, we would expect such a large shock to be unlikely to occur in the subsequent year. Nevertheless, we view this estimation as a useful illustration and robustnes check.

Given the portfolio of within-contract effort that we find theoretically and empirically, another puzzle is why the principal in our model weights performance in every period equally. The principal could induce constant (high) effort by announcing that when the time comes for recontracting, she will use an inverse weighting formula that puts more emphasis on performance in early periods of the contract. There are at least two potential problems with this strategy. First, neither in our model, nor in reality is there commitment by firms to the structure of future contracts; indeed such a commitment would effectively be a contract in itself. In addition, announcements of this kind are not renegotiation proof. Any alternate weighting scheme has the additional effect of reducing the insurance gains to the agent. Thus, such an alternative would erode the benefits inherent in long-term contracts, and lead to fewer observed long-term contracts since agents will be less willing to make the necessary wage concessions.

References

Berri DJ, AC Krautmann. Shirking on the court, *Economic Inquiry* 2006; **44(3)**: 536-546.

Cantor R. Work effort and contract length, *Economica* 1988; **55(219):** 343-353.

Dewatripont M, I Jewitt, and J Tirole. The economics of career concerns, part I: comparing information structures, *The Review of Economic Studies* 1999; **66(1)**: 183-198.

Fair RC. Estimated age effects in athletic events and chess, *Experimental Aging Research* 2007; **33**: 37-57.

Foster AD, MR Rosenzweig. A test for moral hazard in the labor market: contractual arrangements, effort, and health, *The Review of Economics and Statistics* 1994; **76(2):** 213-227.

Fudenberg D, J Tirole. Moral hazard and renegotiation in agency contracts, *Econometrica* 1990; **58(6):** 1279-1319.

Fudenberg D, B Holmström, P Milgrom. Short-term contracts and long-term agency relationships, *Journal of Economic Theory* 1990; **51(1):** 1-31.

Gibbons R and KJ Murphy. Optimal incentive contracts in the presence of career concerns: theory and evidence, *Journal of Political Economy* 1992; **100(31)**: 468-505.

Harris, M and B Holmström. On the duration of agreements, *Institute for Mathematical Studies in the Social Sciences* 1984, Stanford University, Technical Report no. 424.

Hausman JA, and GK Leonard. Superstars in the National Basketball Association: economic value and policy, *Journal of Labor Economics* 1997; **15(4)**: 586-624.

Holmström B. Equilibrium long-term labor contracts, *The Quarterly Journal of Economics* 1983; **98 Supplement:** 23-54.

Holmström B. Managerial incentive problems: a dynamic perspective, *The Review of Economic Studies* 1999; **66(1)**: 169-182. [Originally in *Essays in Economics and Management in Honor of Lars Wahlbeck* 1982 (Helsinki: Swedish School of Economics)].

Holmström B, and P Milgrom. Aggregation and linearity in the provision of intertemporal incentives, *Econometrica* 1987; **55(2)**: 303-328.

Holmström B, and P Milgrom. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design, *Journal of Law, Economics, & Organization* 1991; **7**: 24-52.

Kahn LM. The sports business as a labor market laboratory, *The Journal of Economic Perspectives* 2000; **14(3)**: 75-94.

Lambert RA. Long-term contracts and moral hazard, *The Bell Journal of Economics* 1983; **14(2):** 441-452.

Lazear EP. Performance pay and productivity, *The American Economic Review* 2000; **90(5)**: 1346-1361.

Lehn K. Property rights, risk sharing and player disability in Major League Baseball, in *Sportometrics* 1990. Goff and Tollinson, eds. College Station, Texas (Texas A&M Press): 35-58.

Zeckhauser R. Medical insurance: A case study of the trade-off between risk spreading and appropriate incentives, *Journal of Economic Theory* 1970; **2**: 10-26.

Figure 1. Effort profile under chosen parameter values

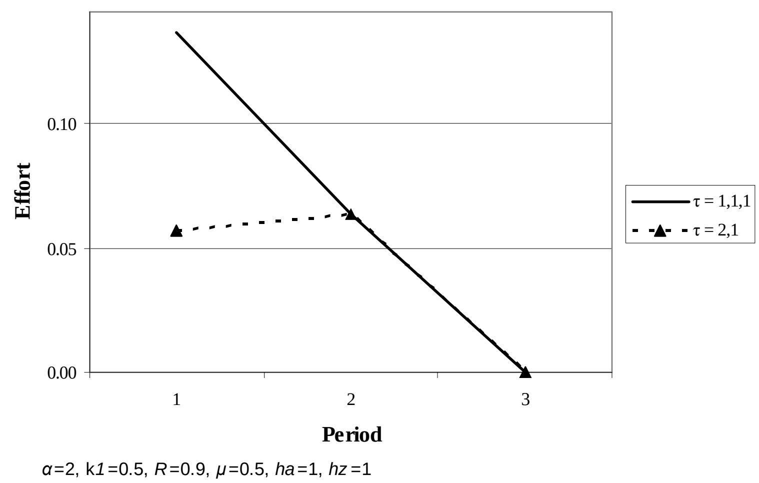

Figure 2. Frequency distribution for years remaining in player contracts

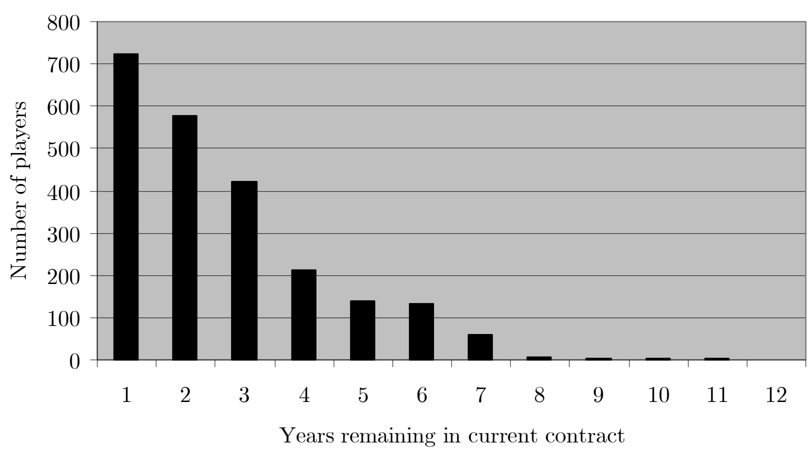

Figure 3. Estimated performance-years remaining elasticity

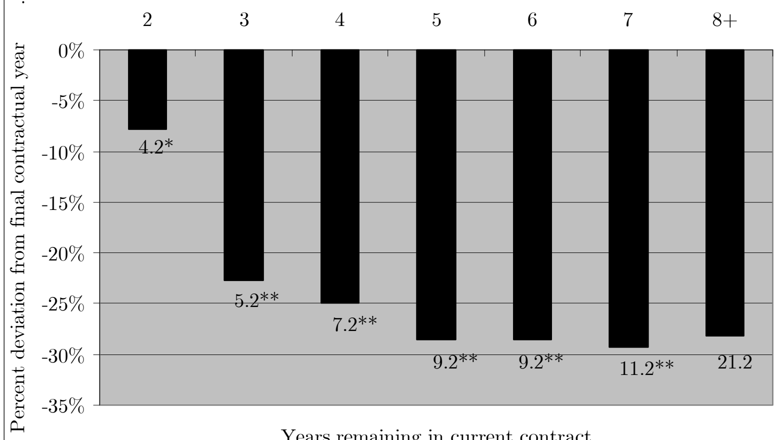

NOTES: The figure depicts the results of a variation of equation (7) which estimates the effort effect of number of years remaining in the current contract on output, allowing the effect of each year to enter individually; the coefficients shown are the combined years remaining and years remaining/experience interaction for a player with 3 years experience; all results are relative to output in the final year of the contract with standard errors reported in italics; ** and * denote p¡ 0.01 and p¡0.1

Table 1. Summary statistics

# of player-year observations	2,260
# of players	654
A. Player Characteristics	
# of years observed: Mean	3.5
Standard deviation	2.3
Mean height (in feet)	6.6
Mean draft position	25.5
Mean age at first observation	25.2
Mean experience at first observation	3.9
Race: American-born, black	77.3%
American-born, white	12.0%
Other	10.7%
Superstar: % of players	1.8%
% of observations	4.4%
B. Contract Statistics	
# of contracts observed: Mean	1.9
Max	5
Length of contract (years): Mean	3.1
Standard Deviation	2.1
Years remaining in current contract: Mean	2.6
Salary (millions of 1995$): Mean	$2.26
Standard deviation	$2.75
C. NBA Efficiency Score Statistics	
Total population:	
Mean	10.0
Standard Deviation	6.4
Superstars	
Mean	23.5
Standard Deviation	6.0
Mean efficiency score with	
One year remaining in current contract	10.9
Two years remaining in current contract	10.6
Three year remaining in current contract	9.5

Table 2. Regression results of ln(Efficiency)[1] on player characteristics and years remaining in current contract.

	Fixed Effects				OLS
Independent Variable	(1)	(2)	(3)	(4)	(5)
Years remaining	-0.0219*	-0.0981**	-0.0870**	-0.2165**	0.0246
	[-2.2]	[-3.9]	[-3.5]	[-6.6]	[0.7]
Years remaining squared	--	0.0106**	0.0106**	0.0161**	-0.0059
		[3.4]	[3.5]	[5.1]	[-1.8]
Experience	-0.0683*	-0.0909*	-0.1026*	-0.1640**	-0.0011
	[-5.1]	[-5.4]	[-6.2]	[-8.0]	[0.1]
Team-specific experience	--	--	0.0655	0.0571	0.0844**
			[6.1]	[5.3]	[7.8]
Years remaining × experience	--	--	--	0.0167**	0.0050*
				[6.6]	[1.9]
Salary	0.0055**	0.0049**	0.0033**	0.0032**	0.0109**
	[4.3]	[3.8]	[2.6]	[2.5]	[8.6]
Total duration of current contract	0.0355*	0.0382*	0.0310*	0.0304*	0.0217
	[2.3]	[2.4]	[2.1]	[2.1]	[1.1]
Superstar	--	--	--	--	0.2124**
					[3.5]
Age	--	--	--	--	-0.0117
					[-1.2]
Height	--	--	--	--	-0.1317**
					[-2.8]
Draft pick	--	--	--	--	-0.0105**
					[-9.4]
Race					
American-born, white	--	--	--	--	-0.1566**
					[-3.6]
Foreign-born	--	--	--	--	-0.0235
					[-0.4]
R^2	0.1168	0.1232	0.1459	0.1692	0.3817
N	2,260	2,260	2,260	2,260	2,260
ρ	0.7863	0.8120	0.8248	0.8579	--

NOTES: t-statistics derived from robust standard errors reported in brackets; All regressions also include team dummies and T-1 year-specific dummies. The OLS results additionally include college-attended dummies; [1] Efficiency score is a measure of player outputas discussed in the text;* and ** denote statistical significance at $p < 0.05$ and $p < 0.01$.

Table 3. Fixed effects egression results of ln(Efficiency)[1] conditional on a player being in the first or last year of the current contract.

	Coef.	Std. Error
A. Full sample (N=2,260)		
First year of a contract	-0.0694**	0.0256
Last year of a contract	0.0645*	0.0286
B. Sample with two consecutive contracts (N=714)		
First year of subsequent contract relative to last year of prior contract	-0.1598**	0.0470

NOTES: The models control for all variables in Table 2, column (3); Robust standard errors reported, [1] Efficiency score is a measure of player output as discussed in the text; * and ** denotes statistical significance at $p < 0.05$ and $p < 0.01$.

Appendix A

A.1 Derivation of optimal effort:

The optimal levels of effort arise through the derivation of equilibrium conditions (1),

$$\begin{aligned} e_1^M &= \frac{R^2}{2}\frac{\partial E_{y_1,y_2}U_3}{\partial e_1^1} + \frac{R}{2}\frac{\partial E_{y_1}e_2^M(y_1)}{\partial e_1^1} + \frac{R^2}{2}\frac{\partial E_{y_1,y_2}e_3^M(y_1)}{\partial e_1^1} \\ e_2^M(y_1) &= \frac{R}{2}\frac{\partial E_{y_2}U_3}{\partial e_2^M} + \frac{R}{2}\frac{\partial E_{y_2}e_3^M(y_1,y_2)}{\partial e_2^M} \\ e_3^M(y_1,y_2) &= 0. \end{aligned}$$

The change in expected third-period utility due to a marginal change in second-period effort is

$$\begin{aligned} \frac{\partial E_{y_2}U_3}{\partial e_2^M} &= \frac{\partial}{\partial e_2^M}E_{y_2}\left(-\frac{1}{\alpha}\exp[-\alpha\kappa_1 y_3(y_1,y_2)]\right) \qquad (10) \\ &= \frac{\partial}{\partial e_2^M}E_{y_2}\left(-\frac{1}{\alpha}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{M*} + y_2 - e_2^{M*})}{h_a+2h_z} + e_3^{M*}\right)\right]\right) \\ &= E_{y_2}\frac{\partial}{\partial e_2^1}\left(-\frac{1}{\alpha}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{M*} + y_2 - e_2^{M*})}{h_a+2h_z} + e_3^{M*}\right)\right]\right) \\ &= E_{y_2}\left(\frac{\kappa_1 h_z}{h_a+2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{M*} + y_2 - e_2^{M*})}{h_a+2h_z} + e_3^{M*}\right)\right]\right) \end{aligned}$$

Using equation (10) and the equilibrium conditions on effort beliefs ($e_t = e_t^*$ in all periods), optimal second-period effort reported in equation (3) follows. Using similar processes, we can derive optimal effort in the first period of a multi-period contract (e_1^{M*}), and first and second-period effort in the case of multiple single-period contracts (e_1^{S*}, and e_2^{S*}).

A.2 Proof for Proposition 1:

Proposition 1 *(Expected) within-contract effort is increasing.*

Proof. *Equations (3) and (2) show that the first-order conditions for e_1^M and e_2^M are*

$$\begin{aligned} e_1^{M*} &= \frac{R^2}{2}E_{y_1,y_2}\left(\frac{\kappa_1 h_z}{h_a+2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1+\bar{z}_2)}{h_a+2h_z}\right)\right]\right) \\ e_2^{M*}(\bar{z}_2) &= \frac{R}{2}E_{y_2}\left(\frac{\kappa_1 h_z}{h_a+2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1+\bar{z}_2)}{h_a+2h_z}\right)\right]\right) \end{aligned}$$

Because the decision to extend or not must come at the start of the game (i.e., before period 1), both the principal and agent are interested in period 1 expectations of second-period effort, $E_{y_1,y_2}(e_2^{M*})$. Comparing the first-order condition for e_2^{1*}, (3), above to the equation for $E_{y_2}U_3$,

$$E_{y_2}U_3 = E_{y_2}\left(-\frac{1}{\alpha}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{1*} + y_2 - e_2^{1*})}{h_a + 2h_z} + e_3^{1*}\right)\right]\right),$$

we can see that $e_2^{M*} = \frac{R}{2}\frac{\partial E_{y_2}U_3}{\partial e_2^M} = \frac{R}{2}\frac{-\alpha\kappa_1 h_z}{(h_a+2h_z)}E_{y_2}U_3 = cE_{y_2}U_3$ *with c representing the constant term,* $\frac{R}{2}\frac{-\alpha\kappa_1 h_z}{(h_a+2h_z)}$*, and in which we apply all equilibrium conditions within expected utilities (i.e.,* $e = e^*$ *in all periods). Likewise, the first order condition for* $e_1^{M*} = \frac{R^2}{2}\frac{\partial Ey_1,y_2U_3}{\partial e_1^M} = RcE_{y_1,y_2}U_3$*. Taking first-period expectations on second-period effort implies*

$$\begin{aligned} E_{y_1,y_2}(e_2^{M*}) &= E_{y_1,y_2}(cE_{y_2}U_3) \\ &= cE_{y_1,y_2}(E_{y_2}U_3) \\ &= cE_{y_1,y_2}U_3 \\ &= \frac{1}{R}e_1^{M*}. \end{aligned} \tag{11}$$

Moving from the second to third equality requires invoking the law of iterated expectations. Since $R < 1$*, we conclude that expected within contract effort is increasing as we move towards the termination date.*■

A.3 Proof for Proposition 2:

Proposition 2 *First-period effort under a single-period contract is greater than first-period effort under a multi-period contract;* $e_1^{S*} \geq e_1^{M*}$*.*

Proof. *The first-order condition for* e_1^S*, we see that*

$$e_1^{S*} = \frac{R}{2}\frac{\partial E_{y_1,y_2}U_2}{\partial e_1^2} + \frac{R}{2}\frac{\partial E_{y_1}e_2^S(y_1)}{\partial e_1^2} + \frac{R^2}{2}\frac{\partial E_{y_1,y_2}U_3}{\partial e_1^2} + \frac{R^2}{2}\frac{\partial E_{y_1,y_2}e_3^S(y_1,y_2)}{\partial e_1^2}.$$

As above,

$$\frac{\partial E_{y_1,y_2}U_3}{\partial e_1^S} = E_{y_1,y_2}\left(\frac{\kappa_1 h_z}{h_a+2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{S*} + y_2 - e_2^{S*})}{h_a+2h_z} + e_3^{S*}\right)\right]\right). \tag{12}$$

Using (12) and applying the equilibrium conditions that $e_1^S = e_1^{S*}$*,* $e_2^S = e_2^{S*}$*, and* $e_3^S = e_3^{S*} = 0$*, the first-order condition for first-period effort becomes*

$$\begin{aligned} e_1^{S*} &= \frac{R}{2}\frac{\partial E_{y_1,y_2}U_2}{\partial e_1^2} + \frac{R^2}{2}E_{y_1,y_2}\left(\frac{\kappa_1 h_z}{h_a+2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1+\bar{z}_2)}{h_a+2h_z}\right)\right]\right) \\ &= \frac{R}{2}\frac{\partial E_{y_1,y_2}U_2}{\partial e_1^2} + e_1^{M*}. \end{aligned}$$

This implies

$$
\begin{aligned}
e_1^{S*} - e_1^{M*} &= \frac{R}{2}\frac{\partial E_{y_1,y_2}U_2}{\partial e_1^S} \\
&= \frac{R}{2}\frac{\partial}{\partial e_1^S}E_{y_1,y_2}\left(-\frac{1}{\alpha}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{S*})}{h_a + h_z} + e_2^{S*}\right)\right]\right) \\
&= \frac{R}{2}\frac{\kappa_1 h_z}{h_a + h_z}E_{y_1,y_2}\left(\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{S*})}{h_a + h_z} + e_2^{S*}\right)\right]\right) \\
&> 0.
\end{aligned}
$$

The last term is true since a positive constant multiplies an exponential function, which is always positive.■

Proof for Proposition 3:

Proposition 3 *Effort in the second period of the agent's life is the same regardless of whether the agent is in a one-period contract or in the second period of a two-period contract;* $e_2^{1*} = e_2^{2*}$.

Proof. *Using the procedure outlined in Case 1, the change in expected third-period utility due to a marginal adjustment in second-period effort is a modified version of equation (10)*

$$
\frac{\partial E_{y_2}U_3}{\partial e_2^S} = E_{y_2}\left(\frac{\kappa_1 h_z}{h_a + 2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(y_1 - e_1^{S*} + y_2 - e_2^{S*})}{h_a + 2h_z} + e_3^{S*}\right)\right]\right). \tag{13}
$$

After applying all equilibrium conditions, the first-order condition for second-period effort is

$$
e_2^{S*}(y_1) = \frac{R}{2}E_{y_2}\left(\frac{\kappa_1 h_z}{h_a + 2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1 + \bar{z}_2)}{h_a + 2h_z}\right)\right]\right)
$$

which is equal to e_2^{M*} *by (3).*■

Appendix B

To find a closed-form solution for effort to be used in the numerical simulation, we must first integrate over a series of random variables. In what follows we give an example of how we avoid some of the complications involved with multiple integration.

Using equation (2),

$$
e_1^{M*} = \frac{R^2}{2}E_{y_1,y_2}\left(\frac{\kappa_1 h_z}{h_a + 2h_z}\exp\left[-\alpha\kappa_1\left(\frac{\mu h_a + h_z(\bar{z}_1 + \bar{z}_2)}{h_a + 2h_z}\right)\right]\right).
$$

we can now explicitly solve for the first-order condition for first-period effort.

Because we must take expectations over 2 periods of outcomes (net of effort), we must integrate over 3 random variables. Let J_1 be a new random variable, defined as the combination of all three so that $J_1 = 2a + z_1 + z_2 \sim N(2\mu, 4\sigma_a^2 + 2\sigma_z^2)$. We now need only to take expectations over J_1

$$e_1^{M*} = \frac{R^2}{2} \int \frac{\kappa_1 h_z}{h_a + 2h_z} \exp\left[-\alpha\kappa_1 \left(\frac{\mu h_a + h_z J_1}{h_a + 2h_z}\right)\right] f(J_1) dJ_1.$$

Splitting the random variable from the rest,

$$e_1^{M*} = \frac{R^2}{2} \frac{\kappa_1 h_z}{h_a + 2h_z} \exp\left[-\alpha\kappa_1 \left(\frac{\mu h_a}{h_a + 2h_z}\right)\right] \int \exp[t \cdot J_1]\, f(J_1) dJ_1$$

in which we define $t = \frac{-\alpha\kappa_1 h_z}{h_a + 2h_z}$. Finally, note that the integral term is the moment generating function for a random variable, J_1. Namely, since J_1 is a normally distributed random variable with mean 2μ and variance $4\sigma_a^2 + 2\sigma_z^2$, $M_{J_1}(t) = E(\exp[tJ_1]) = \int \exp[tJ_1]\, f(J_1) dJ_1 = \exp[2\mu t + \frac{t^2}{2}(4\sigma_a^2 + 2\sigma_z^2)]$. By rejoining the terms under one exponential function we get

$$e_1^{M*} = \frac{R^2}{2} \frac{\kappa_1 h_z}{h_a + 2h_z} \exp\left[-\alpha\kappa_1 \left(\frac{\mu h_a + h_z(2\mu + t(2\sigma_a^2 + \sigma_z^2))}{h_a + 2h_z}\right)\right].$$

We can mimic this procedure to find that a closed-form solution for optimal second-period effort is

$$e_2^{M*}(y_1) = \frac{R}{2} \frac{\kappa_1 h_z}{h_a + 2h_z} \exp\left[-\alpha\kappa_1 \left(\frac{\mu h_a + h_z(\bar{z}_1 + \frac{h_a\mu + h_z\bar{z}_1}{h_a + 2h_z} + \frac{t}{2}\left(\frac{1}{h_a + h_z} + \sigma_z^2\right)}{h_a + 2h_z}\right)\right].$$

We can do the same for effort under a series of 1-period contracts so that we have straightforward equations for effort which depend only on the underlying parameters.