Tulane Economics Working Paper Series

Employees' Performance Variation over Fixed-Term Contracts - Evidence from the National Hockey League

Furmaco, L.
Department of Economics
and the Murphy Institute
Tulane University
IZA
GLO
lfumarco@tulane.edu

Longley, N.
Department of Business
Nevada State College

Palermo, A.
Institute for Labour Law and
Industrial Relations in the
European Union

Rossi, G.
University of Birkbeck
London

Working Paper 2107
May 2021

Abstract

We investigate whether employees vary their performance during fixed-term contracts. We follow National Hockey League players' performance over ten seasons. We use a two-stage least square fixed effect model to address empirical limitations in previous studies. We find that players' performance varies at the end of the contract depending on ability, tenure, and (geographical) willingness to move. In particular, long-tenure and low-ability short-tenure workers vary their performance, depending on their continent of origin; these results might be due to different willingness to move, at different stages of players' career.

Keywords: fixed-term contracts, incentives, shirking behavior, strategic behavior
JEL codes: D82, J24, J33, M52, Z22

Employees' Performance Variation over Fixed-Term Contracts - Evidence from the National Hockey League

Fumarco, L.,[1] Longley, N.,[2] Palermo, A.,[3] Rossi, G.[4]

[1] Department of Economics and the Murphy Institute, Tulane University, IZA, GLO

[2] Department of Business, Nevada State College

[3] Institute for Labour Law and Industrial Relations in the European Union

[4] University of Birkbeck, London

Abstract. We investigate whether employees vary their performance during fixed-term contracts. We follow National Hockey League players' performance over ten seasons. We use a two-stage least square fixed effect model to address empirical limitations in previous studies. We find that players' performance varies at the end of the contract depending on ability, tenure, and (geographical) willingness to move. In particular, long-tenure and low-ability short-tenure workers vary their performance, depending on their continent of origin; these results might be due to different willingness to move, at different stages of players' career.

JEL-Classification: D82, J24, J33, M52, Z22
Keywords: fixed-term contracts, incentives, shirking behavior, strategic behavior

* Corresponding author, e-mail: lfumarco@tulane.edu

Acknowledgements: For their insightful feedback, we thank participants to the IAAEU Workshop on Labour Economics 2019, to the Sport Economics and Sport Management Conference 2019, to the WEAI/NAASE 1st Virtual Conference 2020, and to the SEA conference 2020. We owe a debt of gratitude to Bruno Rodrigues for the excellent data collection work. We thank Clemens Buchen and Bryce VanderBerg for their feedback in early stages of this paper.

1 Introduction

Employers try to align employees' incentives to their goals with formal agreements; that is, contracts. Despite these formal agreements, working relationships may still be characterized by moral hazard (Holmstrom, 1979): although employees' performance is potentially verifiable, their real effort is not observable. Efforts are made to design contracts so that distortions are minimized; however, rent efficiency trade-offs lead to suboptimal allocations, and to efficiency and welfare loss. Moreover, contracts are often fixed-term, do not change and are not renegotiable during their duration. Therefore, overall, contracts provide scope for employees' performance manipulation.

Against this background, we study two policy relevant research questions; that is, whether and how employees' performance varies along fixed-term contracts. Although these questions are not new, past literature provides ambiguous evidence due to some empirical limitations. We reconcile previous results by investigating a large dataset that allows us to follow employees over time and to address endogeneity. Thus, we explicitly address some of the main empirical limitations in the previous literature.

More concretely, we investigate how performance varies both at the beginning and at the end of the contract. We focus mainly on the latter stage, which has been traditionally neglected.

We set up a tailored model with moral hazard, which allows us to separate individuals based on their abilities using data information. Moreover, we account for relevant individual characteristics, namely ability, tenure, and willingness to (geographic) mobility to find a job.

Our results show that, while workers tend to perform with stability over the entire contract, different behaviours emerge at the end of the contract depending on workers' characteristics.

More in detail, our study adds to the literature on the employer-employee relationship in the presence of moral hazard by investigating two theoretical aspects. First, we study the evolution of employees' performance over the length of a fixed-term contract. Whereas theoretical contributions have studied the effect of abilities on performance in the presence of asymmetric information, less emphasis has been given to the dynamics of effort and performance over the entire duration of fixed-term contracts. Most of the literature focuses on *shirking behavior* (i.e., performance decrease at the beginning of a multi-year contract), while less attention is paid to strategic behavior (i.e., performance increase at the end of the contract).

Second, we investigate how strategic behavior varies with workers' incentives and outside options.

To overcome the traditional scarce availability of workers' productivity, we use a unique large panel data set on National Hockey League (NHL) players over ten seasons (2007/08-2016/17).

The prediction that employees who sign fixed-term contracts face low incentives to perform with stability has been extensively studied in professional sports because of the presence of publicly available high-quality data (Bryson et al., 2015; Buraimo et al. 2015; Purcell, 2009; Kahn, 2000). However, whether this prediction is mirrored by actual behavior is disputed (Buraimo et al 2015; Feess et al., 2015; Landry et al., 2015; Frick, 2011; Krautmann & Donly, 2009; Purcell, 2009; Stiroh, 2007; Berri & Krautmann, 2006; Maxcy et al., 2002; Oyer, 1998; Woolway, 1997; Sommers, 1993; Krautmann, 1990; Lehn, 1984).

Maxcy (2018) suggest that contradictory results may be driven by empirical problems. We identify, and address, three main empirical problems. First, contract length is endogenous: the best performers are offered longer contracts, and players with a long contract perform consistently high (Frick, 2007; Buraimo et al., 2015). Second, most studies are based on relatively small samples; so, they face low statistical power and over-weight contract length outliers (Buraimo et al., 2015). Third, to the best of our knowledge, none of the previous studies follows individuals over time. Individual observations are pulled over time: data are investigated with repeated cross-section techniques, in lieu of panel data techniques that control for time-invariant individual characteristics. Our study addresses these three empirical problems.

Studies of moral hazard in North American sports might reach contradictory results due to a fourth endemic problem: differences between different free-agency types are neglected. In leagues like the NHL, the end of the contract leads to two different types of free-agency, depending on the athlete's tenure in the league: restricted or unrestricted free-agency. Based on free-agency rules, long-tenure players are free to move to another team at the end of their contract (i.e., unrestricted free-agents, UFAs), while short-tenure players are not free in this respect (i.e., restricted free-agents, RFAs), as their current team may retain negotiating rights. While both types of free-agents may have an incentive to try to improve their performance in the last year of the contract to secure a new one, the strength of these incentives may be lower for some RFAs. Additionally, there could be differences between players with different potential mobility;

European skaters may benefit from additional outside options: they are more likely to go play in Europe, if they do not obtain a contract they like. Thus, they face lower incentives to improve performance at the end of the contract.

Our investigation proceeds in four steps. First, we study shirking behavior and, second, we investigate strategic behavior. Third, we examine whether strategic behavior differs between UFAs and RFAs, and whether RFAs' behavior differs based on their ability. Fourth, we investigate whether behaviors of UFAs and RFAs vary between North American and European NHL players.

The remainder of the paper proceeds as follows. Section 2 reviews the previous literature on performance variation over the duration of a contract. Section 3 provides an overview of the institutional rules governing free agency in the NHL; this section provides details on the reasons why skaters with different free-agency status might behave differently. Section 4 illustrates a framework used to define a proxy of ability. Section 5 is divided in three subsections to discuss the dataset, the descriptive statistics, and the instrumental variable. Section 6 discusses the methodology and conducts various analyses on shirking behavior as well as strategic behavior. Section 7 concludes by discussing the findings.

2 Literature review

The performance variation over fixed-term contracts is rarely empirically studied in the "standard" labor market, because direct measures of productivity and contracts features are rarely observable. Asch (1990) exploits a points system adopted by the Navy to study workers' incentives on their productivity. This study finds evidence that workers increase their effort when rewards are within reach, while they decrease it when rewards are surely achieved. Oyer (1998) finds that salespeople try harder to reach their quotas when the payment period is approaching, that is, the end of the fiscal year.

The performance variation over fixed-term contracts is more frequently studied in the sports labor market and comes along with the study of free-agency. As illustrated by Maxcy (2018), since the 1970s the advent of free agency in professional sport determined two key phenomena, as team sport athletes were finally free to negotiate their contracts with multiple franchises. First, expected salary inflation affected less restricted players' labor markets, increasing players' compensation. Second, long-term contracts with guaranteed contract earnings became widespread, as player remuneration started being determined ex-ante based on expected performance. Following the seminal articles by Lehn (1982), both phenomena attracted the interest of economists in two distinct areas of research. While the first one is related to the reasons behind the award of long-term contracts to professional athletes, the other one focuses on the impact that this type of contract has on the relationship between performance and incentives. In relation to this last research area, the literature has expanded quite rapidly in the last decades due to the increased availability of more detailed performance and pay data and the knowledge of contract clauses. Within this vast literature, however, the empirical evidence is contrasting (Solow & Von Allmen, 2016; Maxcy, 2018). Additionally, there is a discrepancy between theoretical predictions and empirical observability of the results that is often explained by using a "more" behavioral approach.

Until the mid-2000s, most of the published economic research focused on the Major League Baseball (MLB) in the US (Maxcy, 2018). This is the most suitable professional sport for testing the shirking hypothesis due to the fact that the aggregation of the individual performances of each player determines the team results (Scully, 1974). A baseball player's performance directly reflects his effort during the match, minimizing the interaction effect of teammates on his individual performance. Then, baseball players' individual performances, contract length, and salary are respectively highly quantified, disclosed and available in exceptional detail to the public. Using a sample of the earliest cohort of MLB free agents in 1980, Lehn (1982) examined the effect of long-term labor contracts on player "durability" and found evidence of opportunistic behavior. According to this evidence, long-term contracts increase the amount of time spent on the injured or disabled list because guaranteed multiyear contracts reduce the incentives for players to invest in proper physical conditioning. For Lehn, this effect can be mitigated by the inclusion of incentive bonuses in player contracts. In a later paper, Lehn (1984) distinguished between MLB players that signed for at least three years with the same team and players that signed the same length of contract with a different team. The evidence was that the first group spent more days on the disabled list during the next season. Maxcy et al. (2002) utilized a more recent sample of free agent MLB hitters and pitchers. They looked for ex-ante strategic behavior, defined as improved performance in the year before a new contract was signed, and ex-post strategic behavior (i.e., an alternative term for shirking behavior). In contrast to Lehn (1982, 1984), Maxcy et al. (2002) found that pitchers and hitters spent less time on the disabled list in the season immediately after the end of successful contract negotiations. Moreover, there was no sign of ex-post opportunistic behavior as performance was not statistically different in the year before or after the new contract. Playing time was also found to be above the average in the same season after the contract was signed. The authors concluded that it is necessary to introduce incentive-compatible mechanisms in a contract to minimize the inefficiencies created by shirking behavior.

Using the same sample of MLB hitters, Krautman (1999) and Scoggings (1993) found contrasting evidence of shirking behavior after signing a long-term contract. While Krautman (1999) argued that the observable performance variation was mainly the result of a stochastic process, Scoggings (1993) found that evidence of shirking behavior varied according to the choice of the performance measure. He revealed that over some indicators, players with long-term contracts exhibited lower productivity in the first year of their contract than in the preceding year. Specifically, shirking behavior by MLB hitters was noticed when players' performances were measured in terms of total bases instead of slugging average. Focusing on MLB hitters' offensive production, Marburger (2003) compared players' behavior before and after the end of the reserve system.[1] In contrast to Lehn (1982, 1984), he found evidence of shirking behavior adopted by poorly paid players only before the end of the reserve system and that monetary incentives in the free agency period pushed players to put in more effort. Maxcy (1997) and Fort and Maxcy (1998) also investigated hitters' and pitchers' performances in the MLB. They found that players with long-term contracts and with the option of re-contracting at the end of the season at the time did not adopt shirking behavior—their performances did not deteriorate. For other players, there was no evidence of shirking when their contracts were about to expire—their performances did not change.

Comparing MLB players' marginal productivity, Woolway (1997) found that players with long-term contracts tend to shirk by reducing individual performance a little. Although this decline is minimal, he argued that it might become economically significant at an aggregate level, when many players from the same team shirk. Comparing two different performances of a ballplayer's offensive production, a pure performance metric and an estimation of player's marginal revenue product (MRP), Krautmann and Donley (2009) found evidence that long-term contracts may lead to free agents underperforming in the first season of their new contract only

[1] During the reserve system, the team owners reserved the rights to players, even after their contract expired.

when MRP is used as a proxy. They concluded that the analytical approach used in the analysis can have a significant impact on the shirking tests. This evidence was confirmed by O'Neill (2014), who highlighted the critical importance of both the econometric method and the selection of player's performance measure.

Two studies on long-term contracts in the National Basketball Association (NBA) also reach conflicting conclusions. Berri and Krautmann (2006) looked for evidence that supported ex-post shirking behavior and argued that the MRP metric better represented the productivity of an individual player and it was used to measure player efficiency. Their evidence refuted the shirking hypothesis and concluded that findings on whether opportunistic behavior took place in professional basketball depended largely on the measure of player productivity chosen. They hinted that the conventional methods of measuring player efficiency in the NBA could not fully capture a player's contribution to his team's success. In another paper, Stiroh (2007) used contract data from the 2000-01 NBA season and individual player statistics from 1988 through 2002. He suggested that imperfect information and multi-year contracts created an implicit incentive for employees to strategically alter their effort over the contract cycle. He found strong evidence of opportunistic behavior using a composite measure of player performance and that performance both, increased in the contract signing year and decreased in the following year. His analysis strongly supported the view that there was an increase in performance during the contract year, but the evidence of lower performance after the contract was weaker, instead he showed that player performance regressed to the long run mean. In other words, these results appear to confirm that there are ex-ante strategic behavior hypotheses, but not that players also display shirking behavior after signing a new contract.

As most of the National Football League (NFL) contracts are not guaranteed, franchises are not used to offering long-term contracts. For this reason, this league is not always suitable for testing shirking behavior and the number of studies is limited. Gramm and Schnell (1994, 1997) studied long-term contracts of NFL players at the time of the 1987 strike associated with the demand for free agency by the National Football League Players Association. NFL players' average career duration was rather short and for this reason, players with long-term contracts were less likely to obtain the benefits of free agency. They were therefore more interested in maximizing their current incomes and less likely to go on strike. Conlin (1999) and Conlin and Emerson (1999) used the same database to study long-term contracts in the NFL. Their results indicated that when rookie players signed their first contract after the training camps, the athletic tests over their abilities provided clear performance differentials. However, they also found that effort was influenced by the remaining duration of a contract; the number of games started increased in the last contract year. In another study over long-term contracts, Frick et al. (2002) found evidence that the payment of signing bonuses that were not performance related persuaded free agent players to perform opportunistically; a high percentage of signing bonuses negatively affected team performance.

Finally, to the best of our knowledge, only two studies have investigated the effects of long-term contracts in the NHL. Landry et al. (2015) found suggestive evidence of performance variation, with players under-performing in the first season and improving their performances at the end of the contract. Differently, Purcell (2009) explored only goalies with multiple seasons of service in the NHL and found evidence of shirking in the first season of contract, but no evidence of strategic behavior in the last season.

Beyond North American sports, the effect of longer term contracts on performance is studied in European soccer as well. Frick (2011) and Feess et al. (2015) found evidence of strategic behavior: footballers from the Bundesliga tend to play better when their contract close to expiration than at the beginning. Also Buraimo et al. (2015) investigated the Bundesliga and found evidence that long-term contracts improve performance. This result is in contrast with the principal-agent theory, and they explain it with the fact that if remaining contract length is positively associated with the measure of performance being studied, selection effects dominates moral hazard effects (i.e., negative effects of contract extensions on performance).

3 Free Agency in the NHL

In this section, we describe institutional rules governing free-agency. It serves us to make the case of why players with different free-agency status at the end of the contract might behave differently.

In North American sports, free agents are those players that can escape the monopsonistic control of their current employer and freely sign with another club. In leagues like the NHL, however, not all free agents are created equally.

The rules governing free agency in the NHL are contained in the collective bargaining agreement (CBA) that the league has with its players. These rules are extensive and complex;

however, we can simplify them and focus on the distinction between players who are either unrestricted or restricted free agents.

To explain further, players with at least seven seasons of experience, or who are at least 27 years old, become UFAs and are free (upon the expiration of their contract) to move to other clubs without impediment.[2] These players are able to benefit from the open-market and competitive bidding for their services, allowing them to extract maximum value for their services.

For RFAs, the process is more complex. Players become RFAs if their current team presents them a qualifying offer;[3] this type of proposed contract renewal allows the team to retain negotiating rights.[4] If the qualifying offer is rejected, the player remains a RFA, whereas in the absence of a qualifying offer, the player becomes an UFA.

Can other teams propose a contract to a player who is still an RFA? Yes, an RFA can sign an "offer sheet" with any other NHL club. However, this does not mean that the player will automatically move to that club, since a player's current club can retain the player by simply matching the offer. Thus, RFAs are generally more expensive and their market is shallower; other clubs often question whether it is worth the effort and expense of even making an offer in light of what amounts to a type of "veto-power" retained by players' current club.

RFAs have a shallower market due to unwritten rules as well: traditionally, NHL general managers do not sign other clubs' RFAs. One of the most defining moments in the history of RFAs occurred in July 2007, when Edmonton Oilers general manager Kevin Lowe signed Anaheim Ducks' forward Dustin Penner to an offer sheet. Ducks general manager Brian Burke was livid at Lowe, presumably because Lowe violated the long-standing unwritten rule amongst

[2] Rules for unrestricted free-agency did not change after the 2012-2013 lockout (Staudohar, 2013), that is, within the period in analysis.
[3] A qualifying offer is a contract with a length of at least one season and with a salary that is at least the same as the previous one.
[4] Formally, when a short-tenure player, who is at the end of his contract, receives a qualifying offer by June 29th—of any given year—he becomes a RFA. On the same date, he can start new contract discussions with rival teams, even if he has received a qualifying offer, as long as he did not accept it; however, in that case, he can leave for a rival team only with the current team's permission. Therefore, if the player has not received a qualifying offer by July 1st of the same year, this player becomes an UFA. If the player has not obtained a contract renewal, or a contract with a new team, by December 1st of the same year, he cannot be hired by any other team, and thus will not play in the NHL during that season.

NHL general managers of never signing another team's RFAs. Burke did not match the offer, and Penner moved to the Oilers.[5]

Finally, clubs do not traditionally sign other club's RFAs, because they are more expensive. The compensation, in terms of draft pick, that the new club must pay to the player's current club is often onerous. For example, in the Dustin Penner case, Edmonton had to give up a 1st, 2nd, and 3rd round pick to Anaheim in exchange for Penner. NHL clubs highly value early-round draft picks, and being forced to lose these picks is generally a trade-off general managers are unwilling to make.

Additional CBA rules may counterbalance the veto-power of RFAs' current teams, but these rules pave long and arduous routes. Would-be RFAs can elect to take the club to arbitration, rather than declaring free agency; to simplify, these players can do that only after the completion of four NHL seasons. This means that players who signed the mandated three-year entry level contract upon entering the league[6] are not eligible for arbitration until the expiration of their *second* NHL contract. Thus, most arbitration-eligible players tend to be at least 23 years old. Arbitration is also a very adversarial process, so players are often reluctant to go to arbitration: they know they will play for the same club in the future, and do not want to damage the employment relationship.

In summary, while RFA status seems to confer increased bargaining to players, the institutional and cultural realities of the NHL bring in question whether this is actually true. By extension, if players do not see their RFA status as being particularly valuable, they may have little incentive to strategically increase performance in the final year of their current contract.

4 Theoretical framework

In this section, we develop a model that helps us differentiate players based on their different skills. This is relevant for the heterogeneity analyses on strategic behavior.

A player chooses effort $e \in [0, \bar{e}]$ with $\bar{e} < 1$ at the cost of $c(e; D) = \frac{e^2}{2D}$, where $D > 1$ represents his ability. From the effort, the player generates a value $V(e)$ for the team. We assume

[5] To show the rarity of such an occurrence, Penner is still the only RFA in the past 23 years (since 1997) to change NHL clubs. As revealingly, when Sebastian Aho signed an offer sheet (that was matched) in July 2019, he was the first player to even *sign* an offer sheet in over *six* years, let alone ever change clubs.

[6] Players who entered the league between the ages of 18 and 20 (which includes that majority of players) are required to sign a three-year entry-level contract.

that the generated value is $V(e) = e + \varepsilon$ with ε uniformly distributed $\varepsilon \sim U\left(-\frac{1}{D}, \frac{1}{D}\right)$, representing the component for (bad) luck. Hence, the expected value of ε is zero and its variance depends on the ability of the player, $Var[\varepsilon] = \frac{1}{3D^2}$; thus, the more able is the player, the lower is the relative impact of the (bad) luck component.

The team manager proposes a wage $w(D)$ to the player. Ability is observable by the manager, and we assume that the wage depends on the ability and it is strictly increasing, $w'(D) > 0$. Hence, we consider that market forces drive up the wage of players with higher abilities. Moreover, the manager designs ex-ante a time invariant incentive scheme based on a bonus β, which is paid if the player reaches a target, that is, an observable performance threshold $\hat{V} \in (0,1]$; both β and $\hat{V}$ are endogenously derived.

Formally, we denote with u and π the utility of the player and the team respectively.

$$u = Pr\left(V \geq \hat{V}\right)\beta - \frac{e^2}{2D} \qquad (1)$$

$$\pi = V(e) - Pr\left(V \geq \hat{V}\right)\beta \qquad (2)$$

The manager chooses the contract to maximize the expected value for the team:

$$\max_{\{\beta,\hat{V}\}} \mathbb{E}\left[V(e) - Pr\left(V \geq \hat{V}\right)\beta \mid \mathrm{e}\right] \qquad (3)$$

subject to

$$\tilde{e} = arg\max_{e} \mathbb{E}\left[Pr\left(V \geq \hat{V}\right)\beta - \frac{e^2}{2D} \mid \mathrm{e}\right] \qquad \text{(IC)}$$

The incentive compatible constraint (IC) accounts for the player's incentive to exert the effort that maximizes his expected utility.[7] Given the distribution of (bad) luck, we have $Pr\left(V \geq \hat{V}\right) = Pr\left(\varepsilon \geq \hat{V} - e\right) = \frac{1+D(e-\hat{V})}{2}$ and hence the IC can be reduced to $\hat{e} = \frac{\beta D^2}{2}$.[8]

[7] Note that Equations (1) to (3) do not include $w(D)$ because the fix wage does not affect the incentive compatible constraint.

[8] We assume that D is not too large to avoid that the probability $Pr\left(V \geq \hat{V}\right)$ is equal to one.

Lemma 1 *The optimal incentive contract is such that the target is as large as possible and the performance-related bonus when this target performance is reached is* $\beta = \frac{D(D+\hat{V})-1}{D^3}$.

Proof. See Appendix A ||

Corollary 1 *The optimal performance-related bonus is a decreasing function of ability.*

Proof. See Appendix A ||

We define $\mathcal{N} = \frac{\beta(D)}{w(D)}$ as the player's superstar status; the smaller is $\mathcal{N}$, the higher is the player's ability. In fact,

$$\frac{\partial \mathcal{N}}{\partial D} = \frac{\beta'(D)w(D) - \beta(D)w'(D)}{w(D)^2} < 0, \qquad (4)$$

because $w'(D) > 0$ and as shown in corollary 1 $\beta'(D) < 0$.

It is important to remark that in Section 6.3 we investigate UFAs alone, but we do not divide them into two subsamples, based on their ability. They cannot receive a performance-related bonus, so this model cannot be applied to UFAs, and, thus, we do not conduct heterogeneity analyses on them. Then, one could wonder why we do not split RFAs and UFAs based on other criteria, such as wage only. The reason is that wage does not reflect only performance potential, but additional aspects as well, such as the ability to enrich the team by selling personal merchandise (e.g., jerseys with the player's name on them).

5 Data

5.1 Dataset

We collected players' data from the website nhlnumbers.com. Our data set is extremely large when compared to those of traditional papers that investigate moral hazard in the employer-employee relationship; in fact, our data cover a ten-season period, from 2007/2008 to 2016/2017. We conduct our study on 6,304 player-season observations, that is, 1,618 individual skaters; 93 players appear in the data set for all of the ten seasons, whereas 385 appear only once.

This is the general profile of our players and their performance. We focus on regular season performance of players who hold multi-year contracts. We focus on the regular season only for one main reason: during playoffs, players benefit from additional monetary incentives, so their performance per match could not be comparable with that over regular season. Our sample includes players who changed teams during the season. For instance, they passed through waivers or were traded; these players are assigned a fictitious team that we call "*Change*" and their statistics are cumulated through teams in the same season and contract details are those agreed with the team that owns them in the first part of the season.[9] Moreover, we eliminate players who have not played a single game during the season, as they do have any performance statistics because they lack on-ice time. We include players of all nationalities.[10] Finally, we do not investigate goalkeepers, because their performance is proxied by measures that are not comparable to those of other players.

What measure of performance do we focus on? We investigate three measures. First, *Points per game*, which measures the sum of an individual player's seasonal goals and assists, averaged across the amount out seasonal games; this is perhaps the most direct measure of the skaters' individual performance (Fumarco et al., 2017; Landry et al., 2015; Idson & Kahane, 2000). Second, we study the seasonal *Plus-minus per game,* which is meant to measure the impact of an individual skater on the game;[11] although this measure is not traditionally considered in economics—except for Idson and Kahane (2000), it has for a long time been one of the fans' favorite measures of hockey players' productivity.[12] Third, we study *Penalty minutes per game*; this measure captures the player's intensity in the rink and, thus, his willingness to sacrifice for the team (Idson & Kahane, 2000). All of these outcomes are measured at the end of the season.

[9] For instance, a hockey player started the season with team A, with which he scored 3 points, then he changed to team B and scored 10 points, and eventually changed to team C scoring 5 points; this player has scored 18 points in seasons *t*, while the contract characteristics are those that he agreed to with team A.

[10] By definition, players with a one-season contract are not included. A one-year contract is essentially different from multi-year contracts—the latter can be disaggregated into three parts, namely first and last season as well as the seasons in between.

[11] This is given by the difference between the sum of "pluses", which are awarded to a player each time he is on ice and his team scores without having a "manpower advantage," and the sum of "minuses" which are awarded to a player each time he is on the ice and the opponent team scores a goal while being "shorthanded." When a player commits a penalty, this player has to temporarily leave the rink and go into a penalty box for a set number of minutes—depending on the kind of infraction (National Hockey League, 2017), and his team is said to be "shorthanded;" in this case, his team has a "manpower disadvantage," whereas the opponent team has a "manpower advantage" (see http://www.nhl.com/ice/page.htm?id=26374, June 27th, 2018).

[12] The "NHL Plus-Minus Award" was seasonally awarded from 1982/1983 to 2007/2008.

The focus of our investigation on shirking behavior is on the partial interaction between the discrete variable expressing the number of seasons that are *Remaining* until the end of the contract, and the dummy variable *First_season*; this partial interaction captures how the main effect of *Remaining* changes in the first season. While the variable *Remaining* goes from 0 (i.e., the last season of the contract, which is the reference value) to 14 (i.e., the longest contract length in this data set), the dummy *First_season* equals 1 when *t* corresponds to the first season of the contract. Notice that, since we have excluded hockey players with a one-season contract, *Remaining* cannot equal 0 when *First_season* equals 1; this is the reason why the econometric model does not include the main effect of *First_season*. There is evidence of shirking behavior if the estimated effect of *First_season* is negative.

To investigate strategic behavior in greater detail, we focus on how performance changes in the last contract period. Thus, we construct the dummy variable *Last_season*, that equals 1 when *t* corresponds to the last season of the contract.

5.2 Descriptive Statistics

Table 1 shows that *Remaining* is positively and significantly correlated to two outcomes. That is, with the increase in the number of seasons to the end of the contract (i.e., closer to the beginning of the contract), performance increases, as measured by *Point* and *Plus-minus per game*, whereas performance decreases while approaching the end of the contract. *First_season* is negatively associated with *Point* and *Plus-minus per game*. So, simple associations suggest possible shirking and strategic behavior.

Table 1. Pairwise correlations and descriptive statistics.

	Pairwise correlations							
Variables	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
1 Points	1.000							
2 Plus-minus	0.343***	1.000						
3 Penalty	-0.110***	-0.016	1.000					
4 Remaining	0.308***	0.058***	-0.018	1.000				
5 First_season	-0.136***	-0.050***	0.006	0.285***	1.000			
6 Last_season	-0.069***	0.0079	0.004	-0.545	-0.514	1.000		

7 Age	0.101***	0.075***	0.088***	0.026**	-0.192***	0.095***	1.000	
8 Superstar	0.139***	-0.024	-0.057***	-0.093***	-0.080***	0.045**	-0.421***	1.000
Statistics								
Mean	0.366	-0.030	0.624	1.524	0.461	0.236	27.308	0.268
SD	0.269	0.261	0.647	1.553	0.498	0.424	4.622	0.597
Min	0	-3	0	0	0	0	11	0
Max	1.681	3	17	14	1	1	42	4
Observations	6,304	6,304	6,304	6,304	6,304	6,304	6,279	3,105

Note: Point, Plus-minus, and Penalty Minutes are "per game." SD stands for standard deviation.

Additionally, note the positive correlation between *Remaining* and *First_season*. The direction of the correlation is quite obvious: the first season of the contract corresponds to the largest number of seasons before the end of the contract.

Last_season has a negative and statistically significant correlation only with *Points per game*. This simple correlation suggests the reversal of strategic behavior, that is, a reduction in performance at the end of the contract.

The interpretations of these simple correlations should be considered with a grain of salt for two reasons. First, they do not account for age or experience in the NHL; the first season of a contract corresponds to a younger age—when players are less experienced and/or are not yet well integrated into the team, while the end of the contract corresponds to players' lower physical fitness. Second, there is selection bias: at any given moment, players with long contracts face more seasons to the end of the contract (i.e., the value of *Remaining* is high), and these players are usually the best ones.

Econometric analyses control for additional characteristics. In particular, we control for *Age*, which captures both the effect of biological age and of seasons of service in the NHL. Table 1 provides a statistically significant and positive correlation between *Age* and the performance measures, except for *Plus-minus*. *Age* is positively correlated to remaining seasons as well because more experienced players receive longer contracts. Finally, *Age* is negatively correlated to *First_season*; since older players are offered longer contracts, in any given moment, there are fewer older players in their first season of contract (e.g., it is more likely that a 20 year old player is in his first season of contract than a player who is 23). For the opposite reason, there is a positive correlation between *Age* and *Last_season.* Our econometric analyses account for age squared as well, to capture possible non-linear returns.

When we investigate the role of different incentive structures, we split the sample between UFAs versus RFAs; then, we further divide RFAs between what we call "non-superstar" and "superstar" players (Bryson et al., 2014; Adler, 1985; Rosen, 1981). The superstar status is given by the ratio of performance-related bonus/seasonal wage—both are expressed in millions of dollars. Larger values of this variable are associated with less-skilled players, because the bonus that the team can offer them is relatively larger than their salary. It is important to note that, since only RFAs are split based on their superstar status, the number of observations on this variable is lower: bonuses are offered only to players who would qualify as RFAs at the end of their contract. Table 1 provides statistically significant and positive correlations between *Superstar* and *Points per game*, while the correlation with *Penalty* is negative. It is not surprising to observe a positive correlation between *Remaining* and *Superstar*; superstars tend to have longer contracts.

Finally, econometric analyses account for a variety of fixed effects that are not included in Table 1. Our econometric model controls for unobservable season characteristics.[13] In order to account for the possible effect of the team on individual performance, we control for team fixed effects as well. Additionally, all our analyses control for time-invariant players' characteristics.

5.3 Instrumental variable

To address the endogeneity of *Remaining* (i.e., the remainder of the contract length), we follow Buraimo et al. (2015)'s strategy and adjust it to the recent developments in the relative age effects literature. We use teammates' average contract length as an instrument. The reason is that this instrument reflects two team-specific characteristics that are exogenous to the player *i*: (a) the team contract duration policy; and (b) the combination of skills and contract duration needed by the team. An additional advantage of this instrument is that it controls for free-riding behaviors, since it catches teammates' effect on a player's *i* performance. Moreover, we stratify this instrument by quarter of birth, which provides an exogenous variation in individual performance. In fact, performance varies by quarter of birth depending on two exogenous factors: player's date of birth and cut-off dates used to divide players in different youth categories (Fumarco et al. 2017; Gibbs et al., 2012). The usage of relative age as an instrument was first proposed in the specialized literature (Lenard & Peña, 2017).

[13] For example, there was a lockout in season 2012/2013, which shortened the competition season and created more uncertainty in terms of contract renewals.

More concretely, we construct an instrumental variable which represents the average contract length of skater *i*'s teammates, who are born in the same quarter and we call it simply *Team_length*.[14]

Is this instrument relevant? On theoretical grounds, the answer is yes. It represents a measure of the team's policy toward long-term contracts. What about empirical grounds? Also in this case, the answer is a yes. Table C.1 in the Appendix C shows the first stage results, while Section 6.1 discusses them.

Is this instrument excludable? On theoretical arguments, the answer is yes. It is intrinsically orthogonal to player's *i* productivity. What about empirical grounds? We can test its excludability with unconditional and conditional balance tests, while accounting for the most relevant time-variant skaters' characteristics: age, team, and season. The results are reported in Appendix B, Tables B.1 and B.2, and they suggest that *Team_length* is randomly distributed with respect to these characteristics; Appendix B reports a more detailed discussion of the results. Additionally, Table C.2 in Appendix C shows the reduced form results, while Section 6.1 discusses them; in summary, these results further support to the excludability of this instrument. Overall, theoretical and empirical arguments provide evidence of relevance and excludability of our instrument.

6 Methods and Results

This section is divided into four subsections. First, we investigate shirking behavior. Second, we investigate strategic behavior, regardless of free-agency or superstar status. Third, we investigate strategic behavior on UFAs, non-superstar RFAs, and superstar RFAs. Fourth, we conduct heterogeneity analyses to investigate whether UFAs, and both non-superstar and superstar RFAs from Europe behave differently in the concluding part of the contract.

6.1 Shirking Behavior

We conduct these analyses with a two stage least square (2SLS) regression model and individual fixed effects. Thus, differently from previous literature, we exploit both within and between variation; so, we follow skaters over time, while accounting for their individual fixed effects. Standard errors are clustered at an individual level.

[14] In constructing this variable, the contract length of skater *i* is thus excluded. Basic descriptive statistics on contract length, by season, are reported in Appendix B, Table B.1.

Equation (4) illustrates the second stage:

$$Y_{it} = \beta_0 + \beta_1 \widehat{Remaining_{it}} + \beta_2 \widehat{Remaining_{it} \times First\_season_{it}} + \beta_3 Age_{it} + \beta_4 Age_{it}^2 + \boldsymbol{\delta Season_t} + \boldsymbol{\gamma Team_{it}} + \varepsilon_{it} \tag{4}$$

$Remaining$ is endogenous, and thus also its interaction with *First_season* is endogenous. Therefore, this second stage regresses the outcome variables on the predicted values of *Remaining* and *Remaining* ×*First_season* that are obtained from the two first stages conducted with Equation (5).

$$Endogenous_{it} = \beta_0 + \beta_1 Team\_length_{it} + \beta_2 Team\_length_{it} \times First\_season_{it} + \beta_3 Age_{it} + \beta_4 Age_{it}^2 + \boldsymbol{\delta Season_t} + \boldsymbol{\gamma Team_{it}} + \varepsilon_{it} \tag{5}$$

Where *Endogenous* is *Remaining*, in one first stage, and *Remaining* × *First_season*, in the other first stage.

Both Equation (4) and (5) include control variables for age and its square to capture non-linear returns. Additionally, we control for fixed effects for season and team.

Table 2 reports the results from the second stage of the 2SLS for the three outcomes; first stage and reduced form results are reported in Table C.1 and C.2, in Appendix C.

Table 2. Remaining and first season effects on points, plus-minus, and penalty time per game; second stages 2SLS fixed effects estimates.

Variables	Points per game (1)	Plus-minus per game (2)	Penalty per game (3)
Remaining ×	-0.040	0.004	0.021
First_season	(0.059)	(0.069)	(0.130)
Remaining	0.065	-0.008	-0.036
	(0.096)	(0.111)	(0.210)
Observations	5,895	5,895	5,895

Note: All the analyses include age and its square, as well as team and season fixed effects. Standard errors clustered on skaters in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

This table does not provide evidence of shirking effect, since the interaction effect is never statistically significant. Moreover, these results do not suggest that performance changes while approaching the end of the contract, since *Remaining* does not have a statistically significant effect.

Note that the number of observations falls short of 409 observations, compared to what is discussed in Section 5. There are two reasons for that to happen: (i) as discussed in Section 5, there are 384 singletons (i.e., players that appear for only one season in the dataset) that cannot be used in a panel data analysis; and (ii) there are 25 missing values on age.

Table C.1 in Appendix C provides support for the relevance of the instruments. *Team_length* has a positive association with *Remaining*; the longer skater *i*'s teammates' contracts, the longer skater *i'*s remaining contract. *Team_length* × *First_season* has a negative association with *Remaining* × *First_season*; the more teammates are at the beginning of long contracts, the less likely skater *i* is at the beginning of a long contract as well. Both estimates are highly statistically significant.

Moreover, Table C.2 in Appendix C provide further support to the excludability of our instruments as they do not directly affect the three outcomes.

We use an alternative specification model, without the semi-interaction term and with a control variable for contract length in place of seasons to the end of the contract, we obtain equivalent results. In this alternative specification, contract length is instrumented with *Team_length*. These results can be provided upon request.

6.2 Strategic Behavior: Main Analyses

In this subsection, we investigate the entire sample and study strategic behavior as captured by a dummy variable, which equals 1 if it is the last season of the player's contract, that is, $Last\_season$.

Equation (6) illustrates the fixed effect model:

$$Y_{it} = \beta_0 + \beta_1 LastSeason_{it} + \beta_2 Age_{it} + \beta_3 Age_{it}^2 + \boldsymbol{\delta Season_t} + \boldsymbol{\gamma Team_{it}} + \varepsilon_{it} \quad (6)$$

Model (6) includes control variables for age and its square to capture non-linear returns. Additionally, we control for fixed effects for season and team.

Note that, while the specification model from Equations (4) and (5) already provides some insights on the possible existence of strategic behavior, the model in Equation (6) allows us to obtain a refined picture, on possible non-linearities: *Remaining* includes both the last season of the contract and the seasons between this season and the first one.

Table 3 reports the results.

Table 3. Last season effects on points, plus-minus, and penalty time per game; fixed effects estimates.

Variables	Points per game (1)	Plus-minus per game (2)	Penalty per game (3)
Last_season	0.004	0.013*	0.006
	(0.004)	(0.007)	(0.012)
Observations	5,895	5,895	5,895

Note: All the analyses include age and its square, as well as team and season fixed effects. Standard errors clustered on skaters in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table 3 provides statistical evidence of strategic behavior only on plus-minus per game.

These results could hide heterogeneous effects. Players with a different free-agency status might behave differently, and, among RFAs, behaviors might change depending on notoriety, that is a proxy for ability. Subsection 6.3 focuses on these heterogeneity analyses.

In an alternative specification model, we control for contract length, and we obtain equivalent results. In this specification, contract length is instrumented with *Team_length*. These results can be provided upon request.

6.3 Strategic Behavior: Players' Free-agency Status and Superstar Status

Past literature highlights conflicting results concerning the presence of strategic behavior. These results might be contradictory because players could behave differently based on both free-agency and superstar status. In particular, some free-agents may have countervailing incentives. Hence, we perform a heterogeneity analysis on three subsamples. First, we focus on UFAs; they can move to whatever team they want at the end of their contract. Then, we investigate two subsamples of RFAs. One of these two subsamples is composed of non-superstars, for whom the "superstar level" is larger than its mean (i.e., the bonus is relatively larger than the wage), while the other subsample is composed of superstars, for whom the "superstar level" is lower than its mean (i.e., the bonus is relatively is lower than he wage).[15] Hence, we divide the sample according to ability.

Table 4 reports the results.

Table 4. Last season effects on points, plus-minus, and penalty time per game. Fixed effects estimates, for UFAs, non-superstar RFAs, and superstar RFAs.

Variables	Points per game (1)	Plus-minus per game (2)	Penalty per game (3)
UFAs			
Last_season	0.002	0.021**	-0.011
	(0.005)	(0.009)	(0.019)
Observations	3,082	3,082	3,082
Non-superstar RFAs			
Last_season	-0.027	0.004	-0.097**
	(0.022)	(0.034)	(0.040)
Observations	729	729	729
Superstar RFAs			
Last_season	-0.010	-0.016	0.008
	(0.008)	(0.012)	(0.025)
Observations	2,084	2,084	2,084

Note: All the analyses include age and its square, as well as team and season fixed effects. Standard errors clustered on skaters in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 4 provides partial statistical evidence of UFAs' strategic behavior on *Plus-minus* per game. Moreover, there is partial statistical evidence of the opposite of strategic behavior for non-superstar RFAs, but only when we look at *Penalty minutes per game*. Finally, we do not observe evidence of strategic behavior on superstar RFAs.

Results in this section provide weak evidence of strategic behavior. However, there is still one intersectional aspect that is worth being studied. European skaters could behave differently because they face an alternative option that is characterized by a non-pecuniary aspect: they have higher mobility and could go back to Europe. North American counterparts might be less likely

[15] We obtain equivalent results when we use the median of *Superstar*, that is 0, as the pivotal value to define superstar and non-superstar players.

to consider the option of moving to Europe (i.e., they are characterized by lower international geographic mobility).[16] This is what we explore in the last subsection of these results.

6.4 Strategic Behavior: Players' Free-agency Status and Superstar Status, by Continent of Origin

The model specification used in this section and in Section 6.3 differs with respect to one detail only: here, we insert a semi-interaction between *Last_season* and a dummy for the skater being *European*. The main effect of *European* is missing, since it is eliminated by the within transformation.

Table 5 reports the results.

Table 5. Last season effects on points, plus-minus, and penalty time per game, by continent of origin. Fixed effects estimates, for UFAs, for non-superstar RFAs, and for superstar RFAs.

Variables	Points per game (1)	Plus-minus per game (2)	Penalty per game (3)
UFAs			
Last_season	0.007	0.021**	-0.016
	(0.006)	(0.010)	(0.022)
Last_season ×	-0.022*	0.001	0.026
European	(0.013)	(0.022)	(0.036)
Observations	3,082	3,082	3,082
Non-superstar RFAs			
Last_season	-0.059**	-0.027	-0.077*
	(0.025)	(0.037)	(0.042)
Last_season ×	0.098***	0.092	-0.060
European	(0.033)	(0.058)	(0.055)
Observations	729	729	729
Superstar RFAs			
Last_season	-0.001	0.005	-0.008
	(0.009)	(0.015)	(0.025)
Last_season ×	0.001	-0.012	0.046
European	(0.019)	(0.034)	(0.069)
Observations	2,278	2,278	2,278

Note: All the analyses include age and its square, as well as team and season fixed effects. Standard errors clustered on skaters in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

[16] Only recently, in 2020, North American players have started to move to Europe more frequently and at a younger age, but this is due to the COVID-19 pandemic being more under control in Europe: (https://www.eurohockeyclubs.com/news/europe-presents-another-dimension-for-young-north-americans; accessed on January 10th, 2021)

Table 5 provides several interesting insights. North American UFAs increase their performance in the last season, as measured by *Plus-minus per game*; they are expected to be less mobile than European colleagues and, therefore, they want to obtain a contract renewal. However, North American RFA players who are non-superstars reduce their performance, as measured by both *Points* and *Penalty minutes per game*. Perhaps, they reduce their performance due to demotivation: they are young and not that good, they know they do not have enough leverage to get a contract renewal, and they do not have elsewhere to go—they do not consider going to Europe.

European older players, that is, the UFAs, reduce their performance facing lower incentives in their last season of contract; if they are offered a bad contract, they can go back to Europe. Why might going back to Europe be an appealing option? There is an expected pay cut, but this is compensated for by non-pecuniary aspects, such as playing more and thus increasing the chances of still being selected by the national team. Conversely, younger lower-ability players, that is, non-superstar RFAs, increase their performance; they want to remain in the best league in the world instead of going back to Europe in the prime of their career. This might represent the signaling incentive. That is, they acquire credentials by working in high level institutions.

Finally, this table does not provide statistical evidence of strategic behavior or its opposite for either European or North American superstar RFAs.

In the concluding section, we further discuss these results.

7 Conclusion

This study contributes to the literature on the employer-employee relationship in presence of moral hazard with respect to two aspects. First, we investigate the evolution of employees' performance for the duration of fixed-term contracts; in particular, we explore how employees' performance varies at the beginning and at the end of a multi-year static contract. Second, we argue that preferences and/or signaling incentives might affect workers' performance at the end of the contract.

We use a unique and large panel data set on National Hockey League (NHL) players that are followed over ten seasons, where contracts between employers and employees (i.e., players) do not change for the duration of the contract. This large data set allows to conduct empirical improvements over past studies by exploiting both within and between variation, while accounting for individual fixed-effects. Moreover, we combine panel data techniques with an instrumental variable approach to cope with the endogeneity of contract length.

We find two set of main results. First, no evidence of shirking behavior. Second, while the analysis on the entire sample reveals no strategic behavior, some statistically significant results arise when we disaggregate players by contract status (i.e., UFAs, superstar RFAs, and non-superstar RFAs) and then when we investigate whether players' strategic behavior changes between North American and European players. There are no significant effects for superstar RFAs; this confirms that younger workers with higher abilities perform with stability. Conversely, we find evidence of strategic behavior among older and lower ability players. Non-superstars short-tenure European players improve their points-per-game in the final season of their contract. North American players' points-per-game decline. We argue that one possible explanation is that, although European players are more likely to move back to a European league, they are less willing to do so in the prime of their career, while non-superstar North American players might be willing to move to a lower North American league. In other words, European non-superstar RFAs might have stronger incentives to do their best. For UFAs, we find that European players worsen their performance in the last season of their contract, while North American players improve it. Also in this case, geographic mobility and non-pecuniary aspects might be an explanation, with European players being less motivated, since they could go back to Europe.

The lack of consistent evidence of shirking behavior and some evidence for strategic behavior could come for at least two main reasons. First, teams may value players that perform consistently, so players are disincentivized from engaging in any type of performance manipulation. Second, the relatively short length of the contract (and, overall, short careers) give players no opportunity to be on "cruise-control" through the middle years. In the period of this analysis, only a few star NHL players have signed very long contracts that would allow that.[17]

There are two additional reasons that are peculiar to the NHL labor market: (i) there is a hard salary cap, so being an UFA in hockey may not be advantageous, hence reducing the incentive in hockey for a large spike in effort during the last season—net of differences in

[17] For example, Alexander Ovechkin, 13 years, and Sidney Crosby, 12 years. Two other superstar players, Jonathan Toews and Patrick Kane, became well-known for signing long-term contracts one year before they were to become UFAs.

geographic mobility; (ii) institutional constraints are such that RFAs do not possess as much bargaining power as might be first expected, hence reducing their incentive to manipulate effort over the entire contract.

These results have real life policy implications. Based on these results, high-skilled workers with relatively short tenure and low market power owed to institutional rules and complex contracts are expected to perform with stability and as best as they can. Conversely, for less capable workers different levels of geographic mobility might shift some market power in their favor, because they can move from one local market to another, depending on personal preferences and circumstances. Junior scholars who seek tenure-track positions are a good example as, in this group of workers, only scholars with the highest mobility might be characterized by some market power and thus by a slightly more volatile performance over the contract duration. Finally, it could be possible to study shirking and strategic behavior in the labor market of tenure-track positions in academia and test whether ability, and potential differences in geographic mobility, matter. Details regarding contract length are usually disclosed in job ads, while some productivity measures (e.g., publications, quantity of courses being taught, and participation in conferences) are usually openly available on institutional or personal websites. While the creation of such dataset would be painstaking, it promises interesting results. Such data set would allow researchers to investigate different questions that are related to our study. To put it in context, the academic setting would allow one to study a peculiar scenario created by the ongoing COVID-19 pandemic: anecdotal evidence from specialized websites (e.g. econjobrumors and AEA/JOE) suggests that many young researchers from prestigious universities have crossed oceans and moved to less prestigious ones, following the negative shock in labor demand. Are they more likely to leave their institute at the end of the contract—instead of aiming for a tenured position—if the opportunity rises? Our results suggest that ability, geographic mobility, and incentives might play a role in this choice.

References

Adler, M., (1985). Stardom and talent. *The American Economic Review, 75*(1), 208-212.

Berri, D. J., Krautmann, A. C. (2006). Shirking on the court: Testing for the incentive effects of guaranteed pay. *Economic Inquiry*, *44*(3), 536-546.

Bryson, A., Rossi, G., Simmons, R. (2014). The migrant wage premium in professional football: A superstar effect? *Kyklos*, *67*(1), 12-28.

Bryson, A., Frick, B., Simmons, R. (2015). Sports Economics: It May Be Fun but What's the Point? *National Institute Economic Review, 232*, R1-R3.

Buraimo, B., Frick, B., Hickfang, M., Simmons, R. (2015). The Economics of Long-Term Contracts in the Footballers' Labour Market. *Scottish Journal of Political Economy*, *62*(1), 8-24.

Caliendo, M., Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys, 22*(1), 31-72.

Conlin, M. (1999). Empirical test of a separating equilibrium in National Football League contract negotiations, *Rand Journal of Economics*, 30(2), 289-304.

Conlin, M. Emerson, P. M. (1999). Double moral hazard and multi-dimensional separating equilibrium in contract negotiations: A study of the National Football League. Department of Economics, Cornell University, Mimeo.

Feess, E., Gerfin, M., Muehlheusser, G. (2015). Contracts as Rent-Seeking Devices: Evidence from German Soccer. *Economic Inquiry, 53*(1), 714-730.

Fort, R. Maxcy, J. (1998). Last period problem in sports. Department of Economics, Washington State University, Mimeo.

Frick, B., Dilger, A., and Prinz, J. (2002). Arbeitsmarktregulierung und nachvertraglicher Opportunismus: Die Verhaltensfolgen garantierter Handgeldzahlungen in der National Football League. *Zeitschrift für Betriebswirtschaft*, 72(4), Sportökonomie, 163-180.

Frick, B. (2007). The Football Players' Labor Market: Empirical Evidence from the Major European Leagues. *Scottish Journal of Political Economy, 54*(3), 422-446.

Frick, B. (2011). Performance, salaries, and contract length: Empirical evidence from German soccer. *International Journal of Sport Finance*, *6*(2), 87-118.

Fumarco, L., Gibbs, B. G., Jarvis, J. A., Rossi, G. (2017). The relative age effect reversal among the National Hockey League Elite. *PloS ONE, 12*(8), 1-16.

Gibbs, B. G., Jarvis. J. A., Dufur, M. J. (2012). The Rise of the Underdog? The Relative Age Effect Reversal Among Canadian-Born NHL Hockey Players: A Reply to Nolan and Howell. *International Review for the Sociology of Sport, 47*(5), 644-649.

Gramm, C. L. Schnell, J. F., (1994). The Empirical Relations between Employers' Striker Replacement Strategies and Strike Duration. *Industrial and Labor Relations Review, 47*(2), 189–206.

Gramm, C. L. Schnell, J. F. (1997). Following the leader: Race and player behavior in the 1987 NFL strike. *In W. Hendricks (eds), Advances in the Economics of Sports*, 2, 115-143.

Holmstrom, B. (1979), Moral Hazard and Observability. *The Bell Journal of Economics, 10(1),* 74-91.

Idson, T. L., Kahane, L. H. (2000). Team effects on compensation: an application to salary determination in the National Hockey League. *Economic Inquiry, 38*(2), 345-357.

Landry, J., Edgar, D., Harris, J., Grant, K. (2015) National Hockey League Guaranteed Contracts: A Principal Agent Problem Impacting on Performance. *Management Research Review, 38*(12), 1306 -1330.

Lehn, K. (1982). Property rights, risk sharing, and player disability in Major League Baseball. *Journal of Law and Economics, 25*(2), 343-366.

Lehn, K. (1984). Information Asymmetries in Baseball's Free Agent Market. *Economic Inquiry, 22*(1), 37-44.

Lenard, M. A., Peña, P. A. (2017). Maturity and minorities: the impact of redshirting on achievement gaps. *Education Economics, 26*(6), 593-609.

Linden, A., Samuels, S. J. (2013). Using balance statistics to determine the optimal number of controls in matching studies. *Journal of Evaluation in Clinical Practice, 19*(5), 968-975.

Kahn, L. M. (2000). The Sports Business as a Labor Market Laboratory. *The Journal of Economic Perspectives, 14*(3), 75-94.

Krautmann, A. C. (1990). Shirking or Stochastic Productivity in Major League Baseball? *Southern Economic Journal, 56*(4), 961-968.

Krautman, A. C. (1999). What's Wrong with Scully-Estimates of a Player's Marginal Revenue Product. *Economic Inquiry, 37*(2), 369-381.

Krautmann, A. C. Donley, T. D. (2009). Shirking in Major League Baseball - Revisited. *Journal of Sports Economics, 10*(3), 292-304.

Marburger, D. R. (2003) Does the Assignment of Property Rights Encourage or Discourage Shirking? *Journal of Sports Economics, 4*(1), 19-34.

Maxcy, J. G. (1997). Do long-term contracts influence performance in Major League Baseball? In W. Hendricks (Ed.), *Advances in the economics of sport* (pp. 157-176). Greenwich, CT: JAI.

Maxcy, J. G., Fort, R. D., Krautmann, A. C. (2002). The Effectiveness of Incentive Mechanisms in Major League Baseball. *Journal of Sports Economics, 3*(3), 246- 255.

Maxcy, J. G. (2018). Multi-period contracts as risk management in professional sports. In Longley, N. (Eds.), *Personnel Economics of Sports* (pp. 88-110). Northampton, MA: Edward Elgar Publishing.

National Hockey League (2017). 2017-2018 Official Rules of the NHL. Official Rules of the NHL. Chicago, Illinois: Triumph Books.

O'Neill, H. M. (2014). Do hitters boost their performance during their contract years? *Business and Economics Faculty Publications, 7*, 78-85.

Oyer, P. (1998). Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality. *Quarterly Journal of Economics, 113*(1), 149-185.

Purcell, E. A. (2009). Long-Term Contracts and the Principal-Agent Problem. *Gettysburg Economic Review*, *3*(5), 45-66.

Rosen, S. (1981). The economics of superstars. *American Economic Review, 71*(5), 845-858.

Rosenbaum, P. R., Rubin, D. B. (1985). Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician, 39*(1), 33-38.

Scoggins, J. F. (1993). Shirking or Stochastic Productivity in Major League Baseball: Comment. *Southern Economic Journal, 60*(1), 239-240.

Scully, G. W. (1974). Pay and performance in Major League Baseball. *American Economic Review, 64*(6), 915-930.

Solow, J, von Allmen, P. (2016). Performance expectations, contracts and job security. In Barry, D., Skinner, J., Engelberg, T. (Eds.), *Research handbook in employment relations in sport* (pp. 46-68). Northampton, MA: Edward Elgar Publishing.

Sommers, P. M. (1993). The influence of salary arbitration on player performance. *Social Science Quarterly*, 439-443.

Staudohar, P. D. (2013). The hockey lockout of 2012–2013. *Monthly Labor Review*, U.S. Bureau of Labor Statistics, July 2013, https://doi.org/10.21916/mlr.2013.24

Stiroh, K. J. (2007). Playing for Keeps: Pay and Performance in the NBA. *Economic Inquiry, 45*(1), 145-161.

Woolway, M. D. (1997). Using an empirically estimated production function for Major League Baseball to examine worker disincentives associated with multi-year contracts. *American Economist, 41*(2), 77-83.

Appendix A

Proof of Lemma 1.

The manager's problem can be written as:

$\max_{\{\beta,\hat{V}\}} \mathbb{E}[\pi] = e - \frac{1+D(e-\hat{V})}{2}\beta$

subject to

$\hat{e} = \frac{\beta D^2}{2}$ (IC)

Using (IC) in the objective function and simplifying:

$\max_{\{\beta,\hat{V}\}} \mathbb{E}[\pi] = \frac{\beta D^2}{2} - \frac{\beta}{2} - \frac{\beta^2 D^3}{4} + \frac{D\hat{V}\beta}{2}$

Since, the objective function is linear in $\hat{V}$, the solution requires that $\hat{V}$ is high as possible: $\hat{V} = 1$. Differentiating w.r.t. β:

$$\frac{d\mathbb{E}[\pi]}{d\beta} = \frac{D^2}{2} - \frac{1}{2} - \frac{\beta D^3}{2} + \frac{D\hat{V}}{2}$$

Hence, it is straightforward to see that the maximizer is:

$$\beta = \frac{D^2 + D\hat{V} - 1}{D^3}$$

Proof of Corollary 1

Differentiating the optimal bonus as from Lemma 1:

$\frac{d\beta}{dD} = \frac{-D^2-2D+3}{D^4} < 0$

where the inequality follows because $D > 1$.

Appendix B

Table B.1 Descriptive statistics on contract length, by season.

Season	N	Mean	Standard deviation
2008	619	2.868	0.966
2009	663	3.086	1.213
2010	644	3.250	1.445
2011	633	3.313	1.641
2012	591	3.425	1.780
2013	629	3.480	1.888
2014	662	3.530	1.871
2015	654	3.555	1.878
2016	597	3.709	1.954
2017	612	3.770	2.017

Balance tests

We verify whether *Team_length* is orthogonal with respect to skaters' observable and time-variant characteristics; we do that by using two sets of balance tests. For these tests, the information on *Team_length* is compressed into a dummy variable that we call *Med_iv;* this variable equals 1 for skaters whose value of *Team_length* is equal or lower than its median value (i.e., 3.257 seasons) and 0 otherwise.

For the first set of tests, the sample is divided in two groups, based on the *Med_iv* value. Then, we compare the average value of the observable characteristics across these two groups and assess possible imbalances with three tests. First, we conduct t-tests to assess statistical significance of possible differences across the two groups, and then we check (i) standardized differences and (ii) variances ratios to assess economic significance as well. Based on previous literature (Linden & Samuels, 2013; Rosenbaum & Rubin, 1985), an imbalance is not economically significant if the value of (i) is 0.2 or lower,[18] and if the value of (ii) it is about 1.

Table B.2 reports the results. In order to limit the size of the table, we report only the average values of (i) and (ii) for season and team fixed-effects.

Table B.2 Unconditional balance t-tests on observable characteristics, standardized mean differences, variances ratio.

	Med_iv = 0		Med_iv = 1		T-test, absolute mean difference	Standardized mean difference	Variances ratio
	Mean	Variance	Mean	Variance			
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Age	27.311	22.095	27.305	20.553	0.006	0.050	0.762
Seasons						0.001	1.002
Team						-0.010	1.274

Note: Column (6) and (7) report the average values of the standardized mean difference and the variances ratio for season and team fixed effect. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Results in Table B.2 reassure us on the validity of *Team_length*. The t-test on the absolute mean difference in *Age* between the two groups of skaters, in column (5), is not statistically significant. Moreover, all of the standardized mean differences as well as variance ratios—for *Age, Seasons*, and *Team*, give reassuring values, which are close to the reference ones.

Additionally, we conduct two parametric conditional balance tests where we run two OLSs—with individual fixed effects—with *Med_iv* and the original instrumental variable

[18] According to Caliendo and Kopeining (2008), an imbalance can be neglected if the value of (i) is even smaller, between 0.03 and 0.05. Results for Column (6) are reassuring even if we refer to these stricter thresholds.

Team_length as outcome variables being regressed on all of the control variables. Table B.3 reports the results.

Table B.3 Conditional balance test; fixed effects estimates of observable characteristics on *Med_iv* and *Team_length*.

Variables	Med_iv	Team_length
	(1)	(2)
Age	0.014	-0.016
	(0.020)	(0.020)
Age square	-0.094	0.102
	(0.094)	(0.098)
Season[a]	[0.152] [a]	[0.055] [a]
Team[a]	[0.375] [a]	[0.537] [a]
Observations	6,279	6,279

Note: Standard errors clustered on skaters in parenthesis.
[a] Average p-value is reported in square brackets.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Overall, the results from the unconditional balance tests in Table B.2 are confirmed.

Appendix C

Table C.1 *Remaining* and (*Remaining* × *First_season*) first stages.

Variables	Remaining (1)	Remaining × First_season (2)
Team_length × First_season	-0.102 (0.162)	-0.381*** (0.064)
Team_length	0.421*** (0.008)	0.667*** (0.014)
Observations	5,895	5,895

Note: All the analyses include age and its square, as well as team and season fixed effects. Standard errors clustered on skaters in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table C.2. *Team_length* and (*Team_length* × *First_season*) reduced forms.

Variables	Points per game (1)	Plus-minus per game (2)	Penalty per game (3)
Team_length × First_season	-0.001 (0.009)	-0.001 (0.015)	-0.004 (0.027)
Team_length	0.009 (0.009)	-0.001 (0.002)	-0.001 (0.003)
Observations	5,895	5,895	5,895

Note: All the analyses include age and its square, as well as team and season fixed effects. Standard errors clustered on skaters in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.