2018

Shirking and Remaining Years on Players' Contracts in Major League Baseball

Henry Tizard
Skidmore College, htizard@skidmore.edu

Shirking and Remaining Years on Players' Contracts in Major League Baseball

Student Researcher: Henry Tizard

Faculty Sponsor: Professor Monica Das

May 1st, 2018

This thesis was submitted in partial fulfillment of the requirements for the course Senior Seminar (EC375), during the Spring semester of 2018

Abstract

This undergraduate economics thesis is meant to find statistically significant evidence for shirking behavior in Major League Baseball (MLB). Theory suggests that players shirk on effort when they have recently signed a long-term lucrative contract, since there is little incentive to compete when money is guaranteed to the player regardless of current performance. It is particularly important to understand the MLB labor market, since the firms giving out contracts not only have a copious amount of production information regarding their employees, but this data is widely available to the general public. This study will make use of modern sabermetric statistics in order to further open up the conversation regarding shirking with advanced statistics. It will also seek to control for other motivational forces at play, such as intrinsic motivation of the player based on their own self-confidence, as well as extrinsic motivation regarding the performance of the time in terms of win-loss record. Such motivational factors had not yet been discussed within the shirking conversation.

Introduction

Do Major League Baseball (MLB) players exhibit less effort when they have many years left on their contract? The MLB labor market is set up in such a way that after signing a long-term contract, players really have no reason to play except to earn their next contract. Does this mean that players will only exhibit maximal effort as their contract nears expiration? Research regarding shirking behavior informs economists regarding human behavior. Evidence of shirking would imply that intrinsic motivation is not strong, and that the worker is entirely incentivized by money.

Other articles regarding shirking behavior, such as those by renowned sports economist Anthony Krautmann, set out on numerous occasions to see whether or not shirking behavior can truly be empirically proven. Unfortunately, different models seem to find differing results, depending on the theoretical framework behind the model, the performance measures used, and the type of contracts included within the sample.

This study will seek to be the first to control for intrinsic motivation in the context of testing for shirking behavior in Major League Baseball. Effort is a direct result of motivation, so this article contains multiple regression models that include other motivating or demotivating controls in order to isolate the effect of the years remaining on a player's contract and that player's performance.

The model that this thesis presents brings experimentation and innovation to the literature. This article also is the first to use the sabermetric statistic Wins Above Replacement (WAR) as the performance variable. This makes this thesis unique in that no other paper uses a performance measure that can compare pitchers and hitters.

Unfortunately, the models created in this study had differing results, and shirking behavior based on contract years remaining remains empirically unproven. Differing performance measures yielded different results, which was consistent with other articles in the literature. More studies need to be conducted in order to understand the WAR statistic better, as well as to establish regular measures of forces of motivation competing with the incentive-based motivation, such as extrinsic and intrinsic.

The next section will provide a comprehensive literature review of shirking behavior. Then, the next section will go over the framework for the first model. The third section will analyze the model. The fourth section will introduce a similar model with very slight adjustments, then compare the models. The fifth section will do this once more, slightly adjusting the model and discussing differences. The final section is a brief conclusion to sum up the thoughts, opinions, and lessons that the thesis provides.

Literature Review

Sports are unique. They are among the only markets in which it is straightforward to quantify an individual's productivity. In addition, labor contracts are structured in such a way where players often known almost exactly what they are going to be paid, regardless of current performance. Rather, contracts are negotiated based on past performance, and teams must commit guaranteed money to players based on projections for the future. When it is put so simply, it seems obvious that teams should commit more money and years to superior players. This is how the market has functioned. However, this creates a classic principle-agent dilemma, in which players function as the agents, and team owners function as principles (Pedicelli, 2015). Once owners or executives elect to sign a player to a long-term contract, that player no longer has an incentive to exhibit maximum effort (assuming that exerting effort is unpleasant). Take the case of Pablo Sandoval, for instance. Sandoval was an accomplished player and a playoff hero for the San Francisco Giants, but after signing a 5-year, $90 million contract with the Boston Red Sox after the 2014 season, he was consistently overweight, out of shape, and an underperformer for the Red Sox. He was released by the Red Sox after two and a half disappointing seasons, but he is still being paid by the team. This extreme example exemplifies the importance of evaluating contract structures in the Major League Baseball (MLB) labor market. Could such a circumstance been avoided? Is there a pattern of reduced effort after the signing of a long-term contract? Sports economists have studied this idea of "shirking", or making the decision to exhibit less effort under certain circumstances. This paper will extensively analyze and critique the existing literature on this subject to provide context for my own model for analyzing the controversial relationship between effort and incentives in professional sports.

A fundamental problem that arises in principle-agent relationships, outlined by Holmstrem (1979), is the imperfect information that principles have to deal with when designing contracts for agents. As a result, employers are expected to invest capital into monitoring the actions of their employees in order to get a stronger sense of individual production and to ultimately use this information to create a labor contract (Holmstrem, 1979). Sports, specifically Major League Baseball due to their advanced statistical measures of performance, are the finest example of this performance monitoring, analysis, and subsequent action based on this information. Investing as much capital as is efficient into performance monitoring is described as a second-best solution to perfect information (Holmsrem, 1979). This behavior can easily be observed in Major League Baseball, as advances in analytics are at the forefront of public attention, largely beginning by the "Moneyball" approach most famously used by Billy Beane and the Oakland Athletics in the late 1990's, which has since been used across the league. Although complete monitoring has not yet been achieved, Major League Baseball owners and executives continually strive towards this ideal. This makes the MLB and other sports labor markets unique. Unlike most firms, who are unable to invest capital heavily into monitoring, since it is so prohibitively expensive, MLB franchises have copious amounts of money to monitor and scout players in order to optimize labor contracts. When it comes to monitoring players, all of whom make six, seven, eight, or even nine figures, investing into this function is not only important, it is absolutely necessary to the success of a franchise. Holmstrem (1979) would anticipate that near perfect monitoring would result in contracts that actually penalize dysfunctional behavior (Holmstrem, 1979). However, incentive based contracts are only starting to become popular recently, and contracts that penalize poor performance do not exist. This could be due to the power of the MLB Players Association. The unique amount of information in baseball specifically makes it an important market to understand. As technology, as well as statistical power to predict and value continue to improve, monitoring similar to that in baseball will likely become to the norm in other labor markets as well, so the evolution of labor contracts in baseball could serve as a framework for future contracts in other labor sectors.

In order to set the context for discussing this market inefficiency, it is important to understand where the MLB labor market is in its evolution. Before free agency, players had no choice but to sign contracts under the reserve clause system. This reserve clause was in every contract, and it allowed the owner of the team to renew a player's contract after the previous contract expires indefinitely (Knowles et al. 2003).

The advent of free agency brought about a shift in the balance of risk sharing. As opposed to the team having the right to continue to renew the existing contract of the player, free agency makes it so that teams usually do not have this opportunity to opt out of future payments, unless of course this provision exists within the contract. Essentially, the player is provided with assurance that he is going to paid, and thus it can be inferred that this player will demonstrate moral hazard.

Obviously, due to the long-term contract situation stated above, there is an inefficiency inherent in the labor market. Luckily for MLB teams, players like to play, and there is strong intrinsic motivation within players. If there was not, then Giancarlo Stanton could have stopped exhibiting effort purposely after signing his contract worth over three hundred million dollars, and the Marlins would have had absolutely no option but to pay the man lazily sat on the bench, producing nothing for the franchise, because he simply does not have to. This type of contract structure would fail outright in just about any other type of business. Therefore, there must be other forces, such as intrinsic motivation that make this inefficiency difficult to prove econometrically. These forces must be controlled for in order to isolate the inefficiency.

Anthony Krautmann, who writes extensively on the subject of shirking in Major League Baseball, and John L. Solow created a study that measures changes in performance over the course of long-term contracts (Krautmann and Solow, 2009). In the introduction, they state that there are many reasons to be skeptical about whether athletes shirk (Krautmann and Solow, 2009). This is undoubtedly true. The main argument he brings up is, however, not an argument. They *ask* why owners continue to offer long-term contracts if there are perverse incentives. "Theory suggests that incentive-compatible contracts should evolve to punish opportunistic behavior" (Krautmann and Solow, 2009). In this statement, Krautmann and Solow fail to acknowledge that we are not indeed in the long run. Free agency has only existed for fifty years in the MLB, and in that time, contracts have changed significantly, but the information-gathering phase is far from complete. Theories often assert that there is a long-run equilibrium that we trend towards, but in this case, the long-run equilibrium or trend, perfect information, is impossible to achieve. To ever assume that the market has already achieved long-run equilibrium status is a lofty assumption, and in this case, it is actually impossible, since no one knows what a player will contribute at the time that the team official and the player agree to the terms of the contract, given the current contract structure in which salaries are determined before performance. Perhaps contracts will evolve to punish opportunistic behavior more effectively, but the mere fact that they have not yet is not an argument that disputes the phenomenon.

This faulty assumption is important for the actual model itself in that the model does not control for other motivational factors that guide the athlete to not shirk (or to shirk). As stated previously, the contract structure undoubtedly does, theoretically, leave athletes in a situation in which they do not need to exert unpleasant effort to make large sums of money. This idea needs to be isolated in order to be proven empirically. Other factors, such as intrinsic motivational factors, such as gaining fan support, winning a world series, earning the respect of teammates and the media, etc. are not considered in the model and thus cause it to be biased.

Faulty assumptions aside, it is time to analyze the model for what it is. The model is centralized around the idea that there is evidence of shirking when actual performance deviates significantly below expected performance. The model tests for changes in expected adjusted OPS and actual adjusted OPS (adjusted OPS is simply a common sabermetric statistic used to evaluate the production of hitters) over the course of long-term contracts. Note that this does not include pitchers. The model executes this test in multiple steps. They first estimate the relationship between performance and experience in order to create expectations for how a player will perform when they exert maximum effort. They only consider players in their walk year in order to avoid a bias, as they assume that players not in their walk year will not exert maximal effort due to their hypothesis (Krautmann and Solow, 2009). They also necessarily used a vector of various controls. Next, they calculated the players' shirk value as the difference between expected and actual performance (Krautmann and Solow, 2009). Moving right along, they then estimate the probability that the current contract is the player's last (Krautmann and Solow, 2009). This contrasts with other models that only look at the ex post result of retirement. This idea is better because the player has not necessarily made the decision to retire in the middle of the season, but they are certainly aware of the likelihood that they make this decision in the future. This is one of the particularly brilliant aspects of the paper. Lastly, they estimate the relationship between the shirk variable and the amount of years remaining on the player's contract to ultimately acquire their results (Krautmann and Solow, 2009).

The results assert that players who have a high probability of retiring at the conclusion of the contract exhibit less effort (Krautmann and Solow, 2009). They also found that players in their walk years play better than their projections (Krautmann and Solow, 2009). They also found that the coefficient of the years remaining variable was positive and significant, showing that players do in fact exert more effort when they have less time left on their contracts. Overall, this is a highly complex, successful model, one that sets the standard for all over models based on shirking behavior in Major League Baseball.

Another model involving Krautmann, this time in tandem with Donley, looked at shirking through two different definitions of productivity. One of them was of course marginal physical product, focusing on what the player is able to provide in terms of production on the field (Krautmann and Donley, 2009). The other model within the study was centered on marginal revenue product. The one using marginal physical product also used OPS, but the projections were formed using production from the prior three years, which obviously differs from the projection process in the model already discussed. This model yielded results that were not statistically significant (Krautmann and Donley, 2009). Unfortunately, the authors provide no possible explanations as to why this may be, especially since Krautmann was involved in the paper discussed above that did find statistically significant results. Unfortunately, he offers no comparisons. The performance measure differs between the two studies very slightly, as the former study utilizes a more advanced sabermetric statistic called adjusted OPS, whereas this one uses the normal OPS statistic that has been around for over a hundred years. Perhaps a reason for this is the lack of a variable controlling for the player's retirement probability. A potential problem with this model is that it uses a simple control for age, as a linear relationship between age and performance is not likely to exist in linear form, as there is more likely to have a parabolic relationship between these variables. Therefore, functional form may be a problem throwing the results off. Another potential problem with the model is that the three years prior to the signing of a new contract may be biased downwards, as those years immediately preceding a contract carry high incentive to perform to maximum potential. This potential problem was discussed in the previous paper, but not this one. It is also not the cause of the insignificant coefficient values, since one would expect that bias to be in the direction that would favor more significant results

The model using marginal revenue product as the value standard brings an interesting twist to the literature. Marginal revenue product refers to the amount of revenue that a player generates. Expected marginal revenue product is logically assumed to be the negotiated salary for the player due to free-market bidding, while actual marginal revenue product requires calculation. They did so by estimating the team's run differential (runs scored minus runs allowed), then using this to estimate a quadratic revenue function. Using this approach, the authors were able to find significant results (Krautmann, 1990). The author offered numerous explanations as to why this may be the case, such as synergies between players and teams that consistently overpay. More research needs to be done in order to validate these explanations. In my opinion, the expected marginal revenue product being equal to the salary is true, but assuming that these expectations being realistic and accurate is difficult. These expectations should logically shift as more data is acquired, but contract structures do not allow for these shifts to be reflected in the model. In actuality, outlooks on players can shift within a year or two, so in the context of long-term contracts, which often exceed four years, these expectations are incredibly dated. Also, this assumption is faulty in that teams tend to structure contracts so that the salary is "back loaded". There are numerous reasons that teams back load contracts, such as acquiring the player in their prime for less than they are worth, allowing the team to have more current money to spend on other players, as well as assuming higher future revenues and lower luxury tax limits, etc. The fact of the matter is that team's pay players more than they expect them to be worth at the end of the contract, so this could inflate the shirk variable for the end of contracts that are back loaded.

Another article brings up the randomness of production in professional sports. Krautmann offers an explanation dispelling the lack of incentive hypothesis centered around the randomness of productivity in professional baseball. It is clear that there is some degree of randomness and luck involved (Krautmann, 1990), but this explanation does not sit particularly well. This is especially true in a sport like baseball, where the sample size is normally around 162 games (a full season) and around 600 plate appearances, a large sample size as opposed to say football, which has only 16 game seasons.

In the conclusion, Krautmann cites the complexity of contracts as a reason that the disincentive problem does not exist without providing a theoretical basis for this claim (Krautmann, 1990). What complexity is he referring to? He fails to provide any sort of examples of how certain provisions may cancel out the disincentive problem. He may be referring to team or player options to opt out of the contract, but one could easily leave these contracts out of the sample. He may also be referring to incentive bonuses that are becoming more and more prevalent in contracts. However, these incentives, in almost all cases, only represent a relatively small percentage of the overall value of the contract. Again, these are easy to keep out of the sample.

A response by Scoggins points out a very valid and distinct problem with this Krautmann article. Scoggins decided to run the exact same regression, except altering the performance measure ever so slightly, from slugging average to total bases (Scoggins, 1993). Slugging average is simply total bases per at-bat, so I must stress how small of a change that this is to the original regression model (Scoggins, 1993). This miniscule change allowed Scoggins to reject the null hypothesis that shirking does not occur. This stresses the importance of testing using multiple different performance measures in analysis.

It is frustrating that Krautmann does extensive research in multiple different ideas that are prevalent in sports economics issues without combining them into a model that encompasses multiple factors. He simply does his best to isolate certain effects, and he does not connect his articles in order to gain any conclude-able or useful findings.

Another article, this one by Maxcy et al. (2002) postulates that both ex ante strategic behavior and ex post opportunistic shirking behavior both take place in Major League Baseball. Therefore, they run a regression testing the performance of players who are in close proximity to a new contract, either about to signed, or just having been signed. They use a control for the average performance of the player three years prior as a proxy for expected performance, then compare that to the performance either just before or just after a contract was signed (Maxcy et al., 2002). The performance metric used was slugging average for hitters and strikeout to walk ratio for pitchers. Obviously, these differing measurements mean that the pitchers and hitters must be compared separately. They also use a dummy variable for both the first and last year in the contract, which I found to be very interesting, since it allows them to isolate the effects of the two most extreme years for both opportunistic behavior strategic maximal effort behavior. The study found evidence of ex ante opportunistic behavior, by finding that time spent on the Disabled List and playing time is significantly higher immediately preceding contract negotiations (Maxcy et al., 2002). They were unable to find any evidence that performance was higher before a new contract, or lower after a new contract for either pitchers or hitters.

An economist by the name of Katie Stankiewicz (2009) takes on a different approach. Stankiewicz (2009) decided to compare the production players under multiyear contracts and one-year contracts in order to show that players on one-year contracts are more productive than those with multiple year contracts (Stankewicz, 2009). This is problematic in that players with long-term contracts are inherently better players, as teams will only be willing to commit multiple years to a player if that player is a well above-average player. Her model is very simply and only has controls for age, games played, and coach's success, the latter of which is largely inconsequential in the world of baseball. She made no attempt to control for the talent of the players involved. Unsurprisingly, she found that players with long-term contracts out-performed those on one-year contracts, meaning that the sign of her one-year contract was negative instead of the positive sign that she expected (Stankewicz, 2009).

Stankewicz also expanded upon this paper by looking at average productivity of players with four-year contracts in each individual year of the contract. This means that she averaged all sample players' performance in the first year of the contract, the second year, etc., even including the year right before the contract was signed (Stankewicz, 2009). She concluded that since average performance did not increase in year four, then shirking could not be shown to have taken place. This is problematic for many reasons, the largest of which being that she did not use any controls, not even for age. Since players can only enter free agency once they have six years of major league experience, players are often on the second half of their career and declining by the time these long-term contracts are signed, so age would be expected to have a negative sign.

One thing that Stankewicz did do that was interesting is that she used a sabermetric statistic EqA as her performance variable, which is shorthand for equivalent average per out. This statistic adjusts performance for variables such as home ballpark of the player and level of pitching that the player faced against. The statistic also includes base-running value, which is often ignored in the literature, for example by OPS, adjusted OPS, and slugging average.

The glaring, overarching theme in the literature regarding shirking behavior in Major League Baseball is that economists disagree about whether it is provable. Some believe that the evidence of shirking is sufficient to prove the relationship between contract length and effort exists in the way that the shirking hypothesis expects, while others are unconvinced. The models that have been created are often very susceptible to minor changes, as these minor changes can drastically alter the conclusions that the model makes. These changes can be either controls or statistical measures used to aggregate performance.

Another problem with the existing literature is that all of the studies look exclusively at hitters, while almost none consider pitchers. Pitchers play considerably less games than hitters do, but it is incredibly difficult for pitchers to keep themselves sufficiently healthy and strong to compete at the Major League level. This is due to the repeated torque that a starting pitcher puts on his arm over the course of a season. This is not to say that it is easier for hitters to sustain themselves than it is for pitchers, but ignoring half of the game when it comes to shirking behavior seems to be an oversight.

They also rarely account for other incentive forces that are compelling players to not shirk. The more sophisticated models have been able to find evidence of shirking, at least in certain performance measures. Surprisingly, all of these articles, despite having been written in the twenty-first century, used statistical measures for performance that would be considered archaic strategies for aggregating production. There are numerous sabermetric statistics that are far more generally accepted nowadays in their goal of aggregating player value than slugging average or OPS. The only measures that were at all modern were the marginal revenue product estimate in Krautmann and Donley (2009), and the equivalent average technique used in the Stankewicz (2009) article. None used wins above replacement (which is similar to the marginal revenue product estimate), which was shocking due to its prevalence in sports media coverage of players. It is widely cited by MVP voters as a driving reason behind voters' decision making.

Also, there is no mention as to how the models account for various types of options that are often in MLB contracts. These options include player, team, mutual, and vesting options. Player options are also called opt-outs, so player's can sign contracts in which before a certain year in the contract, they can choose to not exercise the rest of the contract and become free agents. Team options are the same except the team can nullify the contract. In mutual options, both parties must agree to the terms or else the player becomes a free agent. In vesting option, the player will become a free agent unless they have reached a certain threshold; these options are by far the least common of the four.

Since the model about to be discussed also attempts to control for player confidence, it is important to outline past research that has been conducted regarding confidence. Feltz (1988) outlined the major frameworks that sports economists have used to study confidence.

The first theoretical approach to studying self-confidence in sports is through self-efficacy. Bandura's theory of self-efficacy refers to the judgments of what an individual can do with the skills he or she possesses, and this directly impacts the effort and conviction one has to successfully executive a behavior required to produce a desired outcome (Feltz, 1988). This is the exact type of effort that I want to be able to control for in my model. Expectation of person efficacy come from four sources, personal accomplishments, vicarious experiences, verbal persuasion, and physical arousal, and performance accomplishments are the most dependable source of information to determine an individuals self-efficacy (Feltz, 1988). However, it should be mentioned that this brings validity to Stankewicz' (2009) control of coach's success in her model, despite it not being the most effective way to determine an individual player's confidence.

It is also important to look at other sports, when discussing shirking behavior in professional sports. The National Football League has a slightly different labor market, as the contracts are much more incentive based than in Major League Baseball. Marks (2017) created a model in which he looked at all NFL players in the same sample, including different positions with dummy variables to characterize them. He found that players with new, long-term contracts performed worse on a sabermetric point scale created by Pro Football Focus, but not to a statistically significant degree. This model sets the precedent that it is highly difficult to compare players amongst different positions in the same sport. He did not run any regressions only using players of a certain position, giving himself a larger sample than just the 2015 season. Like most models looking at shirking behavior, Marks (2017) used age and age-squared as variables of controls. He also used the performance in the previous year as a basis for comparison for his model. He compared players' performance to their performance the previous year, citing the difference between these two values as his dependent variable. Like other models, Marks (2017) focuses on whether or not the player had just signed a new, long-term contract.

Shirking is also extensively studied in European soccer. The journal article, Feess et al. (2010) assesses the effectiveness of long-term contracts as incentives, much like all of the other relevant articles, except this one looks at the Bundesliga, the top tier soccer league in Germany. Much like my own data, this article collected its information by hand, from a magazine in this case. Also like my own models, Feess et al. (2010) are aware of their own selection bias in selecting their sample. While their issues stem from selecting only long-term contracts, mine stems from only considering free-agent contracts and ignoring contracts of players that are still under the reserve system. They also used a number of interesting controls, such as whether the player was on the national team, age, number of games, whether the previous contract was "renewed" (which cannot happen to players in my samples), or whether the previous contract was allowed to expire. The idea behind this is that a player who has had their contract renewed has earned the favor of ownership for some reason, whereas allowing a contract to expire can signal unfulfilled expectations or low talent (Feess, et al., 2010). Feess, et al. (2010) were actually able to find highly significant results using this framework, showing strong evidence of shirking associated with people who had signed long-term contracts.

Lastly, shirking is also studied in the National Basketbal Association (NBA). Jean (2010) looks at the long-term contract and ex post shirking behavior and ex ante opportunistic behavior, very similar to Maxcy et al. (2002) journal on baseball.

This just goes to show how widespread and well studied this framework is. Like Maxcy et al. (2002), Jean (2010) was unable to find significant results of shirking behavior using this framework. Jean (2010) however, does raise an important point that I am trying to make in his conclusion, that there may be a variety of confounding factors that prevent these shirking behaviors to be detectable. My study will attempt to control for some of these confounding factors within the framework.

The First Model

For the model, the goal was to directly show a correlation between wins above replacement (WAR) for a player and the years left on that player's contract during that season. In order to find this relationship, the sample and the controls for the model are of course extremely important. This next section will detail how the sample and controls were chosen and what they were intended to do in the model.

For the sample, the goal was to be as inclusive as possible. The players selected were those on free agent deals, and the seasons that were considered were the 2016 and 2017 seasons. The sample was fairly large as compared to the samples used in the existing literature. The information was gathered using information from spotrac, baseball reference, and fangraphs. I chose to leave out catchers and relief pitchers, as these players play significantly less than other position players and starting pitchers, even when healthy. They likely could have been included without significantly impacting results, but I chose to leave them out due to the fact that they do play and contribute less. Leaving these players out posed no risk for the results of the paper, but I feared leaving them in the model would cause distribution of the WAR statistic to be too heavily concentrated around zero, leading to strong heteroskedasticity. That said, I was able to use nearly all of the players on veteran contracts for the model that were neither catchers nor relief pitchers. There are a few missing due to difficulty in acquiring certain contract data, which necessitated that I pay for the data. Also, I left out players whose contracts were comprised almost entirely or entirely of option years, such as the case of David Price's contract with the Red Sox and Clayton Kershaw's contract with the Dodgers. I ended up with

188 observations in total, which is more than sufficient based on the sample sizes of the existing literature.

Also, instead of using a SHIRK variable that represents the disparity between actual performance and expected performance, this model simply looks at the relationship between the years remaining on the contract and a performance measure. This is where it becomes important to control for motivational factors and talent, which is the biggest challenge taken on by this model, and these controls will be discussed in detail in a moment. The goal was to keep the framework as simple and easy to analyze as possible, while also planning on acquiring robust results.

Here is what my initial regression equation came out to:

$$WAR_i=\beta_{0i}+\beta_{1i}ADJGAMES+\beta_{2i}YRSLEFT+\beta_{3i}LOGSALARY+\beta_{4i}LOGTMONEY+$$
$$\beta_{5i}PREVREC+\beta_{6i}RECORD+\beta_{6i}LASTYEAR+\beta_{7i}AGE+\beta_{8i}AGE^2+\beta_{9i}IMPROVE$$

The performance measure used in the initial model is different from those in other models. I chose *WAR* (using the fangraphs definition, specifically) as the performance measure for multiple reasons. *WAR* estimates the number of wins that the player contributed over the course of a season compared to a replacement level player for that season. First, *WAR* takes into account performance of all aspects of the game, including base-running, fielding, and hitting for positions players, and pitching and fielding for pitchers. It even adjusts what performance level characterizes a replacement level player year by year, accounting for the evolution of the average replacement level player as the years go on. Since this study only uses 2016 and 2017 statistics, it is unlikely that this benefit is going to make any difference. Nonetheless, it provides a very complete aggregation of a player's performance. Other studies often use measure like batting average, slugging percentage, or on-base percentage plus slugging percentage. These measures only look at hitting, ignoring fielding and base running, which are two very important components of a player's contribution to the game. This is the only piece in the literature that uses *WAR* as its performance measure. *WAR* also controls for other factors, such as the home ballpark factor, which is important since the dimensions and altitude of the ballpark contribute to the player's production and players play half of their games at their home ballpark. EqA, used by Stankewicz (2009), also takes into account a lot of these same variables, so this article uses a similarly modern performance statistic. Her measure did not include defense, however, only offense. The ballpark controls are important because a ball in Denver, Colorado simply goes farther than it will in other ballparks due to thinner air at extreme altitude. Therefore, players on the Colorado Rockies will have statistics that are biased upwards since half of their games were in Coors Field. *WAR* takes a hitter-friendly or pitcher-friendly environment into account on both the hitter side of things and the pitcher. There are endless examples like this one that allow for more advanced, comprehensive player comparison using *WAR*. *WAR* also allows us to include pitchers in the model, which is an important differentiator of the model, since pitchers are essentially the other half of the game and are completely lacking from the literature on shirking behavior.

The other most important variable is the number of years remaining on the player's contract, *YEARSLEFT*. These were collected on spotrac.com, which is a good resource for acquiring contract information in many sports. Similar variables were used in all of the models in the literature review in one step or another, although this regression is done in one step. The years left variable is very simply the total number of years in the contract minus the number of completed seasons on the contract, not including player option seasons, team option seasons, mutual option, and vesting option seasons. Other models did not say how they treated these seasons, but I chose to omit them entirely. These options may present an issue in the model, as they present a whole host of different relationship dynamics between the player and team that may affect the player's effort decisions depending on the player's perception of the salary in the option. Meaning, if a player has the right to opt out of a contract but feels that they are being overpaid, then they may have similar behavior to a player that has many years left on the contract, but a player with a player option in which they believe the salary to be too low, then they may behave as though they are in a contract year. It may have been conceivably possible to determine a probability that certain options would be exercised, and so these years could be considered in the model, but this model does not do so. Such an adjustment would be similar to the Krautmann and Donley (2009) retirement probability metric that they used instead of an ex post retirement variable.

The next variable is adjusted games played, *ADJGAMES*. The variable is there in order to control for players that missed games over the course of the season, which is usually due to injury or personal family emergencies. Adjusted games played is simply games played for position players. For starting pitchers, I multiplied games started by five, since starting pitchers pitch every five games for the most part. There are a few pitchers who got extra starts due to timely off-days in the team's schedule, but having a few adjusted games played numbers as higher than the 162 game season is unlikely to create any significant problems for the model. It is also possible that a hitter can play more innings than the amount of games they played times 9 due to extra innings, which are the baseball equivalent to overtime, so this more or less evens out due to this aspect of the game.

I also felt it necessary to control for the player's *SALARY* for the season. Since the player's in this study vary greatly in their overall production, there needs to be a way to control for how talented the player is. Since salary is determined by past production, we can assume that players with higher salaries have higher expected performance, so this variable attempts to put players on even ground. This may be problematic in that players with higher salaries have more incentive to shirk, as players who are not being paid as much of an astronomical sum may need to earn further contracts in the future in order to retire comfortably, since players likely have the goal to live out the rest of their lives without again entering the workforce, and professional athletes retire at a far earlier age than nearly all other professions.

Also, the initial model has a dummy variable that equals one if the player retires after the season, *LASTYEAR*. These players do not have the incentive to earn their next contract, since they are not going to be playing the next year. It would have been ideal to have a variable for retirement probability as in the model from Krautmann and Solow (2009) but for the purposes of this study, the ex post dummy variable is likely sufficient.

The model also controls for the team's record in both the season before and the season itself, *PREVREC* and *RECORD*. The record from the season before is used in order to adjust for the team's confidence going into the season. Presumably, teams that were good the previous year will feel confidence in their ability to improve and to ultimately compete for a title. Thus, a higher record will result in higher production. Record from the season itself works similarly, but it also accounts for how the team has changed since the previous year. Controls for extrinsic motivation are key here, extrinsic motivation is a key force opposing the lack of motivation associated with shirking.

Similar to other studies, this one also controls for *AGE* and AGE^2. As players get older their performance diminishes, so these variables control for these effects. AGE^2 is important since the relationship between age and performance is quadratic. As a player ages, they will improve for a time until they reach peak performance, also known as the player's prime, then they will regress after this point as they age and gradually lose their physical tools.

The regression also controls for talent with salary and total money committed to the player, *LOGSALARY* and *LOGTMONEY*. Players with more money committed and higher salaries are historically better player, so one would assume that higher salaries produce more. It is extremely important to control for talent because the study is not concerning fluctuations in performance over time, rather it tests for indicators of performance that are consistent between a wide variety of players. I took the log of salary and the log of total money committed so that the coefficients would represent the change in performance based on a percent increase in salary and money earned. This is more appropriate because the incremental value of a dollar is negligible and much more difficult to read as a miniscule coefficient.

Since Stankewicz (2009) found that players with longer term contracts perform better on the whole than players on one-year contracts and Krautmann and Oppenheiemr (2002) found that players with longer contracts generally have higher salaries, it can be inferred that higher salaries and higher total money can be associated with higher talent levels. Obviously, this is only true amongst players that have already hit free agency, since players that have not yet hit free agency are paid far less than they are worth if they have superstar level talent. This is not a problem in this study, since we only consider players veteran players that have been able to negotiate their contracts after establishing their talent.

The last variable is a variable that accounts for the player's individual confidence level and overall career trend called *IMPROVE*. The variable is a dummy that equals one if the player's *WAR* in the season preceding the observation is higher than their average *WAR*. This is a key variable in the model, since we are trying to control for other types of motivation for a player in order to isolate the effects on motivation and effort of the amount of years left on the contract. I have chosen to focus the confidence measure on self-efficacy, since it can be most easily measured by making this comparison between past performance and average performance. It is also easiest to measure, since it is largely based on past performance. Some may argue that this value will almost always be negative if a player got injured the previous year, and thus the player did not really underperform in that year. However, I think this circumstance makes sense in the framework, as players may not trust their bodies as much, particularly the part that was injured, in the season following an injury. This player may not lose confidence in their abilities, exactly, but they still may lose confidence in their bodies' ability to maintain when exerting maximal effort. This would cause them to be cautious about exerting maximum effort.

Analysis of the First Model

Despite these controls, the model did not find the relationship between years left on the contract and the performance of the player that was expected. In fact, the years left on the contract variable had a positive sign, which was significant at the one-percent level, which was not at all expected and is extremely perplexing. This could be due to a variety of factors to be discussed in this section.

Surprisingly, neither of the logs of salary nor total money committed were statistically significant. This suggests that this variable did not do what it was set out to do, which was to control for player talent. Obviously, a player with more talent is expected to outperform a player with inferior talent. The only things that would explain a player with more talent having lower productivity are that they played fewer games, yet there was not a lot of variation in this variable in the model, they were shirking, which there was no evidence of based on the model, on the contrary, there is evidence of the opposite, or if the player had begun to regress, yet we find that age and age-squared both had insignificant results. The total money coefficient even had a negative sign, suggesting that this relationship flat out does not exist, yet no conclusions can be drawn due to a lack of significant results. Other studies did not need to control for talent, but this study clearly does, as it does not create a shirk variable based on the difference between expected and realized production.

It also needs a talent control because there are so few restrictions on the sample size. Just because a player has no remaining years on his contract of course does not mean that they are going to outperform the very finest players in baseball who just so happen to have more years left on their contract. There needs to be an effective control that can account for this talent differential, so that we can observe the isolated effect of the amount of years remaining on the contract. Having more years left on the contract certainly would not increase a player's performance, but since the model fails to control for talent differentials between players on long-term and short-term contracts, we observe this positive sign. This leads me to believe that it would be prudent to scrap these variables and to find one that better controls for player talent. It is clear that salary and guaranteed money will not work, so something based on past production seems to be a better idea.

The finding that total guaranteed money and salary are not effective controls for talent is surprising. Maxcy (2004) determined that not only do superior players get longer contracts, but also bigger salaries. Therefore higher salaries and more guaranteed money would presumably be correlated with talent and performance, yet this is not found to be the case in this study. Krautmann and Oppenheimer (2002) can explain this phenomenon with their empirical conclusion that contract length and return on performance are inversely related. This suggests a trade off that occurs, in which players are willing to accept less money for their performance when negotiating a long-term deal (Krautmann and Oppenheimer, 2002). Meaning, if a player's performance is held constant, they will receive more money per year on a short-term deal than on a long-term deal. I may have created a bias in my own model here. If players with high salaries also have longer term contracts, and are thus more likely to have more years left on their contracts, then those players may be more likely to shirk as a whole, biasing the entire talent control downwards.

AGE and AGE^2 were also surprisingly insignificant. This could be due to a fairly small sample size. Many players in the sample were able to produce for a longer time than most players do. At the same time, there were many players who regressed in their early thirties. This is not to say that any specific player improved as they aged deep into his thirties, but rather that some players started to regress earlier than others. Since there is no set age where players will start to regress, we observe insignificant results. If the model used a data set with a thousand or more players, then the significance level would likely rise. Despite the insignificant p-value, the age variable should remain in the model, if for no other reason, because it is a pervasive variable in the literature on shirking behavior. It is possible that this variable is insignificant because pitchers and hitters begin to regress at different times. Pitchers may regress earlier on the whole than hitters, so hitters with higher ages and WAR values may be skewing the results. Also, it makes sense that the coefficients of the two *AGE* variables are negative, since all of the players in the study have all hit the free agency market already. This means that they have six years of experience in Major League Baseball already. Therefore, we would not expect them to continue to improve, since they have likely already hit their prime.

ADJGAMES had a coefficient that was highly significant and positive, so this variable likely did the job of controlling for the amount of time on the field during the course of the season. This is encouraging, since this is the only model to include both pitchers and hitters in the same model, and controlling for the amount of time on the field is pivotally important in models trying to explain performance with counting statistics.

The *IMPROVE* dummy variable proved to be significant and positive, as expected. This can be interpreted as the confidence from the strong previous year carried into the next season and caused helped that player continue upon their above-average performance. This suggests that the self-efficacy control that I created based on the analysis of Bandura's self-efficacy theory found in Feltz (2002) turned out to be effective, which is a major victory for the model. More research should be done regarding self-efficacy as a driver for productivity, especially so in order to find even better controls for intrinsic motivation to be used in studies such as this one.

I also conducted robustness checks in the model. Multicollinearity was a slight issue between *LOGSALARY* and *LOGTMONEY*, but neither vif score exceeded five, so the problem was not significant enough to fear a significant skewing of results. Consult Table 2 to see the vif results from STATA. Heteroskedasticity was a more significant issue, as the chi^2 value was significant on a one-percent level. Based on the other models that I ran for this thesis, I postulate that this heteroskedasticity had something to do with the performance measure, *WAR*. I am not entirely sure why this is, but both of my regressions using *WAR* had this issue, while the one I did using OPS did not have this issue whatsoever.

My study is far more general than the other studies. What I mean by this is that the study does not look at any specific circumstance, like at players who just signed a new contract or players who are about to hit free agency. Rather, the study simply looks at all players who have hit free agency and how their performance is effected by how many years they have remaining on their contract. For this reason, this study requires some creative controls that have not been tested before to my knowledge. Unfortunately, the existing published literature does not go into detail about what their vectors of controls entail, they just simply state that they used vectors of controls in the model.

I would feel more confident in asserting that years left remaining on the contract and production do not have an inverse relationship if there were far more data points, despite the fact that the p-value is significant at a one-percent level. Also, some of the variables have the unanticipated sign, suggesting that they fail to control for the phenomenon that they are intended to control for. The previous record variable, for example, actually has a negative sign. One would expect that a stronger previous record would have the effect of improving morale of the team, thus improving performance. However, there are arguments that would suggest that this sign makes perfect sense. Perhaps a strong record the previous year could provoke complacency in players, for example. This argument is very basic and not fleshed-out enough to really satisfy me, however. Also, the negative coefficient on *PREVREC* is not statistically significant, so this is likely just a statistical fluke.

All imperfections of the model aside, it is also entirely possible that the relationship between years left on the contract and a player's performance really is insignificant, since there is a lot of disagreement regarding previous results. Perhaps players simply shirk when they sign a lucrative new contract, as was suggested by Scoggins' (1993) response to Krautmann (1990), and therefore players with high salaries do not perform statistically significantly better than players who are not paid as much. Perhaps Krautmann and Solow (2009) found that players in their walk years outperformed those not in their walk years, but only because of a different phenomenon than consciously shirking. Perhaps players are able to push themselves into unsustainable production when they know that it counts the most, similar to a mother lifting a car off a child when she knows she has to. The mother cannot sustain the amazing strength, but it can be achieved for a relatively short period of time. Really, there is no strong evidence that a player begins to work harder as they near the end of their contract, which would need to be the case in order for my model to find statistically significant results.

While this model did not yield the results that were expected, I did learn a lot from this model based on what I did wrong. Also, the model contributes to the literature in that it introduces the idea of controlling for factors that are inherently difficult to control for, such as intrinsic and team motivation levels. Such work will hopefully spark others to search for more effective controls in these respects. Since I do not have the confidence to conclude that shirking does not exist, I took it upon myself to try to improve the model and to find better controls in order to find this relationship.

The Second Model

In order to improve upon the model, I made a minor adjustment to the aforementioned regression. Unfortunately, it is evident that the logs of salary and total money were unable to control for player talent. If they were, then we would be able to observe positive and significant results, which is not the case. In fact, we can observe in the first table of results that the log of total money actually had a negative coefficient. The adjustment was to replace these two variables with an average WAR variable. Average WAR would in theory be able to control for the talent exhibited by the player over the course of their contract.

The second model is as follows:

$$WAR_i=\beta_0+\beta_{1i}ADJGAMES+\beta_{2i}YRSLEFT+\beta_{3i}AVGWAR+\beta_{4i}IMPROVE+\beta_{5i}RECORD+\beta_{5i}PREVREC+\beta_{6i}LASTYEAR+\beta_{7i}AGE+\beta_{8i}AGE^2$$

Refer to Table 4 for complete results. Note that the only difference here is the omission of the log of salary and the log of total money variables.

The model revealed positive and significant results to a one-percent significance level for the new control, *AVGWAR* (average WAR), which was encouraging. However, the model again fails to garner a negative sign for the *YRSLEFT* variable, which was highly disappointing, but not totally unexpected. It seems that even when we have an effective variable of performance, we are unable to find the results that we expect.

The with problem average WAR is that it could be skewed downward due to missed time on the field. If a player spends a significant amount of time on the Disabled List, then they are very likely to have a low WAR value for that season. This contrasts with statistics like OPS, where this is not a problem. In fact, OPS could be far too high or low than is representative of the player, since the sample size is so low and thus the variability is much higher than statistics usually are. WAR per game or WAR per plate appearance would have been better, but I did not consider this until time ran too thin.

<u>The Third and Final Model</u>

After the results still did not come out in a way consistent with the literature, it became evident that I needed to change my performance variable to a more conventional metric used in the literature. There were also considerable heteroskedasticity problems in these first two models including *WAR*, see Table 2 and Table 5 for the chi^2 values and significance levels. I landed on OPS, which I think is a more complete representation of value than slugging average, which was also taken under strong consideration. OPS refers to slugging average added to on base percentage, which I think is better than simply slugging average because walks are so important in the game of baseball, and they reflect advanced abilities of the hitters to tell balls from strikes. Slugging average of course encompasses a hitter's ability to get extra base hits, also known as power. It was time to go back to basics, as the WAR variable did not yield expected results, so I wanted to see if the model's significance levels would change like Krautmann's (1990) did when Scoggins (1993) changed the performance metric. I ended up with this equation:

$$OPS_i=\beta_0+\beta_{1i}Games+\beta_{2i}Yrsleft+\beta_{3i}Cops+\beta_{4i}Confidence+\beta_{5i}Record+\beta_{5i}Prevrec+$$
$$\beta_{6i}LastYear+\beta_{7i}Age+\beta_{8i}Age^2$$

Since I am now using a hitting statistic, *OPS*, I had to drop all of the pitchers from my sample, leaving me with a still sufficient 102 observations. It was not possible to compare the batters' OPS with the pitchers' opposing OPS, since a high OPS would be considered good performance for a hitter and a bad performance for a pitcher. *COPS* (career OPS) replaces *AVGWAR*. *CONFIDENCE* is much the same as *IMPROVE* in the previous regressions in that *CONFIDENCE* equals zero if the previous year's OPS is lower than their career average, and it is one if the OPS in the previous year was higher than their career average. This controls for self-confidence in much the same way as *IMPROVE*.

OPS unfortunately had the same heteroskedasticity issues that the WAR model had; see Table 8 for the results of the test. Luckily, heteroskedasticity does not bias the results, so any incorrect signs cannot be attributed to this issue with the error terms being correlated with a variable. Multicollinearity was not an issue whatsoever, and this can be confirmed in Table 9.

Unfortunately, only two of the variables had statistically significant results in the final model, *COPS* and *RECORD*. They had the positive signs that were anticipated. One other positive from this model was that the *YEARSLEFT* variable had a negative coefficient. It was extremely small and insignificant, but it was the only "correct" or expected negative coefficient that I was able to achieve in any of the models.

Conclusion

This study was unfortunately unable to provide statistically significant evidence of shirking behavior in Major League Baseball. This is likely due to my lack of information regarding what is standard in the literature when it comes to vectors of controls. It also may have been due with the WAR statistic in some fashion, but there must be a way to find the shirking relationship (if it really exists, there is no consensus on this matter) with WAR as the performance measure. WAR is widely regarded as a successful, representative measure of player value of Major League Baseball players, so there is no reason that this measure is simply immune to the phenomenon. The results that I got were very frustrating, and I would welcome any critique regarding a more apt way to find this relationship.

These models also raise some interesting questions, for example, how player options and team options affect a player's performance. It would be fascinating to have a dummy variable for if the player was over or underpaid, going into the option year and how that affects performance and likelihood that the player's option would be picked up. A player who believes himself to be underpaid during an option season would perceive themselves to essentially be in a contract season, whereas someone who perceives themselves to be overpaid in an option season in which they retain control of the option would be essentially the same as a long-term contract. Proving this empirically would be quite interesting. This would be relatively easy to do, as many different models have estimated the value of wins, so one can compare this number to their WAR and salary to create a dummy variable for the player is overpaid or underpaid.

More research needs to be done regarding pitchers as well. Since Marks (2017) was able to compare different positions in football, including comparisons between offensive, defensive, and special teams players in the same model, it seems like a logical next step to compare hitters and pitchers in baseball as well, regardless of their differences in skill set. Perhaps using a WAR per game played statistic would have been a more apt performance measure than either of the two performance measures used in my first two models. WAR is among the most important and most well cited statistics in baseball, so more work needs to be done to incorporate this statistic and understand its place in sports economics. It is very curious to me that simply changing my performance measure from WAR to OPS changed the sign of my years left variable. Why was I getting the wrong sign in the first place? It would have been interesting to have included dummy variables if it was the final year in the player's contract, and the first year of the player's contract, just like the Maxcy et al. (2002) paper.

On the positive side of this thesis, I was able to find a statistically significant relationship between a team's record during that year and the performance of player's on free agent deals during that year. I was also able to find a positive relationship between improvement over the average performance in the previous season and that player's performance the following year. This self-efficacy control is unique, and the control variable, or perhaps a similar variable with the same idea behind it, could become a significant part of sports economics moving forward.

Works Cited

Fees, E. Gerfin, M. Muehlheusser, G. (2010). “The Incentive Effects of Long-Term Contracts on Performance: Evidence from a Natural Experiment in European Soccer”. *Swiss National Science Foundation.* Pp. 1-23

Feltz, D. (1988). “Self-Confidence and Sports Performance”. *Exercise and Sport Science*

Reviews. Vol. 16. Pp 278-294.

Jean, N. (2010). “Performance Variation in the NBA: Guaranteed Contracts and the Contract Year Phenomenon”. *Trinity College of Duke University.* Pp. 1-47.

Holmstrem, B., (1979). “Moral Hazard and Observability”. *The Bell Journal of Economics* Vol. 10. No. 1. Pp 74-91.

Knowles, G.; Murray, J.; Sherony, K., Haupert, M. (2003). “Shirking in Major League

Baseball in the Era of the Reserve Clause". *A Journal of Baseball History and Culture.* Pp. 59-71.

Krautmann, A. (April ,1990). "Shirking or Stochastic Productivity in Major League Baseball"*. Southern Economic Journal*. Vol. 56. No. 4. Pp 961-968.

Krautmann, A. Donley, T (2009). "Shirking in Major League Baseball Revisited". *Journal of Sports Economics*. Pp 292- 303.

Krautmann, A. Oppenheimer. M. (2002). "Contract Length and the Return to Performance In Major League Baseball". *Journal of Sports Economics.* Vol. 3. No. 1. Pp 6-16.

Krautmann, A. Solow, J. (2009). "The Dynamics of Performance Over the Duration of Major League Baseball Contracts". *Journal of Sports Economics*. Vol. 10 No. 1

Marks, G. (2017). "Pay for Play: Shirking in the NFL". *Undergraduate Honors Thesis.* 1391*.* pp 1-19

Maxcy, J. (2004). "Motivating Long-Term Employment Contracts: Risk Management in Major League Baseball". *Managerial and Decision Economics*. Pp 109-120.

Maxcy, J. Fort, R., Krautmann, A. (2002). "The Effectiveness of Incentive Mechanism in Major League Baseball". *Journal of Sports Economics.* Vol. 3. No. 3. Pp. 246-255.

Pedicelli, L. (2015). "Long-Term Guaranteed Contracts in the NBA: A Case for Moral Hazard". *Issues in Political Economy*. Vol. 24. Pp. 85-100.

Scoggins, J. (1993). "Shirking or Stochoastic Productivity in Major League Baseball: Comment". *Southern Economic Journal*. Vol. 60. No. 1. Pp 239-240.

Stankiewicz, K. (2009). "Length of Contracts and the Effect on the Performance of MLB Players". *The Park Place Economist*. Vol 17. No. 1. Pp 76-83.

Appendix

Key:

*=10 percent significance level

**=5 percent significance level

***=1 percent significance level

Table 1: First Model Results

Source	SS
Model	237.276249
Residual	395.227029
Total	632.503278

war	Coef.
adjgames	0.02***
yrsleft	0.24***
logsalary	0.19
logtmoney	-0.08
record	7.99***
prevrec	-2.25
age	-0.01
age2	-0.00
lastyear	-0.31
improve	.40
_cons	-5.24
N	180
R^2_{adj}	.34

Table 2: First Model Heteroskedasticty Test: Breusch-Pagan

chi2(1) = 6.56**

Table 3: First Model Multicollinearity Test

Variable	VIF	1/VIF
logtmoney	4.91	0.21
logsalary	3.80	0.26
yrsleft	1.71	0.59
prevrec	1.22	0.82
adjgames	1.24	0.85
record	1.17	0.86
lastyear	1.13	0.89
improve	1.130	0.91
Mean VIF	2.15	

Table 4: Second Model Results

Source	SS
Model	278.448392
Residual	354.054885

Total	632.503278

war	Coef.
adjgames	0.02***
yrsleft	0.09
avgwar	0.46***
improve	0.50**
record	7.97***
prevrec	-2.37
lastyear	-0.20
age	0.21
age2	-0.00
_cons	-6.55
N	180
R^2_{adj}	.41

Table 5: Second Model Heteroskedasticity Test: Breusch-Pagan

Breusch-Pagan / Cook-Weisberg test for heteroskedasticity
Ho: Constant variance
Variables: fitted values of war

chi2(1) = 10.07***

Table 6: Second Model Multicollinearity Test

Variable	VIF	1/VIF
yrsleft	1.42	0.70
avgwar	1.39	0.72
adjgames	1.21	0.82
prevrec	1.21	0.83
record	1.16	0.86
lastyear	1.12	0.89
improve	1.11	0.90
Mean VIF	1.23	

Table 7: Third Model Results

ops	Coef.
cops	0.81***
age2	0.00
age	-0.05
yrsleft	-0.00
record	0.34**
prevrec	-0.12
lastyear	-0.01
confidence	0.01
_cons	1.06
N	102
R^2_{adj}	.30

Table 8: Third Model Heteroskedasticity Test: Breusch-Pagan

chi2(1)	=	7.00***

Table 9: Third Model Multicollinearity Test

Variable	VIF	1/VIF
yrsleft	1.31	0.76
cops	1.28	0.78
prevrec	1.2	0.83
record	1.16	0.87
lastyear	1.13	0.89
confidence	1.07	0.93
Mean VIF	1.19	