Estimating the Effects of Age on NHL Player Performance

James A. Brander*

Sauder School of Business, University of British Columbia

Edward J. Egan

National Bureau of Economic Research (NBER)

Louisa Yeung

Sauder School of Business, University of British Columbia

August 2013

Acknowledgements: We thank Keith Head and Joanna Zhu for helpful comments.

* corresponding author
Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, Canada V6T 1Z2.
email: James.Brander@sauder.ubc.ca.

Abstract:

Using NHL data covering the NHL seasons from 1997-98 through 2011-12, we examine the effect of age on scoring performance and on plus-minus for NHL skaters (non-goalies). We also consider the effect of age on save percentage for goaltenders. We focus primarily on parametric and nonparametric fixed-effects regression methods that we believe provide the best available correction for selection bias. We also consider three other methods for the purposes of comparison: a method based on the best performances over time, a method based on the age distribution of players in the NHL, and what we call a "naïve" specification that does not correct for selection bias. All methods except the naïve specification provide reasonably consistent results (with some interesting but understandable differences). The naïve specification is very different and gives implausible results, indicating that correcting for selection bias is very important. Our best estimate of the peak age for both scoring and plus-minus for forwards is about 26. For defencemen the peak age for both scoring and plus-minus is later – at about 28. However, both forwards and defencemen exhibit near peak performance over a fairly wide age range going from about 23 to the early 30s for forwards to the mid-30s for defencemen. Strikingly, goaltenders display very little systematic variation in save percentage performance over most of the age range between 20 and 40.

Keywords: NHL, fixed-effects regression, age and performance, selection bias

1. INTRODUCTION

The effect of age on performance in sports is a subject of longstanding interest that has attracted considerable academic attention. The age-performance relationship in major league baseball (MLB) has been extensively studied, starting with the work of sabermetric[1] pioneer Bill James (1982) and including considerable academic research, such as Albert (2002), Fair (2008), and Bradbury (2009). Individual timed sports such as running, cycling, swimming and triathlon have also been extensively studied, as described in a recent review article by Lepers, Knechtle, and Stapley (2013), and many other sports have received at least some academic attention, particularly golf, as in Tiruneh (2010) .

The primary objective of this paper is to provide an assessment of the age-performance relationship in the National Hockey League (NHL). We assess and compare the age-performance relationship for the three different position categories – forwards, defencemen, and goaltenders. We focus primarily on parametric and nonparametric fixed effects regression methods, which we believe are the best approaches, but we also compare the resulting estimates with information from participation data, with what we call the "elite performance" method, and with a "naïve" method that does not correct for selection bias. Therefore, in addition to providing what we believe is the most reliable assessment of age effects on performance in the NHL currently available, we also seek to provide methodological innovations to the literature on age effects in sports more generally by comparing a set of relevant methods. Among other things, comparing the elite performance method with our fixed effects regressions allows us to assess whether elite players differ from the NHL average in their age-performance profile.

There is a conventional wisdom about the age-performance relationship in the NHL that can be inferred from many NHL "blogs", bulletin boards, and from the general hockey media. See, for example, Chen (2010), who quotes an NHL general manager as saying that a 27-year old forward has "his best hockey years ahead of him". Chen goes on to do his own assessment based on the careers of 11 star players and concludes that "peak performance does seem to follow the general notion of 27-32" as the peak period. Many other blogs provide similar comments. However, as originally pointed out by James (1982) and emphasized by many others, assessing the age-performance relationship is susceptible to the problem of selection bias. Although this problem is well-known, correcting for it effectively is not easy and there is much to be gained from a careful analysis of the data using formal statistical methods that seek to explicitly correct for selection bias.

The selection bias problem can be understood by considering the effect of looking at how players of different ages perform in a given season. It would, for example, be very misleading to take the difference between the average NHL performance of 19 year-olds and the average performance

[1] Sabermetrics refers to the statistical analysis of baseball. The term was originally coined by Bill James and is derived from the acronym SABR, which stands for the Society for American Baseball Research.,

of 26 year-olds in a given year as a measure of how much players improve between the ages of 19 and 26. Relatively few 19 years-olds play in the NHL as only the very best 19 year-olds get significant NHL playing time. Therefore, this comparison based on cross-sectional variation in performance creates bias by selecting only the very best 19 year-olds to compare with the full range of ability levels at age 26 (the modal age). This selection bias could lead to a considerable understatement of the gains in performance that occur in a player's early 20s. Similarly, only the very best players play into their late 30s, so comparing their performance with the full range of players at age 26 would understate the extent of age-related decline. We refer to the method of simply comparing the performance of NHL players of different ages without correcting for selection bias as the "naïve" method of estimating the age-performance relationship.

A method of correcting for selection bias used by James (1982) and Albert (2002) for baseball, and by many others for a variety of sports, is what we call the *elite player* method. This method selects a small number of elite players with long careers and tracks how their performance varies with age. If each of these players reaches his or her peak performance at age 26, for example, then we might reasonably conclude that 26 is the general age of peak performance. The logic of this method is that estimates of age-performance relationships are based on comparisons over time for a given player (within-player comparisons) rather than on comparisons across players.

Modern panel data fixed-effect regression methods are a generalization of this elite player approach, except that all players in the data are used, greatly increasing the reliability of the results as a guide to typical player performance. The effects of age on performance are based entirely on within-player comparisons – how each player's performance evolves over time. We use both parametric and nonparametric methods. A parametric method assumes a particular functional form for the aging function – the dependence of performance on age – and then estimates the parameters of this function. A non-parametric method does not assume a specific functional form and instead estimates the predicted performance level at each age numerically. Using both approaches is useful as the comparison of the two methods is instructive and because information learned from a nonparametric approach can be used to inform how to proceed in parametric analysis, improving the confidence we can place in parametric methods.

Another method that is sometimes used to assess the effect of age on performance is the *elite performance method* which compares the best performances at different age levels as, for example, in Berthelot et. al. (2012) for a range of activities.[2] We use an elite performance method here looking, in particular, at the top 10 players for each position in each age category. This method corrects for the primary selection bias problem, and it provides additional information of interest. Specifically, it assesses the age performance relationship for the very

[2] We also note related work by Kovalchik and Stefani (2013) that uses performances of Olympic medalists over time to assess time-related improvements in athletic performance.

best players, which might differ from the average age-performance relationship estimated by regression methods using the full sample of NHL players.

Throughout this paper we refer to player "performance". For NHL "skaters" (forwards and defencemen) we use points scored and plus-minus. We acknowledge that these measures, while of great interest, have well-known limitations. Using points scored (goals plus assists) has the limitation that it considers only the offensive half of the game even though preventing goals is just as important as scoring. This measure is particularly incomplete for defencemen, as preventing goals is their primary function. The plus-minus statistic is the number of even strength and shorthanded goals scored by a player's team while that player is on the ice minus the number of even strength and shorthanded goals scored by the opposing team while that player on the ice. Thus plus-minus attempts to incorporate both offensive and defensive contributions. However, it is subject to the limitation that it depends in large part on the quality of other players on the ice – both teammates and opponents. For goaltenders we use save percentage – the ratio of saves to shots on goal, which is widely accepted as a measure of performance.

Considerable effort has gone into developing more accurate measures of player performance than points scored and plus-minus. See, in particular, Gramacy, Jensen and Taddy (2013), Macdonald (2011), and the various metrics described at www.behindthenet.ca and elsewhere. However, we focus here on scoring and plus-minus for a variety of reasons. First, these metrics remain by far the most transparent and widely used measures of player performance. A related point is that extensive data on these measures is readily available. And they certainly are closely related to performance, even if not perfect measures.

Furthermore, as described more fully in section 4, we believe that our implementation of fixed-effects regression largely corrects for deficiencies in points scored and plus-minus as performance measures because this method looks at within-player comparisons. As long as overall performance for a given player over time is correlated with scoring and/or plus-minus for that player, as seems likely, then our estimates of the age profile of scoring and plus-minus performance would also apply to overall performance. In any case, we would suggest that the effects of age on scoring and on plus-minus (and on save percentage) are of interest in themselves, even if these metrics are not the best possible measures of performance.

The remainder of this paper is organized as follows. Section 2 describes the data. Section 3 presents evidence of performance based on participation data. Section 4 describes our fixed-effects regression methodology and presents our main results. Section 5 provides comparative results based on the elite performance method and the naïve method and assesses the importance of selection bias in NHL data. Section 6 provides concluding remarks.

2. Data

Our primary data source is the widely used publicly available data set provided by the NHL at www.nhl.com. We use data from the 1997-98 season through the 2011-12 season. As the 2004-05 season was lost due to a labor dispute, this comprises 14 years of data. We judged this to be a long enough period to contain sufficient time series information while being sufficiently recent to be reflective of the current situation. In particular, in the NHL as in other professional sports leagues, there are long run variations in scoring and other variables that make comparing players across different eras difficult.[3] The years covered in our analysis should be comparable.

We define age to be a player's age on January 1 – approximately mid-season. For some purposes we treat age as a continuous variable—in which case we use the player's exact age on January 1.[4] For example, a player who turned 27 on November 27, exactly 35 days (10% of a year) before January 1, would be counted as 27.1 – that player's age on January 1. For some purposes we round off age to the nearest whole number as of January 1 to create convenient age categories: 26 year-olds, 27 year-olds, etc.

In assessing scoring, it is important to decide whether to use points scored in a season, points per game played, or points per minute played. A large source of variation in total points per season is due to variation in games played. We take the view that we should correct for time lost due to injury (or for other reasons) and therefore focus on points per game and points per minute instead of points per season.

One important source of variation in points per game is minutes played. Players who play more minutes will have more opportunities to get points. We can correct for this effect by using points per minute played. It could be argued that players who get more playing time per game are those who are playing better and that increased scoring resulting from more playing time per game could therefore be viewed as a legitimate part of performance. It can also be argued that the ability to play a large number of minutes at a high level of proficiency is itself an age-related aspect of performance. These arguments would suggest using points per game.

Our view is that points per minute played is a better measure of performance but we use both points per game played and points per minute in our analysis. These two scoring measures lead to similar results, but there is a small but noteworthy difference regarding the age of peak performance for forwards. Using points per minute implies that the age of peak performance for

[3] One very thoughtful effort to use statistical methods to compare players across eras in a variety of sports, including in the NHL, is provided by Berry, Reese and Larkey (1999).

[4] We note in passing the very interesting paper by Addona and A. Yates (2010) investigating the effects of birth month on NHL performance. They investigate the hypothesis that players born early in the year have an advantage due to being grouped with children several months younger than themselves throughout childhood and therefore exhibiting better relative performance and getting more attention, better coaching, etc.

forwards is somewhat earlier than suggested by using points per game – reflecting the fact that relatively young players tend to play fewer minutes.

There is also some question as to whether we should adjust plus-minus on the basis of games played or minutes played. We have done the analysis with and without per game and per minute adjustments for plus-minus and find that it makes very little difference. This is not surprising as no bias is created in plus-minus by varying the number of games played. While the expected value of points scored is increasing in games played, the expected value of aggregate plus-minus is always nearly zero.[5] The variance of plus-minus does increase with games played, however. As the season progresses we get larger positive and negative values. A player who plays only 20 or 30 games, for example, would typically have a smaller absolute value of plus-minus than a player who plays a full season of 82 games. By using a player's plus-minus for all games played in a season (total plus-minus), our regression methods effectively put more weight on players who play more games. We view this as a desirable property. As there is no offsetting bias to be concerned about, we report results only for total plus-minus in this paper.

We do a certain amount of data-cleaning, including dropping age categories with very small sample sizes. Specifically, we drop players under 19 and over 40 from the data. Cleaning the data in this way has very little effect on the results, but it drops age categories for which the results might be dominated by one or two idiosyncratic observations and for which the sample sizes are so small that performance estimates for that age group have large standard errors and are therefore very imprecise. For example, there is only one defenceman under the age of 19 in the initial data set, along with only four 41 year-olds, three 42 year-olds, etc. Sample sizes for ages 19 to 40 are sufficient to draw meaningful statistical inferences for most purposes.

We also dropped player-years in which a player played less than 20 NHL games on the grounds that performance assessments based on that small a sample of games likely adds more noise than useful information to the analysis. There are, for example, cases where a young player called up late in the season to get some experience plays only a few minutes in two or three games but is fortunate enough to get a couple of points. Such a player could lead a team in "points per game" or "points per minute". It seems clear that such observations should be dropped from the data set. This small amount of data-cleaning has very little effect on the results, but we believe it yields more reliable analysis. After this data-cleaning process we are left with 2,033 players and a total of 9,901 player years.

We look only at regular season performance. Playoffs have the disadvantage that only about half the players in the league play in the playoffs in any given year and most of those players play a

[5] With even-strength play (also called five-on-five play) every goal results in an addition of 1 for all five skaters on the scoring team and a subtraction of 1 for all skaters on the other team, so the average is always zero. However, with shorthanded goals the four players on the penalty kill that scores the shorthanded goal record +1 while all five players on the powerplay scored against record -1. Therefore the average plus-minus is slightly negative as more players get a negative than get a positive for shorthanded goals. Given the very small relative number of shorthanded goals, this effect is essentially negligible.

relatively small number of games. And of course the quality of the opponents varies by round, creating significant difficulties in interpreting performance. Table 1 provides summary data on player age and performance. We report points for an entire season in the table even though we use derived performance measures (points per game and points per minute) in the formal analysis as ‘full-season points’ is a more familiar number for the purposes of understanding context.

Table 1: Descriptive Statistics for Player Age and Performance by Position

		25th percentile	**median**	**75th percentile**	**average**	**standard dev.**
Forwards	age	21.2	27.2	30.8	27.7	4.4
	points	18.9	33.0	51.9	30.9	23.1
	plus-minus	-7	-1	5	-0.6	10.5
Defence	age	24.7	27.6	31.3	28.1	4.5
	points	11.7	19.2	30.4	18.3	14.4
	plus-minus	-7	0	7	0.5	11.5
Goalies	age	26.2	29.0	32.6	29.4	4.3
	save %	90.0	90.7	91.6	90.7	0.012

Even these summary statistics provide some interesting information and suggest significant differences between forwards and defence. Defencemen are, in particularly, somewhat older than forwards at the 25th percentile (age 21.2 for forwards and 24.7 for defencemen), suggesting that defencemen have a somewhat different age-performance profile than forwards. And goaltenders are older still.

It is also worth emphasizing that scoring by defencemen is not all that much less important than scoring by forwards. The median point total for defencemen is 19.2, which is about 60% of the median point total for forwards. Even for defencemen, scoring is a very important aspect of performance.

Our data is set up as panel data. The unit of observation is the individual player and that player is tracked through time. The panel is unbalanced in the sense that different players are in the league for different time periods. This data provides the opportunity to apply a variety of statistical techniques.

We use five different methods, although two are very simple. In order to keep the logical development of our exposition clear, we start with the participation method, then discuss fixed effects regression methods, and finally consider elite performance method and the naïve method for the purposes of comparison and to assess the relative importance of selection bias.

3. Player Participation as an Indicator of Performance

The most common ages in the data should indicate the ages of peak performance. Many players who play during these most common years are not good enough to play in the NHL when younger or when older, but presumably just manage to make the NHL when playing at their peak performance levels. Therefore, the age categories when we observe these players in the NHL should be the age of peak performance. Higher quality players would play during their peak years *and* for significant parts of their non-peak years. And only the very best players would be in the NHL for age categories a long way from peak performance. Therefore this *participation method* would use the relative frequency of different player ages as an indicator of performance at different ages.

The player participation method is very easy to implement as the method does not use performance measures at all, apart from the binary indicator of whether the player is good enough to play in the NHL. All it does is to identify, by position, the number of players of each age in the NHL. The relative frequencies for forwards and defencemen are shown in figure 1 and for goalies in figure 2.

FIGURE 1 HERE

Figure 1: NHL Player Age Distributions for Forwards and Defencemen – 1997-2012

FIGURE 2 HERE

Figure 2: The NHL Player Age Distribution for Goaltenders – 1997 - 2012

The age distributions for the three positional categories are similar, but there are noticeable differences, with forwards being the youngest group, defencemen in the middle, and goaltenders the oldest. The modal or most common age for forwards is 25, but all four years in the range 24 through 27 are very similar. The peak frequency for defencemen is a year later, at age 26, in the middle of a three year peak range covering ages 25 through 27. The age distributions of both forwards and defencemen are skewed to the right. Goaltenders have their peak frequency at age 28 and a clear four year peak period that runs from 26 through 29. Therefore, if we use peak participation as an indicator of peak performance we would predict peak performance for forwards at age 25, for defencemen at age 26, and for goalies at age 28 with near peak performance for one or two years on either side.

This simple analysis based on figures 1 and 2 contains useful information but is incomplete in several ways. First, the precise peak for each position is determined by the marginal players – players just barely good enough to make the NHL for only a few years. Quite possibly higher quality players might peak slightly later. Indeed it is likely that the players who have long careers are precisely those who continue to learn and improve for longer periods. Second, some players drop out due to injury rather than declining underlying ability and, in addition, many marginal players leave the NHL voluntarily once they reach their late 20s rather than bounce back and forth between the NHL and the minor leagues even though they might still be good enough to continue in this marginal role. [6]

It seems likely that these sources of bias would have only a small effect, so perhaps the participation method would be effective in identifying the age of peak performance. However, it also seems unlikely that participation frequencies would track the relative level of performance for the full range of relevant ages, even if it does pick out the age of peak performance. In any case, it is instructive to compare inferences drawn from participation data with the results obtained from regression methods.

4. Fixed-Effects Regression Estimates of the Age-Performance Relationship

A fixed-effects regression is based on a regression specification of the following type:

$$y_{it} = \beta_0 + f(x_{it}) + u_i + e_{it} \tag{1}$$

where y_{it} is the performance of player i at time t, β_0 is a constant, x_{it} is player age, $f(x_{it})$ is some function of player age, u_i is a fixed effect for player i, and e_{it} is a random error applying to player i at time t. If player i is of above average general ability he will have a positive fixed effect: $u_i > 0$.

The function $f(x)$ is sometimes referred to as the aging function as it shows the effects of age on performance while β_0 and u_i are independent of age. This function is assumed to be the same for all players, with variation across players being captured by the fixed effect and by the error term. If we assume that we know the functional form f (often taken to be quadratic) then the resulting regression is parametric. If we assume that the form of f is unknown and we seek to construct f numerically using an estimation procedure, the approach is nonparametric.[7]

[6] As of 2013 a typical marginal player who will divide a season between the NHL and the minor league American Hockey League (AHL) has a "two-way contract" that might pay $500,000 to $900,000 per year on a pro-rated basis for games played in the NHL, but perhaps only $70,000 or $80,000 per year on a pro-rated basis for games played in AHL. Salary data can be found on the website capgeek.com maintained by Matthew Wuest.

[7] A semi-parametric method would include an unknown functional form *f* over some explanatory variables and some other specified functional form over other explanatory variables. It could be argued that specifying linear fixed effects makes our approach semi-parametric but it is more common to refer to this model as a nonparametric model with fixed effects. See, for example, Henderson, Carroll and Li (2008).

Parametric methods have the disadvantage that the estimates and the significance levels are subject to the assumption that the assumed functional form is correct. An incorrect functional form can lead to significant biases and other errors. Nonparametric estimation has the advantage that the form of the regression function is determined by the data. However, much more data is required in order to achieve statistical significance than with parametric methods, precisely because the data must determine the form of regression model. We view both parametric and nonparametric methods as being useful in this case. We use the panel data fixed-effects procedures available in the statistical program STATA 11 and would recommend the textbook by Wooldridge (2008) as a good source on panel data methods.

The logic of the fixed-effects approach is that estimates of the effect of player age are determined by looking at changes over time for individual players. Suppose player A scores 30 points at age 22 and 40 points at age 26, while player B scores 60 points at age 22 and 70 points at age 26. Using a fixed effects estimator we would infer that 26 year-olds tend to score about 10 points more than 22 year-olds. The fact that player B scores more as a 22 year-old than player A scores as a 26-year old would not matter. The fixed effects estimator is sometimes called a *within* estimator because it is based on variation within each unit. The fact that some players do not play in the NHL early in their careers and that some retire early even when they are good enough to continue playing does not bias the analysis in any way. Thus, the fixed effects estimator is precisely what we want to use to solve the selection bias problem discussed in the introduction.

Fixed effects estimators are widely used in general, although we have found only a few examples in sports. Such applications in sports include Fair (2008) for baseball, Arkes (2010) for golf, and Broadie & Rendleman Jr. (2013) for basketball. In all three cases fixed-effects are used to identify and control for individual player ability, as in this paper. It is also possible to use a *random effects* approach, which incorporates cross-sectional as well as within-player variation over time to estimate the age-performance relationship. Computationally, the random effects estimator is a weighted average of the fixed effects estimator, which is unbiased in our case, and an estimator based on cross-sectional variation, which is biased in our case. It therefore seems clear that the fixed effects estimator is preferred. However, there is relatively little difference in the results if random effects estimation is used, reflecting the fact that the data for this case imply putting little weight on the cross-sectional aspect of variation by age. We therefore report only the fixed-effects panel regressions.

We have seen only one paper that uses random effects to estimate the effect of age on performance in a sport, which is Bradbury (2009) for baseball. He reports the random effects results rather than the fixed effects results on the grounds that when the fixed effects estimator was used in combination with a correction for possible serial correlation in the errors, the results were implausible. We note that computational problems can arise when estimating fixed effects models with a correction for serial correlation. However, we have tried the serial correlation correction without difficulty in our case and find very similar results to those we report. We actually use a more general correction that deals with both possible serial correlation and heteroscedasticity by using the cluster option in STATA 11.

The fixed effects estimator should correct for the major selection bias problems. In addition it mitigates other limitations in our performance measures. Consider, for example, points scored as a measure of performance. As previously noted, scoring looks only at offensive performance and neglects defensive performance. However, our assessment of the effect of age on performance is based on each player's changes in performance over time. Therefore, as long as a player's scoring performance is positively correlated with that player's defensive performance, as seems likely in most cases, focusing on changes in scoring will also provide an indication of changes in overall performance. The relationship between points scored and plus-minus is in fact strong enough to imply a positive correlation between improvements in scoring and improvements in defensive play.

The fixed-effects approach should also help with the problem that plus-minus is influenced by the quality of other players – teammates and opponents. Generally, speaking, players on weak teams will tend to have negative plus-minus numbers while players on good teams will tend to have positive plus minus numbers. However, if a player on a bad team improves from -10 to -5 that increment of 5 counts just as much an increment of 5 for a player on a good team whose plus-minus statistic moves from, for example, 5 to 10. And more generally, if we can assume that the net quality of other players on the ice (teammate quality minus opponent quality) is uncorrelated with age, then leaving this net quality out of the regression would not bias the estimated effect of age. As a check, we include some specifications with team goal differential as a control for the parametric fixed effects estimation and find that it has no meaningful effect on the estimated age of peak performance.

4.1 Nonparametric Regression

For our nonparametric regressions we use the simplest available method for estimating *f*. Specifically, we continue to treat age in one-year intervals as we did to estimate participation frequencies (19 year-olds, 20 year-olds, etc.). We then estimate the expected performance for each age level by regressing a measure of performance (such as points scored) on age-specific dummy variables, also incorporating player-specific fixed effects by using the STATA fixed effects estimator. We treat 19-year olds as the base category, so the estimated coefficient for each age shows the expected difference at that age relative to age 19. For each age category, the sum of the coefficient on the age-specific dummy variable and the constant (which provides the coefficient for 19 year olds) then provides the expected performance at that age for an average NHL player (i.e. for $u_i = 0$).

We carry out this exercise for forwards, defencemen and goalies separately. The results are shown in table 2. We multiply our points per game measure by 82 (the number of games in a season) and we multiply the points per minute number by 60 (the number of minutes in a non-overtime game) to make those numbers easier to interpret. Obviously this scaling process has no effect on the age-related pattern of performance.

Table 2: Age-Performance Relationships for NHL Forwards

Age	Pts. per 82 games	Stand. error	Pts. per 60 min.	Stand. Error	Plus-Minus	Standard Error
19	13.3	2.65	1.18	0.082	-7.48	1.76
20	23.8	2.40	1.49	0.080	-3.68	1.81
21	28.8	2.65	1.64	0.084	-1.16	1.82
22	32.5	2.62	**1.70**	0.082	-0.25	1.75
23	36.2	2.67	**1.79**	0.084	**1.27**	1.81
24	**39.4**	2.70	**1.85**	0.083	**1.40**	1.86
25	**40.0**	2.72	**1.83**	0.085	**1.12**	1.84
26	**40.1**	2.75	**1.81**	0.085	0.22	1.84
27	**41.6**	2.80	**1.84**	0.088	1.07	1.86
28	**41.9**	2.84	**1.84**	0.090	0.38	1.88
29	**41.6**	2.78	**1.83**	0.088	0.07	1.87
30	**39.8**	2.82	**1.75**	0.089	-1.09	1.86
31	**39.1**	2.84	**1.76**	0.089	-1.47	1.90
32	36.9	2.80	**1.70**	0.088	-1.90	1.89
33	35.7	2.93	1.66	0.092	-1.39	1.93
34	34.0	2.93	1.62	0.092	-2.68	1.96
35	30.6	3.02	1.50	0.095	-3.28	2.02
36	26.9	3.07	1.42	0.098	-4.55	2.07
37	25.5	3.18	1.39	0.101	-5.40	2.17
38	20.5	3.22	1.24	0.106	-11.90	2.40
39	15.7	3.55	1.16	0.113	-10.64	2.39
40	13.5	5.00	1.21	0.154	-10.66	3.71

For forwards, peak performance in points per game is reached at age 28, but the entire 8-year period covering ages 24 through 31 (shown in bold) exhibits very similar scoring performance (within 90% of the peak performance level). If, however, we focus on points per minute played, we implicitly adjust for the fact that younger players tend to play fewer minutes. This correction reduces the estimated age of peak performance. However, there is very little variation over a fairly long period of near-peak performance (defined as within 90% of the peak) – covering 11 years from age 22 through 32. Standard errors are shown to the right of the point estimates.

Forwards exhibit a fairly clear pattern in plus-minus. The initial plus-minus at age 19 is not very good, but forwards improve consistently for the 5-year period from 19 through 23. At age 23 they enter a 3-year peak period, then have some decline in the late 20s and a more significant decline through their 30s. The analysis implies that the average forward would have a large negative plus-minus in his late 30s and at age 40. Of course the actual players still playing at ages 38 to 40 are not average players. They are much better than average players so they can still compete effectively. Therefore, the actual plus-minus scores of the players still playing at that age are much better than shown in table 2. Even so, such players are not nearly as good in their late 30s as they were seven or eight years earlier and that decline in performance contributes to the estimates in table 2. Plus-minus peaks earlier than points per game but at the same age as points per minute. This finding supports the use of points per minute (or points per 60 minutes) as the relevant scoring metric. Table 3 shows the results for defencemen and goaltenders.

Table 3: Age-Performance Relationships for NHL Defencemen and Goaltenders

	Defencemen						**Goaltenders**	
Age	**Pts. per 82 games**	**Stand. error**	**Pts. per 60 min.**	**Stand. Error**	**Plus-Minus**	**Stand. Error**	**Save%**	**Stand. Error**
19	11.0	2.49	0.63	0.076	-5.12	2.96	89.32	0.54
20	17.9	2.61	0.79	0.082	0.21	2.93	90.20	1.03
21	16.8	2.33	0.72	0.076	-0.84	2.98	90.08	0.82
22	19.9	2.43	0.79	0.077	-0.99	3.13	90.91	0.66
23	22.0	2.53	**0.83**	0.078	1.21	3.16	**91.02**	0.56
24	22.3	2.53	**0.83**	0.078	0.35	3.04	90.60	0.56
25	22.5	2.52	**0.82**	0.077	1.58	3.01	90.99	0.58
26	**24.2**	2.54	**0.86**	0.078	1.20	3.03	90.97	0.55
27	**24.0**	2.53	**0.84**	0.077	**1.73**	3.08	90.89	0.56
28	**24.6**	2.59	**0.87**	0.079	0.04	3.08	**91.01**	0.56
29	**25.0**	2.68	**0.87**	0.081	**1.67**	3.12	90.81	0.54
30	**24.9**	2.71	**0.87**	0.082	0.59	3.15	90.76	0.57
31	**24.4**	2.71	**0.84**	0.081	1.33	3.12	90.64	0.57
32	**24.4**	2.77	**0.86**	0.083	0.49	3.16	90.46	0.58
33	**24.3**	2.81	**0.85**	0.084	-0.77	3.22	90.64	0.58
34	22.2	2.83	0.81	0.085	-1.91	3.24	90.70	0.58
35	20.3	2.82	0.77	0.086	0.83	3.32	90.36	0.60
36	17.4	2.80	0.70	0.085	-1.08	3.28	90.10	0.63
37	18.6	2.96	0.77	0.091	-1.27	3.57	90.33	0.65
38	15.0	3.17	0.71	0.101	-0.78	3.66	89.92	0.73
39	16.8	3.40	0.79	0.103	-1.92	4.14	90.40	0.70
40	12.5	3.44	0.66	0.104	-4.32	5.49	89.61	0.64

Defencemen enter their scoring peak later than forwards – at about 26. The highest level of scoring per game occurs at age 29, but points per minute are virtually the same for ages 28, 29, and 30. Age 29 is the best year for plus-minus. Once again, however, there is a fairly long period of near-peak performance within which year to year differences are small for scoring and somewhat erratic for plus-minus. Putting scoring and plus-minus information together suggests a period of peak performance lasting from about age 26 through to about age 32.

The standard errors are fairly large compared to the differences in performance between adjacent years, reflecting the fact that the data requirements for nonparametric methods are high if we hope to obtain high levels of statistical significance.

For goaltenders the data exhibit relatively little age-related change in performance. There is some indication that very young goaltenders (19-year olds) and goaltenders at the upper limit (age 40) operate below peak performance, but the period in between is remarkably stable.

The performance inference for the age of peak performance implied by these nonparametric regressions has some similarity to the participation method suggested by figure 1 depicting NHL participation rates by age. However, the regression-based analysis suggests that players hold their peak performance longer. Thus participation rates overstate age-related decline, as is consistent with the biases inherent in the participation method discussed earlier, particularly retirement due to injury-related concerns or for voluntary reasons rather than declining performance – biases that are not shared by fixed effects regression methods.

Figure 3 illustrates a number of important aspects of the analysis. The figure shows the age-scoring relationship for points per 60 minutes played for forwards and defencemen. It also shows 90% confidence intervals and it includes quadratic trend-lines fitted to the estimated relationships using ordinary least squares.

FIGURE 3 HERE

Figure 3: Nonparametric fixed effects estimates for age effects on scoring and quadratic trend-lines

The overall pattern of the age-scoring relationships is statistically significant, but the year-to-year differences between any two adjacent years are not statistically significant except near the limits of the age range.

Furthermore, the actual nonparametric estimates are somewhat erratic. We do not really believe that expected scoring per minute for defencemen would fall from 0.79 to 0.72 as the player ages from age 20 to 21 and then rise back 0.79 at age 22. We do not have enough data to separately estimate scoring for each age category with a high degree of precision, as is a common problem with nonparametric methods. This up and down pattern could be smoothed using more sophisticated nonparametric measures based on kernel density estimation, but we do not really have enough data to use such methods effectively. An alternative method of smoothing is to fit a trend-line to the nonparametric estimates as shown in figure 3, where we fit quadratic trend-lines.

These trend-lines show that a quadratic approximation fits scoring by defencemen very well, but for forwards there are small but systematic deviations caused by the skewness in scoring performance. In fact a cubic trendline (not shown) provides a slightly better (and very good) fit to the data. However, if we are going to smooth the results in this way, it makes sense to estimate a parametric model directly, which we do in the next section.

To save space we do not show the corresponding diagrams for plus-minus or for save percentage (for goaltenders). The plus-minus estimates are, as noted earlier, similar in pattern to the scoring relationship but are more erratic, especially for defencemen. The goaltender save percentage exhibits no meaningful age-based trajectory.

4.2 Parametric Regression

The difficulty in obtaining statistical significance without very large sample sizes is often a problem with nonparametric methods. If we are prepared to assume a functional form then, conditional on that assumption, less data is normally required. Parametric methods also often have the advantage of providing a closed form algebraic representation of the relationship that can be used for a variety of purposes, including providing a compact summary of the age-performance relationship. We can think of imposing a functional form as a way of incorporating (and taking advantage of) prior information. In this case we have a lot of prior information about athletic performance implying that it rises and falls and is single-peaked, and there is considerable prior information that performance over time in many areas is well-approximated by a quadratic function.

In this case the shape of the estimated nonparametric relationship is useful in suggesting an appropriate functional form for parametric estimation. The commonly used quadratic form fits the data quite well. However, quadratic functional forms imply symmetry – that performance should decline in later years at the same rate as it improves in earlier years. Our nonparametric analysis suggests that a quadratic approximation should work well for defencemen. However, scoring by forwards is slightly but distinctly skewed to the right as a function of age, suggesting a cubic approximation. In a cubic specification, the significance of the coefficient on the cubic term provides a test of whether the skewness is significant. We estimate cubic and quadratic specifications for both forwards and defencemen. In these regressions we treat age as a continuous measure (such as 26.1 or whatever) rather than relying on integer approximations. We do not report results for goaltenders as there is no discernible age-related pattern of performance for goaltenders.

For a cubic specification we estimate an equation of the form

$$y_{it} = \beta_0 + \beta_1 x_{it} + \beta_2 x_{it}^{2} + \beta_3 x_{it}^{3} + u_i + e_{it} \qquad (2)$$

where β_0, β_1, β_2, and β_3 are regression coefficients and, as in equation (1), y_{it} is the performance of player i at time t, x_{it} is player age u_i is a fixed effect for player i, and e_{it} is a random error applying to player i at time t. To obtain a quadratic specification we constrain β_3 to be zero. Table 4 shows the major results for forwards for points per minute played and plus-minus.

Table 4: Age-Performance Regressions for NHL Forwards

	(1)	(2)	(3)	(4)	(5)
VARIABLES	**pts. per 60 min.**	**pts. per 60 min.**	**plus-minus**	**plus-minus**	**plus-minus**
age	0.282*** (9.372)	0.798*** (4.025)	3.449*** (5.844)	5.624 (1.65)	8.496** (2.024)
agesq	-0.00519*** (-10.02)	-0.0233*** (-3.383)	-0.0658*** (-6.386)	-0.1562 (-1.30)	-0.243* (-1.659)
agecub		0.000209*** (2.645)			0.00204 (1.212)
team goal differential				10.67*** (28.69)	
Constant	-2.000*** (-4.631)	-6.790*** (-3.619)	-44.38*** (-5.322)	10.675*** (28.7)	-91.28** (-2.311)
Obs	6,017	6,017	6,017	6,017	6,017
R-squared	0.73	0.73	0.33	0.38	0.33
Optimal Age	27.2	**26.6**	**26.2**	26.0	25.9
90% of peak	22-33	22-32			

***, **, and * indicate the coefficient is significantly different from zero at the 1%, 5% or 10% levels of statistically significance, respectively. t-statistics are in parentheses.

In the quadratic regression for points (result column 1) the dependent variable is points per minute and both age and age squared are highly significant with the expected signs. An optimal age of 27.2 is implied. However, performance for forwards has a statistically significant skew to the right: improvement occurs more quickly than decline. Thus the cubic term in the cubic regression is statistically significant at the 0.01 level (result column 2). Players reach near peak performance fairly early, hold the peak for quite a while, and then decline slowly. The implied age of optimal scoring performance declines slightly to 26.6. However, the range for which scoring performance is within 90% of the peak level is large – stretching from 22 to 32 or 33.

As for plus-minus performance, the cubic term is not statistically significant, suggesting that the quadratic specification is preferred. The implied optimal age for plus-minus is 26.2, just slightly less than the optimal age for scoring. We do not report a peak range as it is difficult to assess what 90% of peak performance would mean for plus-minus but, as with points scored, players maintain a performance close to peak levels for a significant period. Including team goal differential as a control variable has almost no effect on the estimated age of peak performance. Table 5 shows the age-performance regressions for defencemen.

Table 5: Age-Performance Regressions for NHL Defencemen

	(1)	(2)	(3)	(4)	(5)
VARIABLES	**pts. per 60 min.**	**pts. per 60 min.**	**plus-minus**	**plus-minus**	**plus-minus**
age	0.104*** (4.262)	0.184 (1.206)	1.866** (2.266)	1.074 (1.59)	7.517 (1.328)
agesq	-0.00181*** (-4.396)	-0.00462 (-0.869)	-0.0336** (-2.375)	-0.0192 (-1.64)	-0.231 (-1.172)
agecub		3.21e-05 (0.527)			0.00225 (1.002)
team goal differential				13.59*** (24.36)	
Constant	-0.623* (-1.751)	-1.377 (-0.955)	-24.75** (-2.088)	-23.98 (-1.45)	-77.74 (-1.452)
Obs	3,179	3,179	3,179	3,179	3,179
R-squared	0.664	0.664	0.283	0.283	0.284
Optimal Age	0.104***	0.184	1.866**	1.866**	7.517
90% of peak	**28.6**	28.3	**27.7**	27.9	27.7

***, **, and * indicate the coefficient is significantly different from zero at the 1%, 5% or 10% levels of statistically significance, respectively. t-statistics are in parentheses.

For defencemen, a cubic form adds nothing to the explanatory power of the regression and the coefficients lose their statistical significance. The improvement and decline of defencemen is very symmetric and relatively slow. Also, as suggested by our nonparametric analysis and by the participation rates by age, defencemen peak a bit later than forwards, with a point-scoring peak and plus-minus peak about 28. And defencemen stay close to peak performance later – up to about age 35 rather than only 32 or 33.

We do not report age-performance regressions for goaltenders, as there is no meaningful effect of age on performance except at the extremes of the age range. There is a lot of individual variation across goaltenders. Some do well early and fade in their 30s. Many others do better in their 30s than in their 20s. However, on average, if we compare the early 20s, late 20s, early 30s and late 30s, there is very little systematic difference in performance.

5. The Elite Performance Method, the Naïve Method and Selection Bias

5.1 The Elite Performance Method

This method focuses on the best performances at each level. Before providing formal analysis we show the top 10 performances in the data for scoring, plus-minus, and save percentage.

Table 6: Top 10 Scoring and Plus-Minus Performances

	Player	Age	Pts per 60 min		Player	Age	Plus-Minus
1	Sidney Crosby	24	5.46	1	Peter Forsberg	29	52
2	Peter Forsberg	30	4.41	2	Chris Pronger	25	52
3	Sidney Crosby	23	4.41	3	Milan Hejduk	27	52
4	Sidney Crosby	19	4.39	4	Jeff Schultz	24	50
5	Peter Forsberg	29	4.39	5	Chris Chelios	38	48
6	Mario Lemieux	35	4.36	6	Chris Pronger	23	47
7	Joe Thornton	27	4.34	7	Alex Ovechkin	24	45
8	Daniel Sedin	29	4.23	8	Joe Sakic	31	45
9	Jason Spezza	23	4.18	9	Patrik Elias	25	45
10	Alex Ovechkin	24	4.17	10	Thomas Vanek	23	47

Table 7: Top 10 Goaltending Performances

	Goaltender	Age	Save %
1	Brian Elliott	27	0.940
2	Tim Thomas	37	0.938
3	Dominik Hasek	34	0.937
4	Cory Schneider	26	0.937
5	Dwayne Roloson	34	0.933
6	Miikka Kiprusoff	27	0.933
7	Tim Thomas	35	0.933
8	Marty Turco	27	0.932
9	Dominik Hasek	33	0.932
10	David Aebischer	24	0.931

Tables 6 and 7 illustrate the most elite performances in the data. Even for elite performance assessments we will use much more data. However, these elite performances are some interest. First, it is clear that elite performances can occur over a wide age range. Among the top 10 scoring performances (points per 60 minutes played) are Sidney Crosby at age 19 and Mario Lemiuex at age 35. For plus-minus we have Chris Pronger (a defenceman) at age 23 and Chris

Chelios (also a defenceman) at age 38. Still, for both scoring and plus-minus, there is a clustering in the mid-to-late 20s for these elite performances. For goalies, on the other hand, there is no discernible clustering. Half of these top 10 performances occur over the age of 30 – and well into the 30s, while goalies of 24 and 26 are also in the top 10.

A more formal approach to using the elite performance method is to take the top 10 performances for each integer age level between 19 and 40 and then take the average of those top 10 performances. Figure 4 provides the results in a diagram for forwards for both scoring (points per 60 minutes played) and plus minus. Figure 5 shows the results for defencemen. Trend-lines fitted to these averages are also shown. Cubic trend-lines are used as they fit the data slightly better than quadratic trend-lines.

FIGURE 4 HERE

Figure 4: Elite (Top 10) Scoring and Plus-Minus Performance Average for Forwards

FIGURE 5 HERE

Figure 5: Elite (Top 10) Scoring and Plus-Minus Performance Average for Defencemen

Using the elite performance method generates very similar results to fixed-effect regression methods. The peak performance for forwards occurs about age 27 for both scoring and plus-minus and the peak age for defencemen occurs about age 29. This is slightly later (by about one year) than using the fixed effects regression method. Also, elite performances are slightly more skewed – with a longer period of near peak performance. Elite players achieve 90% of peak performance about age 21 for forwards and 22 for defencemen and maintain near peak performance until about age 37.

We interpret the results to mean that elite players improve faster initially, continue to improve for slightly longer and experience slower age-related decline. They do not experience a major drop-off in performance until their late 30s.

We do not provide a diagram for goaltenders, in large part because there is almost no variation to see. In table 8 we show the average save percentage for the top 10 goalies of each age. Only ages 22 through 37 are covered as the other age categories did not have 10 goaltenders. Elite goaltender performance from age 23 to age 35 is remarkably consistent and exhibits remarkably little variation.

Age	Save %
22	0.913
23	0.920
24	0.920
25	0.925
26	0.925
27	0.925
28	0.924
29	0.923
30	0.923
31	0.922
32	0.922
33	0.921
34	0.923
35	0.922
36	0.917
37	0.916

Table 8: Elite (top 10) save percentages for goaltenders by age

5.2 The Naïve Method and Selection Bias

The naïve method simply calculates average performance, by position, for each age category in the data. For example, there are 98 forwards of age 19 in the data. The naïve estimate of the relevant performance for a 19-year old is the average performance (points per game, points per minute, or plus-minus.) taken over these 98 observations. We determine the average performance for each age category in the same way. In the case of 26 year-olds, there are 1,134 players in the data, so the average is taken over a much larger set. The average performance by age is then the estimated age-performance relationship. An extension of the naïve method would be to fit an estimated regression line to these age-average performance combinations. As noted in the introduction this method suffers from selection bias in that only the best 19 years olds are compared with a full range of players at age 26, overstating the relative performance of 19 year olds and thereby understating the gains due to age between 19 and 26.

Figure 6 shows the result of using the naïve model to assess the age performance relationship. We show only points scored per 60 minutes to save space. The pattern for plus-minus is very similar.

FIGURE 6 HERE

Figure 6: NHL Scoring Averages by Age – Naïve Model

As can be seen in figure 6, there is no discernible single-peaked pattern relating performance to age using the naïve method. Selection bias fully masks the age-performance relationship. If anything, scoring performance seems to increase gradually with age up to the maximum age of 40. The highest scoring average is in fact achieved by 40 year old forwards! Over the entire 14 year period covered by the data there are only nineteen 40 year-old players in the sample, including players such as Mario Lemieux, Jaromir Jagr, Mark Messier, Teemu Selanne, Mike Modano, and other players of that level – the greatest scoring stars in the game over this period. At age 40 only the very best scorers remain in the game, creating a very substantial selection bias.

This selection operates at two levels. Coaches and managers will select only the best players of that age to play and, in addition, the players themselves will exercise self-selection. Many players who could play effectively at age 40 choose not to – possibly because of fear of injury, possibly because they do not want to go the effort of staying in shape (which gets increasingly difficult with age) and possibly because they want to spend time with their families or pursue other objectives. At age 40 playing is worthwhile only for the most gifted athletes – those who can command high enough salaries to make playing worthwhile and who are still good enough to make the game enjoyable to play. This self-selection effect likely explains why older players actually score more than younger players (rather than just being about the same). The fixed-effects regression methods we use in section 4 corrects for both kinds of selection bias.

Comparing the naïve model with the results from the fixed effects regression methods, or with the participation information, suggests that failure to correct for selection bias renders any analysis of the age-performance relationship essentially meaningless.

6. Concluding Remarks

The primary objective of this paper is to assess the relationship between age and performance among NHL players. We have placed considerable emphasis on solving the classic selection bias problem and we believe that we have obtained a more reliable estimate of the age performance relationship than has previously been obtained for NHL players.

Several of our results are consistent with what might be described as "conventional wisdom" and some are not. Results that fall in the "no surprise" category would include the estimated age of peak performance for both forwards and defencemen. Specifically, we believe that our best estimates are those obtained from the parametric fixed effects regression model, and we believe that the best scoring metric is points per minute played. For forwards a cubic specification is the best fit to the data and implies an optimal age of 26.6. The optimal age for plus-minus is obtained from a quadratic specification and is 26.2.For defencemen the best estimate of the optimal age for both scoring and plus-minus is obtained from a quadratic specification and is 28.6 for scoring and 27.7 for plus-minus.

These parametric estimates are consistent with the estimates obtained from nonparametric fixed-effects regression methods, which cannot meaningfully distinguish between the years in the peak period of mid to late 20 twenties for forwards and late twenties for defencemen. They are also consistent with participation data if we make an allowance for the small amount of selection bias in participation data due to injury and voluntary retirement.

Therefore, we feel that the analysis strongly suggests peak performance at about age 26 for forwards and about age 28 for defencemen. This would not come as surprise to people familiar with NHL hockey – fans, players, coaches, journalists, or researchers. The numbers are perhaps on the early side of the conventional wisdom, but not far off, and the finding that defencemen peak later than forwards would also not be surprising.

There are however some surprises in the data. One surprise is that goaltenders exhibit very little systematic variation in performance by age. It might seem that sample size could be an issue. After all, in a given season a typical team will have 14 or 15 forwards that meet the 20 game threshold for inclusion in the data, along with 7 or 8 defencemen, but normally only two goaltenders and sometimes only one. Overall there are 1,246 forwards in the data, 635 defencemen, and only 152 goalies. Still if there is a consistent pattern of performance relative to age, it should show up with 152 players observed over their NHL careers, at least using parametric regression methods. The fact is that goaltenders can perform well (or poorly) at any age within a broad range. This fact is particularly striking for elite goaltenders.

Another surprise in the data is length of the period of near-peak performance. Our sense is that the conventional wisdom exaggerates both the time taken for player development, particularly for forwards, and the importance of age-related decline. Using points per minute player we estimate that forwards reach near peak performance fairly early – about age 22 -- and hold that peak for about 10 years. And defencemen also reach something close to peak performance by age 22 or 23 and stay within 90% of peak performance up to about age 35. The effect of age on year to year variation in performance is relatively modest. It is much more important to be a good player than to be a prime age player. Much of the talk about "player development" and "upside" is misplaced. Thus the many fans, coaches, and journalists who talk about "expected improvement" for 25 or 26 year-old forwards are probably guilty of wishful thinking. Forwards who do not develop into consistent scorers by age 23 or 24 in most cases never will.

One minor surprise relates to the skewness in the pattern of development for forwards. Forwards improve rapidly in their early 20s and get very close to peak performance early. The decline from this peak occurs more slowly. For defencemen, on the other hand, the pattern is closer to being symmetric, with slower development, a later peak, and a slower age-related decline. We also note one other point that we do not believe was previously known – that plus-minus performance peaks earlier than scoring performance. This result suggests that defensive skills develop more quickly than scoring skills for both defencemen and forwards.

An additional significant contribution of our analysis is the comparison of elite players – using the elite performance method – with the overall or average pattern of NHL performance obtained using fixed effects regression. Elite players appear to peak slightly later and have a longer peak period. This is not just a matter of elite players being better and therefore being able to play effectively for a longer period. They also have a slower rate of age-related decline relative to their individual peak performance.

Finally, we would emphasize two methodological points. First, we believe that panel data statistical methods incorporating fixed-effects estimators – in both parametric and nonparametric specifications – are very good methods for dealing with selection bias in assessing the relationship between age and performance in any area. Second, by comparing the results with and without corrections for selection bias we provide an indication of the importance of selection bias is in assessing age-performance relationships and find that such problems are likely to be very serious. However, once appropriate corrections for selection bias are made, the resulting estimates of the age-performance relationship appear to be very reliable and robust.

References:

Addona, Vittorio and Philip. A. Yates. 2010. “A Closer Look at the Relative Age Effect in the National Hockey League.” *Journal of Quantitative Analysis in Sports* 6(4), Article 9: 1-17.

Albert, Jim. 2002. “Smoothing Career Trajectories of Baseball Hitters.” Unpublished manuscript, Bowling Green State University, at bayes.bgsu.edu/papers/career_trajectory.pdf.

Arkes, Jeremey. 2010. “Revisiting the Hot Hand Theory with Free Throw Data in a Multivariate Framework.” *Journal of Quantitative Analysis in Sports* 6(1): Article 2.

Berthelot, Geoffroy, Stéphane Len, Philippe Hellard, Muriel Tafflet, Marion Guillaume, Jean-Claude Vollmer, Bruno Gager, Laurent Quinquis, Andy Marc and Jean-François Toussaint. 2013. “Exponential growth combined with exponential decline explains lifetime performance evolution in individual and human species.” *Age* 34: 1001-1009.

Berry, Scott, C. Shane Reese, and Patrick D. Larkey. 1999. “Bridging Different Eras in Sports” *Journal of the American Statistical Association* 94(447): 661-676.

Bradbury, J.C. 2009. “Peak Athletic Performance and Ageing: Evidence from Baseball.” *Journal of Sports Science* 27(6): 599-610.

Broadie, Mark and Richard J. Rendleman Jr. 2013. “Are the official world golf rankings biased?” *Journal of Quantitative Analysis in Sports* 9(2): 127-140.

Chen, Mike. 2010. “When is an NHL player’s prime age?” SB Nation Special Feature at fromtherink.sbnprivate.com/2010/7/16/1572579/when-is-an-nhl-players-prime-age.

Fair, Ray C. 2008. “Estimated Age Effects in Baseball.” Journal of Quantitative Analysis in Sports 4(1): Article 1., 1-39.

Gramacy , Robert B., Shane T. Jensen and Matt Taddy. 2013. “Estimating player contribution in hockey with regularized logistic regression.” *Journal of Quantitative Analysis in Sports* 9(1):97-111.

Henderson, Daniel J., Raymond J. Carroll and Qi L. 2008. “Nonparametric estimation and testing of fixed effects panel data models.” *Journal of Econometrics* 144: 257–275.

James, Bill. 1982. *The Bill James Baseball Abstract*. New York: Ballantine Books.

Kovalchik, Stephanie Ann and Ray Stefani. 2013. “Longitudinal analyses of Olympic athletics and swimming events find no gender gap in performance improvement.” *Journal of Quantitative Analysis in Sports* 9(1): 15–24.

Lepers, Romuald, Beat Knechtle, and Paul J. Stapley. 2013. "Trends in Triathlon Performance: Effects of Sex and Age." *Sports Medicine*, http://link.springer.com/article/10.1007/s40279-013-0067-4, published online June 2013, forthcoming in print.

Macdonald, Brian. 2011. "A Regression-Based Adjusted Plus-Minus Statistic for NHL Players," *Journal of Quantitative Analysis in Sports* 7(3): Article 4, 1-29.

NHL.com. 2013: The official website of the National Hockey League, www.nhl.com.

Tiruneh, Gizachew. 2010. "Age and Winning Professional Golf Tournaments." *Journal of Quantitative Analysis in Sports* 6(1), Article 5.

Wooldridge, Jeffrey M. 2008. *Econometric Analysis of Cross Section and Panel Data*, 2nd ed. Cambridge Mass.: MIT Pess

Wuest, Matthew. 2013. CapGeek.com.

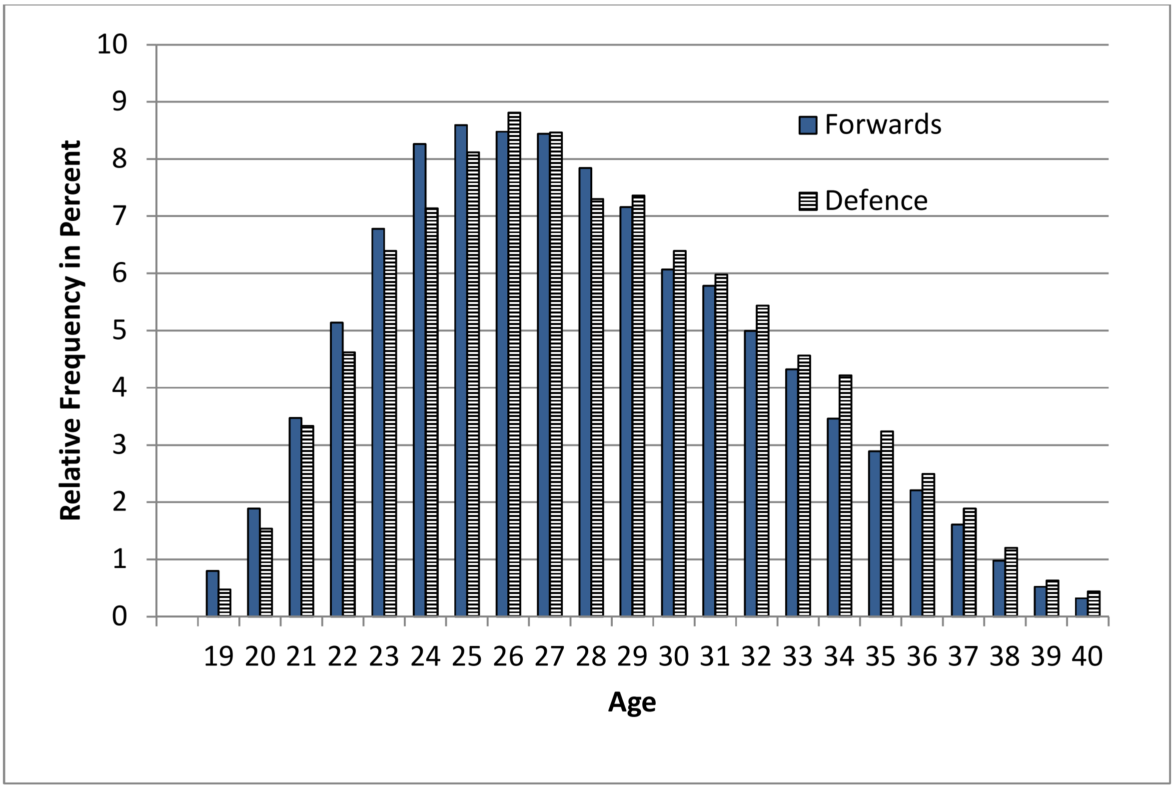

Figure 1: NHL Player Age Distributions for Forwards and Defence - 1997-2012

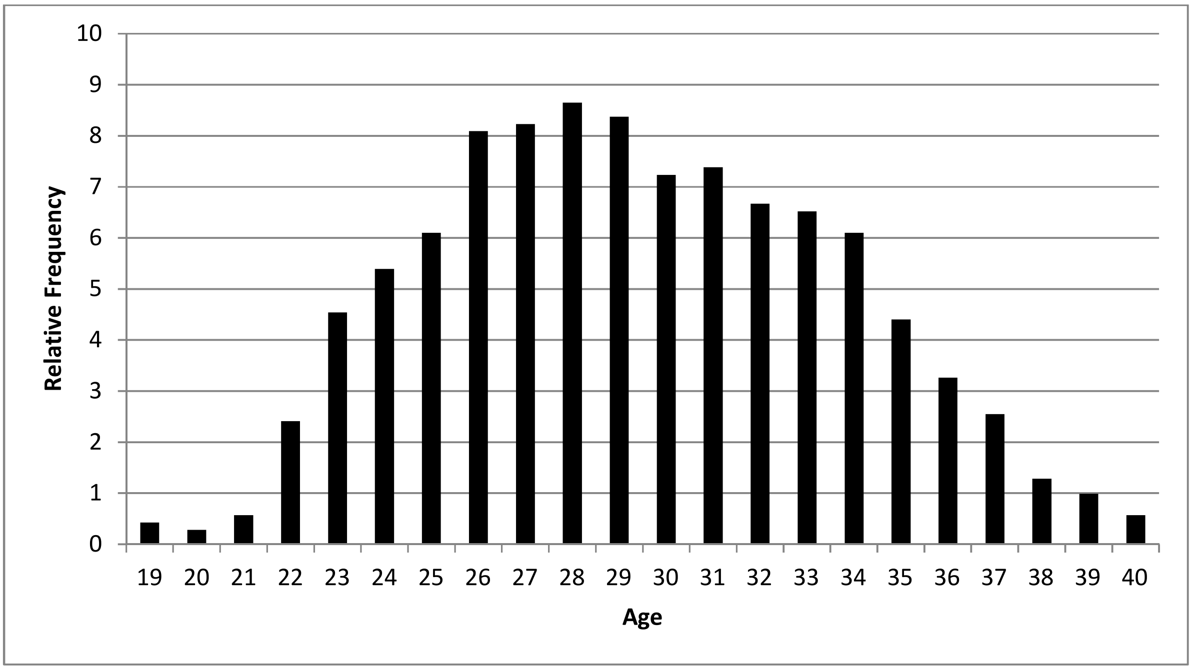

Figure 2: The NHL Player Age Distribution for Goaltenders - 1997-2012

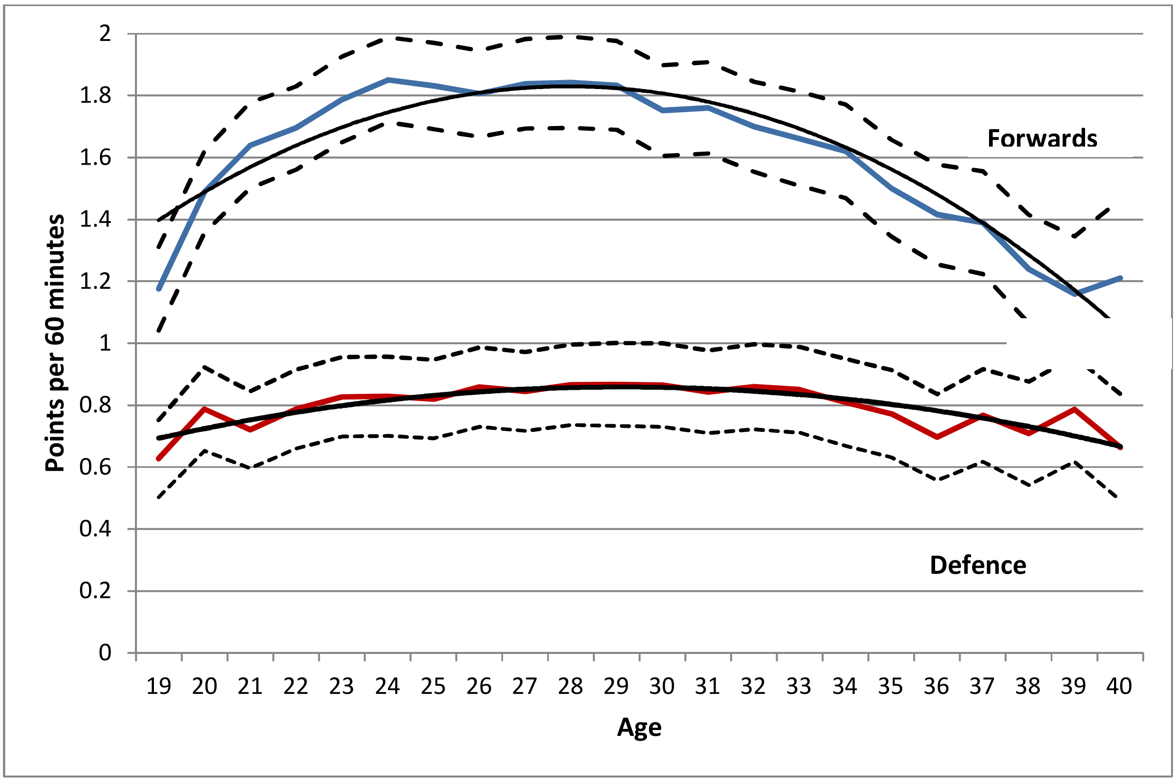

Figure 3: Nonparametric fixed effects estimates for age effects on scoring

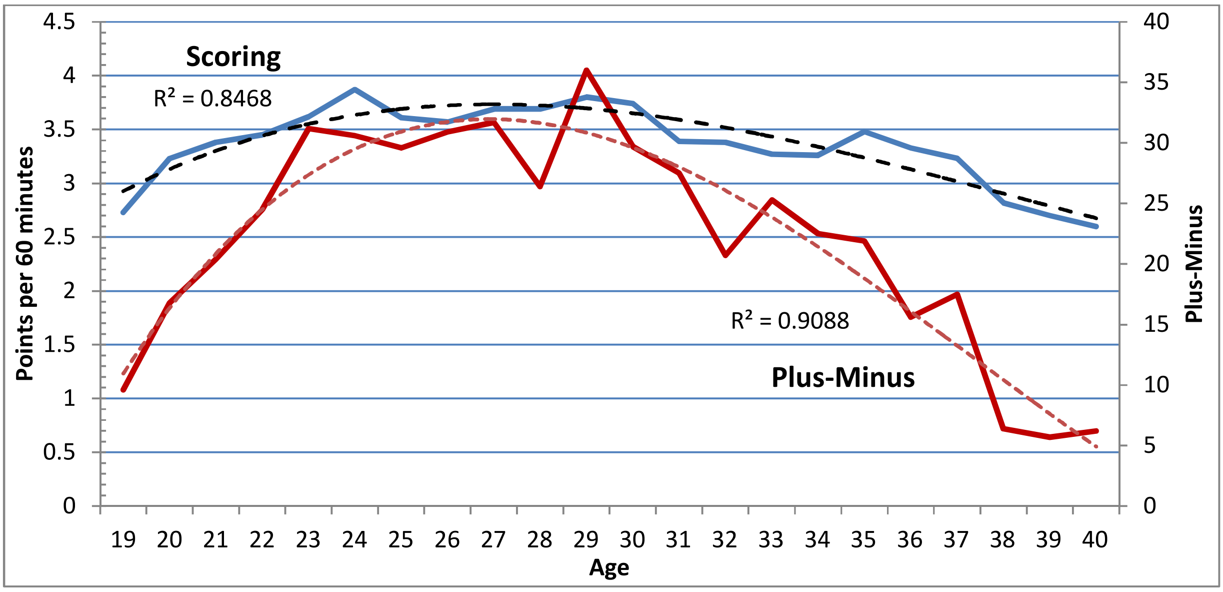

Figure 4: Elite (Top 10) Scoring and Plus-Minus Performance Average for Forwards by Age

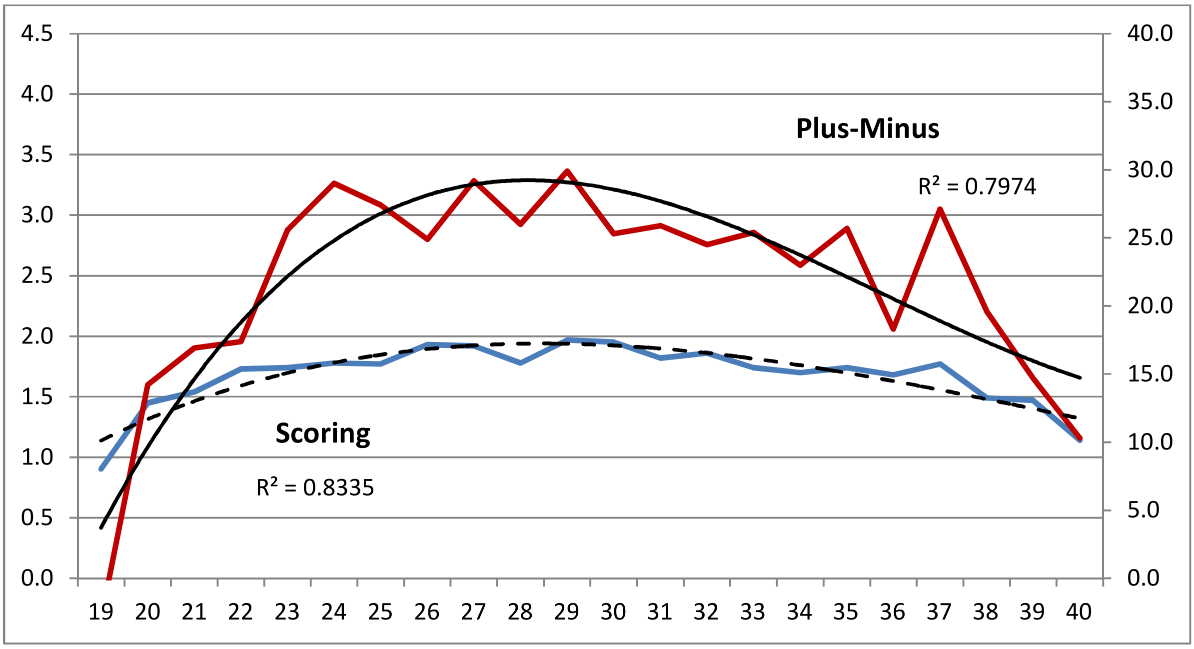

Figure 5: Elite (Top 10) Scoring and Plus-Minus Performance Average for Defencemen by Age

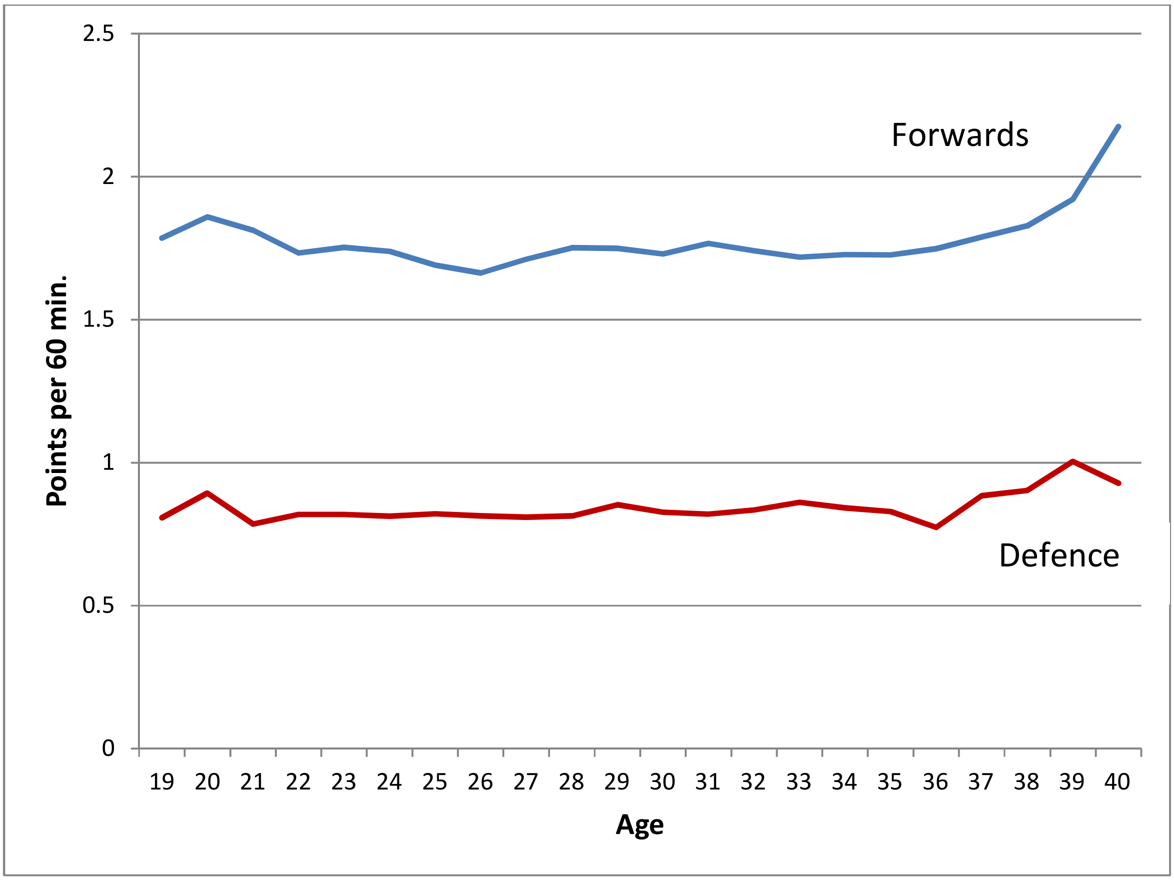

Figure 6: NHL Scoring Averages by Age – Naïve Model