

Reaction to the NHL Salary Cap

Mathieu Parisien and Professor Louis-Philippe Morin

University of Ottawa, Faculty of Social Sciences, Department of Economics. 3rd year Honours Degree: Specialization in Economics

Abstract

On February 16th, 2005, the NHL Commissioner Gary Bettman officially announced the cancellation of the 2004-2005 NHL season. The cancellation was due to dead end negotiations for a new Collective Bargaining Agreement (CBA) between the NHL Players Associations (NHLPA) and the NHL. The main issues surrounding the negotiations were the subject of league revenue distributions and team salary structure. On July 22nd, 2005, a new CBA was reached. The biggest impact the new CBA would have is the introduction of a salary cap ceiling as well as a salary cap floor. This study focuses on these new salary cap restrictions and how it has affected the distribution of player salaries across the NHL. We analyse league statistics and player salaries from the 2000-2001 season to the 2011-2012 season, to identify how teams adjusted their strategies to accommodate for to the salary cap. The study focuses on budget variation among teams, salary distribution within teams, as well as changes in team construction. The results of this research will provide insight for the optimum strategy of salary distribution to be able to construct a successful hockey team within the NHL.

Introduction

The world of professional sports has been used at times to shed light on important labour market research questions such as estimating the effect of employer monopsony on wages, investigating the extent of racial discrimination (e.g. Price and Wolfers (2010)), investigating the relevance of important theorems in economics (e.g., the Coase Theorem in Fort and Quirk (1995)), and estimating the effect of supervision and incentives on worker performance (Kahn, 2000). In this study we investigate whether the 2005 Collective Bargaining Agreement (CBA) affected the way general managers (GM) allocate wages among their players, particularly with the introduction of salary caps in the NHL. We first show that the introduction of salary caps was binding; that teams that were either below or above the predicted salary caps saw their team total salary change drastically. Then we show that teams that were above (below) the salary cap decreased (increased) the share of their teams total salary allocated to their highest paid players.

The research most similar to ours is Fort and Quirk (1995) who examine the impact of the introduction of a salary cap in the NBA (in 1984/85) on the competitive balance using data from the 1975/76 to the 1992/93 seasons. They argue that the many exceptions to the NBA salary cap weakened its potential power over the competitive balance, and in fact, seemed to have worsened it in the short run. To illustrate their point, they note that the Lakers' total payroll was well above the salary cap (by 50-80 percent) during most of the 1985-1989 seasons. Contrary to Fort and Quirk (1995), we show that the NHL salary cap was overall binding.

Methodology

To investigate the effects of the NHL salary cap, we collected public data from USA Today of every player's salary across the NHL from the year 2000 to 2011, for a total of 8654 (player-year) observations. We collapsed these data at the team-year level to construct measures of team-level salary distributions (e.g., average salary, variance, coefficient of variation, and share of team salary devoted to the five highest paid players). Our sample therefore consists of 330 observations (30 teams over 11 years). To account for some of the potential channels explaining our results, we've also identified the general managers of all 30 teams across the observed period, for a total of 64 different observations.

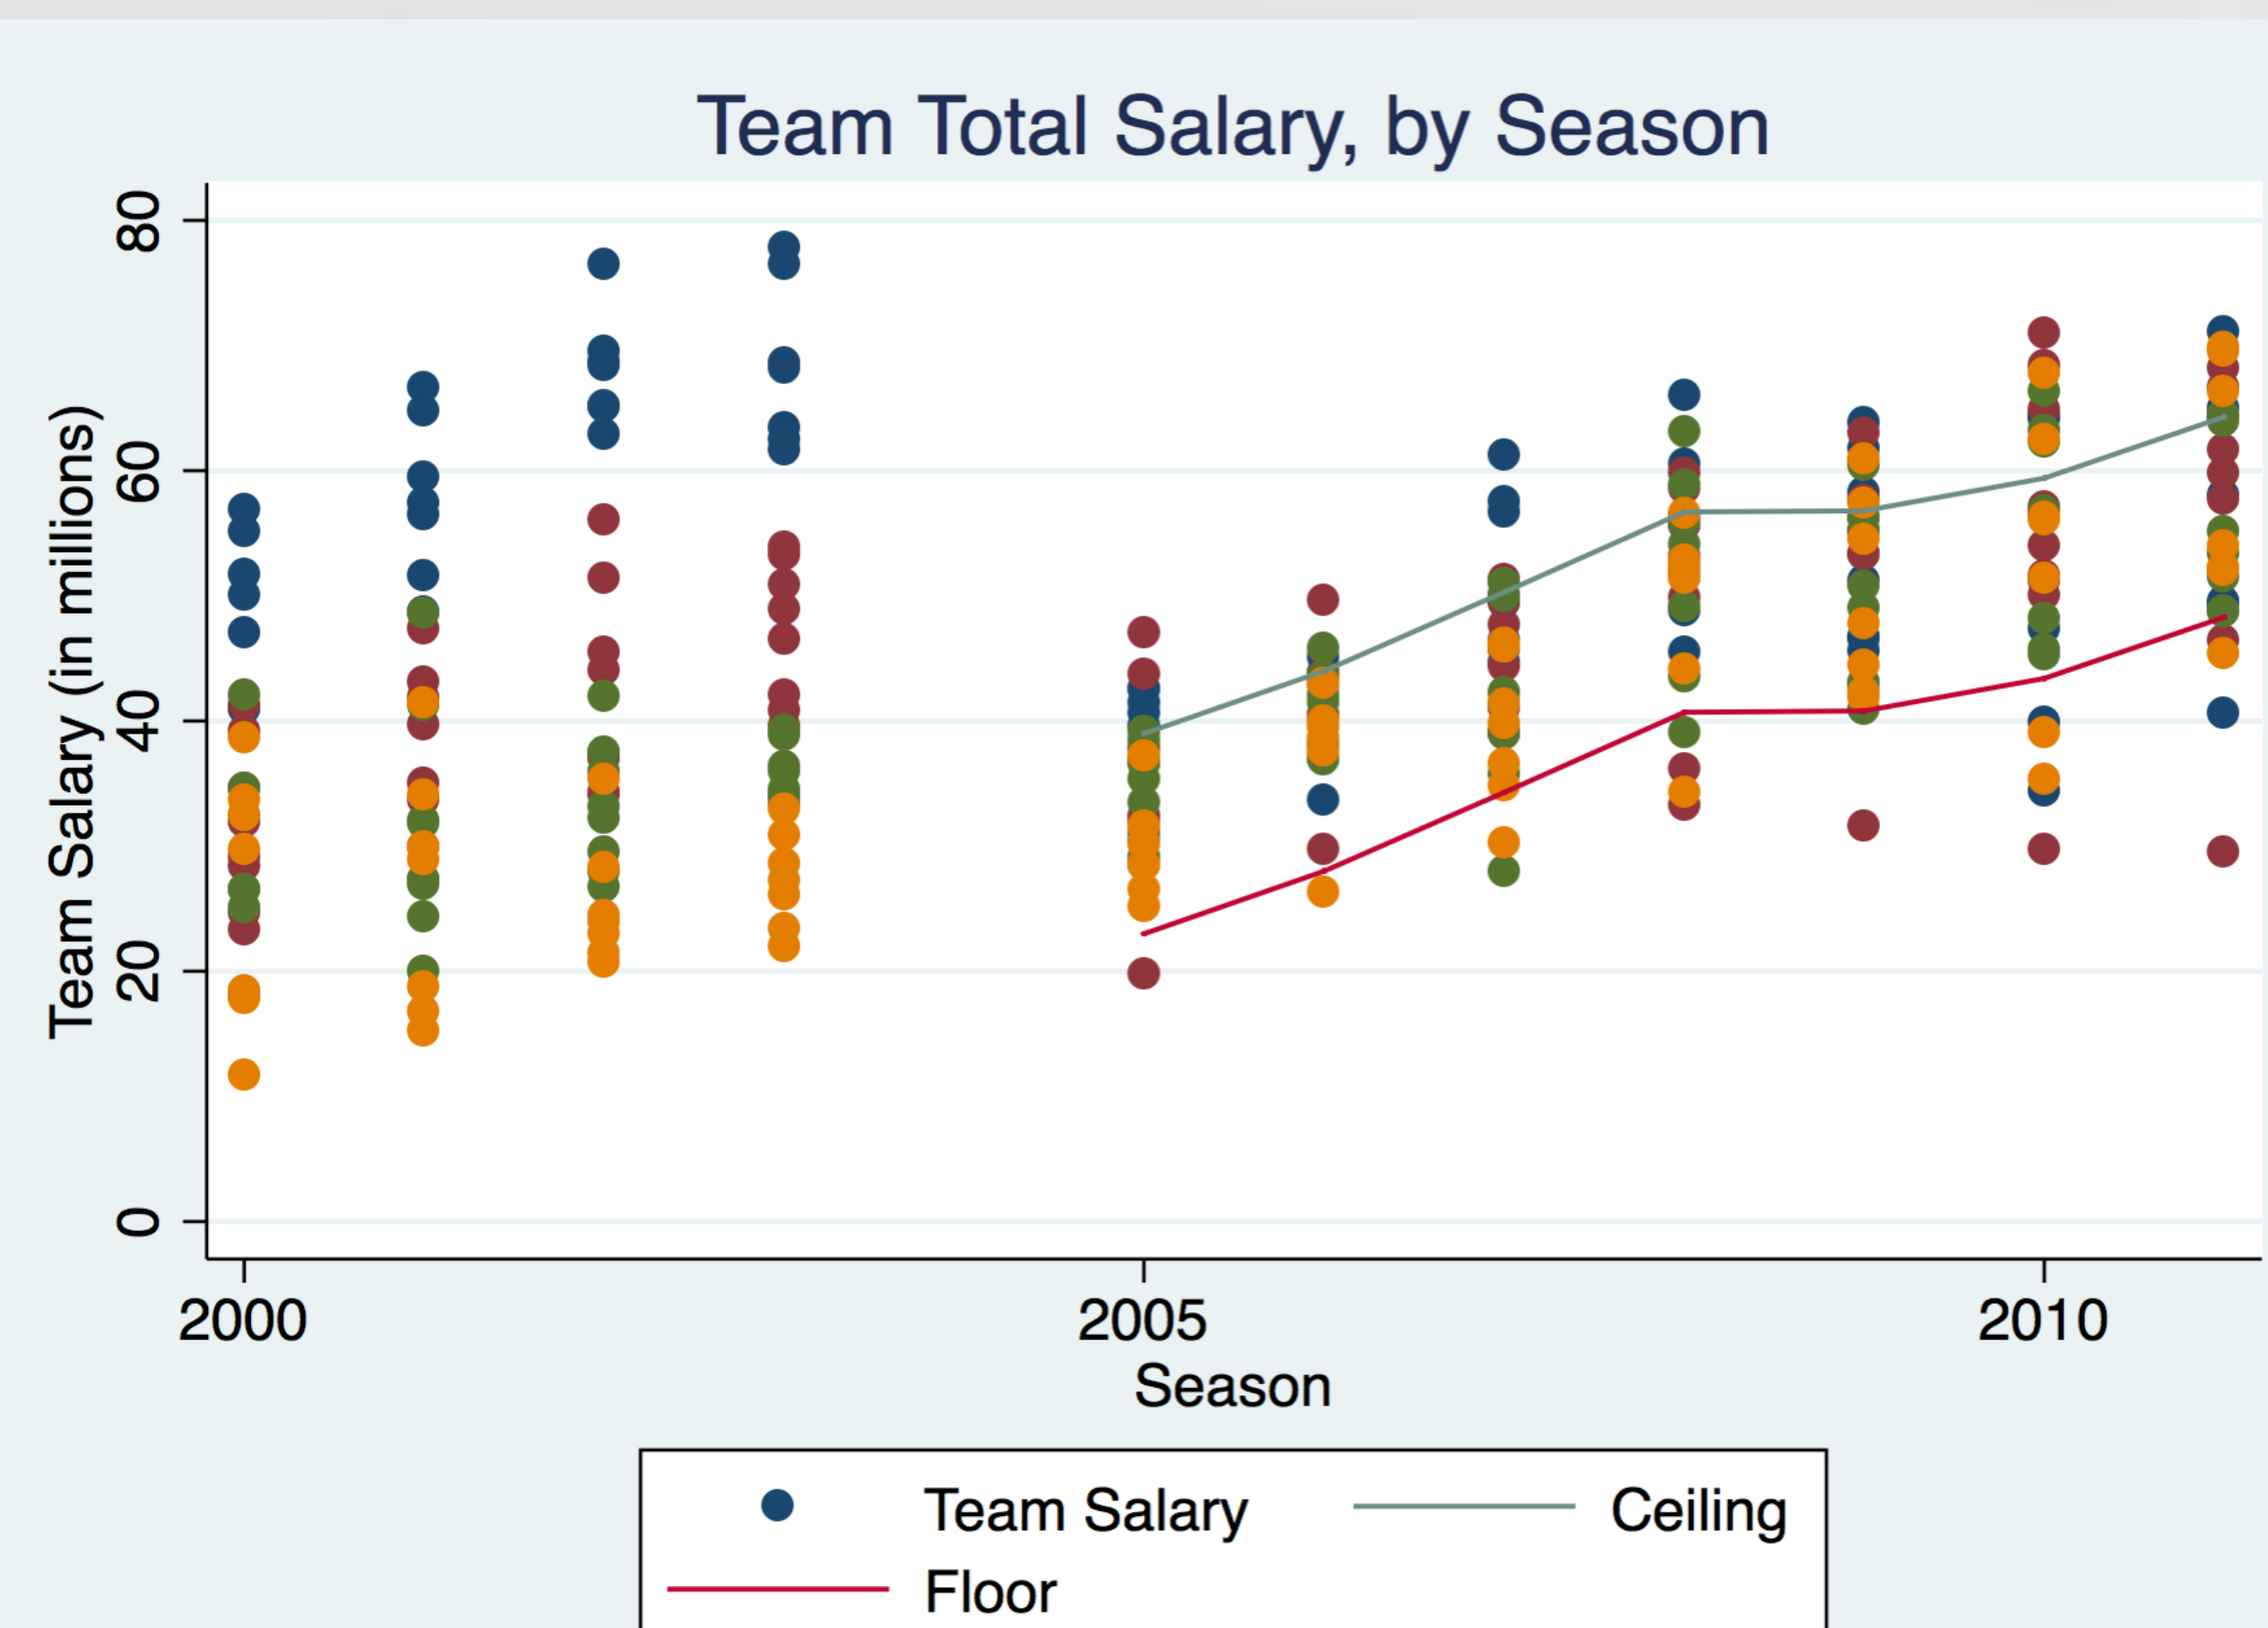
References

Fort, Rodney and James Quirk, "Cross-Subsidization, Incentives, and Outcomes in Professional Team Sports Leagues," *Journal of Economic Literature*, 1995, 33 (3), 1265–1299.

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

Price, Joseph and Justin Wolfers, "Racial Discrimination Among NBA Referees," *The Quarterly Journal of Economics*, 2010, 125 (4), 1859–1887.

Figure 1



	(1)	(2)	(3)	(4)
Post 2004	-0.300*** (0.065)			
Post × Above		0.041 (0.044)	-0.450*** (0.081)	-0.561*** (0.100)
Post × Below		-0.071** (0.026)	0.252*** (0.050)	0.196** (0.079)
Year	0.062*** (0.007)			
Year Fixed Effects	No	Yes	Yes	Yes
Team Fixed Effects	No	No	Yes	Yes
GM Fixed Effects	No	No	No	Yes
Number of Observations	330	330	330	330
Number of Team Clusters	30	30	30	30

Notes: The dependent variable is the natural log of the team salary. Post × Above is a binary variable equal to one if the 2003 team salary was above the 2005 salary ceiling and this team is observed after 2004, and zero otherwise. Post × Below is a binary variable equal to one if the 2003 team salary was below the 2005 salary floor and this team is observed after 2004, and zero otherwise. Cluster-robust standard errors (at the team level) are shown in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2: Salary Caps and Share of Team Salary Devoted to its Five Highest Paid Players, 2000-2011

	(1)	(2)	(3)	(4)
Post 2004	0.021 (0.016)			
Post × Above		-0.000 (0.013)	-0.064** (0.025)	-0.055 (0.033)
Post × Below		-0.025** (0.011)	0.049** (0.019)	0.055* (0.028)
Year	-0.002 (0.002)			
Year Fixed Effects	No	Yes	Yes	Yes
Team Fixed Effects	No	No	Yes	Yes
GM Fixed Effects	No	No	No	Yes
Number of Observations	330	330	330	330
Number of Team Clusters	30	30	30	30

Notes: The dependent variable is the share of the team salary devoted to its five highest paid players. Post × Above is a binary variable equal to one if the 2003 team salary was above the 2005 salary ceiling and this team is observed after 2004, and zero otherwise. Post × Below is a binary variable equal to one if the 2003 team salary was below the 2005 salary floor and this team is observed after 2004, and zero otherwise. Cluster-robust standard errors (at the team level) are shown in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Results

Regression equation

The main econometric model takes the following form:

$$y_{igt} = \beta_1 + \beta_2 Post_t \times Above_i + \beta_3 Post_t \times Below_i + \gamma_i + \delta_g + \pi_t + \varepsilon_{igt}, \quad (1)$$

where y_{igt} is the dependent variable for team i with general manager g in year t . $Post_t \times Above_i$ is a binary variable equal to one if the 2003 team salary was above the 2005 salary ceiling and this team is observed after 2004 and zero otherwise. Similarly $Post_t \times Below_i$ is a binary variable equal to one if the 2003 team salary was below the 2005 salary floor and this team is observed after 2004, and zero otherwise. γ_i and δ_g represent team and general manager fixed effects, respectively. Finally, π_t accounts for year fixed effects. Note that $Post_t$ is absorbed by the year fixed effects once they are included in the regressions.

Table 1 and Figure 1: Salary Cap and Team salaries

Figure 1 simply plots the teams' yearly total salaries (in millions of dollars) over the 2000-2001 to 2011-2012 seasons. There are no points for the 2004-2005 season as it was cancelled. The figure shows a visible reduction in the variation of team spending across the league. The implementation of the salary cap creates a more equal spending pattern across all of the teams, which coincides with the purpose of the cap which was to make teams as even and competitive as possible. In order to investigate whether there was a significant change in the ranking of teams' total salary, we divided teams in four groups based on their 2003-2004 total salary spending. The blue, red, green and yellow dots refer to teams that were respectively, in the top, second, third and bottom quarters of team total salary.

An interesting observation from this figure is the tracking of specific teams. With a color pattern separating teams spending in four groups, we see a clear change in spending habits. This is apparent as the color trends vary more and more with every additional season under the salary cap ruling. This observation further confirms the purpose of the cap, which is to bring a more level playing field within the league.

Table 1 quantifies the effect of the salary cap on teams' total salary. More specifically, it shows the OLS results when the dependent variable is the natural log of the team salary. Columns (2) through (4) estimate equation (1) where we add controls sequentially. Column (1) shows the results of a slightly different equation, where the main explanatory variables are a dummy variable identifying the post CBA years and a time trend variable. This regression simply looks at whether there was a break in the teams' total salary trend.

Column (1) shows that team salaries have **decreased by 30%** with the implementation of the salary cap and it is statistically significant at the 1% level. The effect of spending either above or below the 2005 implemented cap as of 2003 has very little effect in the regression with only the year fixed effect (column (2)), and the signs are counterintuitive. Once we add team fixed effects (column (3)) and also GM fixed effects (column (4)), the results are much larger and have the expected sign. For example, in the preferred specification which has the full set of controls, we've found that team salary has **decreased by 56%** for teams with total spending in 2003 above the 2005 implemented cap ceiling. Also, team salary has **increased by 19.6%** for teams with total spending in 2003 below the 2005 implemented cap floor. Both effects are economically and statistically significant at the 5% level. The contrast of the results between columns (2) and (3) can be explained by the fact that teams which were above (below) the league average total salary before the salary cap were also above (below) after the introduction of the cap—these teams were affected drastically differently by the cap.

Table 2 top 5 salary

Table 2 shows the OLS results when the dependent variable is the share of the team salary devoted to its five highest paid players. The approach used in table 2 is the same sequential method used in table 1 with the year, team and GM fixed effects.

Columns (1) and (2) show that, if we ignore the pre-existing (and 'fixed') team differences, the team share of salary devoted to its five highest paid players has not changed significantly (economically or statistically). The effect of spending above or below the 2005 implemented cap as of 2003 has very little effect in the regression with only the year fixed effect (column (2)), with the result of zero for teams originally above the cap, and a 2.5 percent point decrease for teams originally below the cap. Once we control for team fixed effects (column (3)) and GM fixed effects (column (4)), the results become statistically and economically significant and point to a general conclusion: when teams are forced to reduce (increase) their total team salary, they will spend less (more) on their star players. For example, in the preferred specification which has the year and team controls, we've found that the share of team salary devoted to the top five highest paid players has **decreased by 6.4 percentage points** for teams with total spending in 2003 above the 2005 implemented cap ceiling, and **increased by 4.9 percentage points** for teams with total spending in 2003 below the 2005 implemented cap floor. Both effects are economically and statistically significant at the 5% level.

Conclusion

In conclusion, it is quite evident that the introduction of the salary caps in the NHL affected the way general managers allocate player wages. We have shown that by introducing a maximum cap and a minimum floor, teams on the outside of these limits had to drastically adjust to the new regulations. We have also observed that team spending changes were caused by the general managers themselves changing their behavior and decision making on salary allocation. We see this in the results of the share of salary devoted to the top five highest paid players, incorporating the GM effect. In the end, the cap has caused drastic changes in wage distribution, which has caused teams to rely on more than just deep pockets.

Acknowledgements

I would like to thank my supervising Professor Louis-Philippe Morin, for accepting to guide me through this research and for making this entire experience enjoyable. I would also like to thank Professor Pierre Brochu for his input and guidance and Professor Victoria Barham for introducing me to the UROP project.