The Value of Pace in the NBA

Tim Xin

Undergraduate Economics Honors Thesis

University of California, Berkeley

Advisor: Professor David Ahn

Spring 2012

Acknowledgements: I would like to thank Professor David Ahn for his time and guidance throughout the research and writing process. I would also like to thank Professor Roger Craine for his helpful comments during the development of my research topic.

Abstract

Previous analysis of decision making in the National Basketball Association has shown players to use the shot clock efficiently in scoring opportunities. This paper extends the analysis to late-game situations by examining how teams trade off the value of controlling the length of the game and the expected value of shot attempts. I find that teams strategically adjust their pace with the score differential at the end of games and in certain cases pay for the adjustment with worse shots.

Table of Contents

1. Introduction

Economists have long known that players within a game will strategically use the information that the game is about to end; for example, for the finitely iterated Prisoner's Dilemma, players will defect on the final turn, knowing there is no possibility for punishment. Using backward induction, it is easily demonstrated that every Nash Equilibrium of such a game yields outcomes of deviation on each turn. In this paper, I seek to apply this intuition to analyze decision-making in the National Basketball Association; in particular, how teams adjust their strategies in response to changing time and score situations.

Every basketball fan is familiar with strategic late-game fouling; in the waning seconds of the game, the losing team will foul and regain possession rather than letting the winning team run out the clock. Before the NBA featured a shot clock, which forced a team to attempt a shot in 24 seconds, teams with the lead often employed the "Four Corners" offense, a stalling tactic in which four offensive players stand in the corners of the halfcourt while the point guard dribbles the ball in the middle. In this way, the team with the lead aimed to run out the game clock while minimizing the risk of a turnover, only attempting a shot when an excellent opportunity arose.

With the advent of the 24-second shot clock, however, teams faced a more difficult trade-off. All else equal, the team with the lead prefers fewer total possessions in the game; basic probability theory predicts this will yield a lower probability of the losing team tying the game; similarly, the team without the lead prefers more total possessions in the game. However, both teams also prefer spending the possession on good shots with a high expected value. Since both deliberately waiting until the shot clock runs down and taking shots more quickly than usual tends to yield worse shots on expectation, teams face a trade-off. This paper seeks to better understand this trade-off in the late stages of the game.

It should be noted that while one team (the team in the lead) prefers the game to be slower, the other prefers it to be faster. Therefore, the team playing defense finds value in adjusting the pace of the possession in the opposite direction as the team on offense, and should adjust its team defense accordingly. However, I expect this defensive adjustment to have a small effect on the pace relative to the offensive adjustment. Intuitively, strong offensive players have a variety of options with which to attempt to score; analogously, offensive plays typically feature several players as options. Therefore, it is near impossible for a defense to discourage an offensive team from exercising one of its options quickly. Similarly, there is little a defense can do to discourage an offensive team from letting its point guard run down the shot clock by dribbling in place.

In *Tick Tock Shot Clock: Optimal Stopping in the NBA* (Rao & Goldman, 2009), the authors examined a similar problem: whether NBA players' shooting decisions are optimal. The paper examined two kinds of efficiency. Dynamic efficiency is the condition that the expected value of the shot exceeds the continuation value of the possession, and allocative efficiency is the condition that each player in the line-up has equal marginal efficiency. The authors find that nearly all players and teams achieve dynamic efficiency; that is, players optimize well when it comes to taking shots at certain times in the shot clock.

In *The Problem of Shot Selection in Basketball* (Skinner, 2012), the author posited an alternate model for the continuation value of the possession which takes as parameters the distribution of future shot opportunities, the rate of future shot opportunity encounters, and the turnover rate. The author estimates the parameters using game data and finds that players tend to be overly hesitant to shoot early in the shot clock. The author's explanation for this phenomenon is that players believe the continuation value of the possession to be higher than it is because they are overconfident about their team's ability to protect the basketball; that is, they underestimate the probability of turnovers.

2. Data

The data used in this paper consist of 859 games played by the 30 teams in the NBA during the 2008-2009 regular season, which is publicly available on http://www.basketballgeek.com/data/. Data from games after March 3, 2009 have been omitted. The reason for the omission is that late in the regular season, teams that have been mathematically eliminated from the playoffs are sometimes less committed to winning games in order to secure better draft lottery position and obtain a better player for next season, a strategy known as tanking. Additionally, playoff teams occasionally tank games in order to obtain a more favorable playoff matchup.

From these games, 148433 shooting decisions are observed. I distinguish between a shooting decision and a possession, which is the unit more commonly used to analyze basketball games. By definition, each team has roughly the same number of possessions per game, since one team's possession of the basketball begins when the other's ends. The only exception to this rule comes at the beginning of quarters, when possession is determined by the jump ball at the beginning of the game. In contrast, each team does not necessarily have the same number of shooting decisions. I excluded possessions that conclude with a turnover, since a turnover does not necessarily represent a shooting decision. If a team gains multiple shooting opportunities during one possession—for example, an offensive rebound—I recorded these as multiple shooting decisions. I chose to include free throws resulting from shooting fouls as part of my analysis (in the case of and-ones, I considered them together with the associated made shot) since a shooting foul represents a shooting decision. However, I did not include free throws resulting from non-shooting defensive fouls. I chose to calculate game time remaining independently from any overtime periods, since when teams incorporate game time into their decision-making, they do not know beforehand whether there will be any extra periods.

Table 1: Relevant variables from raw game logs	
Variable	Description
period	Quarter or overtime period
time	Amount of time left in the period
etype	Type of action (e.g. shot)
points	Number of points action yielded
result	Result of action (e.g. make)
type	Type of shot (e.g. layup)

Via some coding, I constructed variables that I directly used in my analysis from the raw game logs.

Table 2: Data overview and description	
Variable	Description
time	Time left in game
diff	Score differential at time of action
shot_time	Number of seconds used from shot clock
outcome	Number of points scored on shot

3. Results

3.1 Relationship of shot time with differential and game time

In order to test the prediction that shot time is decreasing in score differential during late-game situations, I regressed shot_time on time, diff, and time*diff.

$$\text{shot\_time}_i = \text{ß}_0 + \text{ß}_1 * \text{time}_i + \text{ß}_2 * \text{diff}_i + \text{ß}_3 * \text{time} * \text{diff}_i + \varepsilon_i \tag{1}$$

Because I want to examine late-game situations, I excluded all shooting decisions except those in the final 1-5 minutes of the game or the final 1-5 minutes of an overtime period. The reason for excluding the final minute of the game is that teams that are trailing in the final minute often foul intentionally (and hope for missed free throws) in order to regain possession. Additionally, at the end of the game teams are further constrained by the game clock. I also excluded shooting decisions when the absolute value of the score differential was greater than 15; in games that are considered to be out of reach, coaches often concede the game by substituting in bench players to rest the team’s stars.

Table 3: Results of regression (1) in late-game situations				
Variable	Coefficient	SE	t Statistic	p-level
Intercept	13.8676	0.19271	71.96069	0
time	0.2717	0.06042	4.49696	0.00001
diff	0.2883	0.02506	11.50738	0
time*diff	-0.0432	0.00788	-5.4849	0

$R^2 = 0.03774$
N = 9189

To compare, I regressed shot_time on time, diff, and time*diff for non-late-game situations. I examined each shooting decision in the first three quarters. I did not exclude any data on the basis of score differential, since teams are typically still trying to win by the end of the 3rd quarter.

Table 3: Results of regression (1) in non-late-game situations				
Variable	Coefficient	SE	t Statistic	p-level
Intercept	14.07934	0.0618558	227.62	0.000
time	-0.0088498	0.0019452	-4.55	0.000
diff	0.0643476	0.0066098	9.74	0.000
time*diff	-0.0029741	0.0002538	-11.72	0.000

$R^2 = 0.0016$
N = 111426

In the late-game situations, the coefficients conform to expectations. The coefficient on the diff variable is significant and positive. For example, when a team is down by 10 points, the team will, on expectation, attempt a shot 5.8 seconds earlier in the shot clock than a team up by 10 points. Additionally, the coefficient on the interaction term, diff*time, is significant and negative, which is opposite in sign from the coefficient of the diff variable. That is, given a certain score differential, the more time is left on the game clock, the less "extreme" a team's pace adjustment would be. This aligns with the intuition that as the game progresses toward the end, teams find more value in adjusting pace. In non-late-game situations, the coefficient on the diff variable is still significant and positive (0.064), but it is much smaller than the coefficient on the diff variable in late-game situations (0.2883).

Table 4: Shot times given game time and differential intervals

	1-2 min.	2-3 min.	3-4 min.	4-5 min.
[-15, -11]	11.51	13.10	12.92	13.38
[-10, -6]	11.08	13.18	13.37	14.30
[-5, -1]	13.12	14.43	14.43	14.29
[1, 5]	15.99	15.75	16.23	15.42
[6, 10]	15.99	16.77	15.59	15.60
[11, 15]	16.13	16.01	15.77	16.32

Table 4 reinforces this conclusion. There is a clear trend that as one moves down the table (increasing differential) shot time tends to increase. It is less clear that as one moves across the table (earlier in the game) shot time tends to increase for negative differential and decrease for positive differential. However, this relationship seems to hold for negative differential situations.

3.2 Trade-off between expected points and shot time

Since teams that are ahead value holding on to the ball and preventing the other team from having as many possessions, and vice versa, I anticipate teams will trade off expected outcome for pace adjustment. As Rao and Goldman showed, in general, the number of points scored on a shot declines monotonically with time elapsed; intuitively, the less time is left on the shot clock, the higher the continuation value of the possession. To check whether teams trade points for pace adjustment, I compare late-game situations with 3rd-quarter situations, during which I assume teams are indifferent to pace. I consider positive and negative differential situations separately, as I believe the trade-off will produce opposing effects: a positive-differential team will sacrifice points on shots taken late in the shot clock, and a negative-differential team will sacrifice points on shots taken early in the shot clock. Teams without the lead will prefer to take worse shots earlier in the shot clock compared to earlier in the game. On the other hand, teams with the lead will run the clock down at the expense of running a more effective offensive set that yields multiple opportunities throughout the shot clock. In order to test this prediction, I regressed outcome on shot_time, diff, and shot_time*diff.

$$\text{outcome}_i = ß_0 + ß_1 \text{ shot\_time}_i + ß_2 \text{ diff}_i + ß_3 \text{ shot\_time*diff}_i + \varepsilon_i \qquad (2)$$

Table 5: Results of regression (2) (negative differential, late-game)				
Variable	Coefficient	SE	t Statistic	p-level
Intercept	1.48999	0.14971	9.95273	0.E+0
shot_time	-0.02952	0.01037	-2.84579	0.00446
diff	0.01543	0.01543	1.00045	0.31718
shot_time*diff	-0.00117	0.00108	-1.07983	0.28031

R^2=0.01184
N=2748

Table 6: Results of regression (2) (negative differential, 3rd quarter)				
Variable	Coefficient	SE	t Statistic	p-level
Intercept	1.35875	0.08782	15.4714	0.E+0
shot_time	-0.01364	0.00559	-2.43794	0.01479
diff	0.00533	0.00916	0.58177	0.56073
shot_time*diff	0.00046	0.00058	0.79428	0.42705

R^2=0.01258
N=8889

Table 7: Results of regression (2) (positive differential, late-game)				
Variable	Coefficient	SE	t Statistic	p-level
Intercept	1.36996	0.14971	9.95273	0.E+0
shot_time	-0.02952	0.01037	-2.84579	0.00446
diff	-0.01819	0.00999	-1.82056	0.06878
shot_time*diff	-0.00022	0.00101	-0.21977	0.82607

R^2=0.01691
N=2756

Table 8: Results of regression (2) (positive differential, 3^{rd} quarter)				
Variable	Coefficient	SE	t Statistic	p-level
Intercept	1.48873	0.08684	17.14378	0.E+0
shot_time	-0.01758	0.00553	0.18113	0.85627
diff	0.00162	0.00894	-3.17868	0.00148
shot_time*diff	0.0001	0.00057	0.17942	0.85761

R^2=0.00961
N=9460

The results of the regressions are largely inconclusive. For the negative differential cases, the shot_time coefficient is significant for both the 3rd quarter and the end of the game. The absolute value of the coefficient is larger at the end of the game, suggesting the decrease in points scored with respect to shot time is steeper at the end of the game than in the 3rd quarter. This contradicts what I expected for negative differential situations. For the positive differential cases, the shot_time coefficient is significant only for the end of the game, so I cannot make a meaningful comparison. Because I expect the decrease in points during the end of the game to only occur as a result of the trade-off with pace, I expect the decrease to only occur early and late in the shot clock (depending on whether positive or negative differential). Therefore, linear regression may not be the best tool to examine the situation.

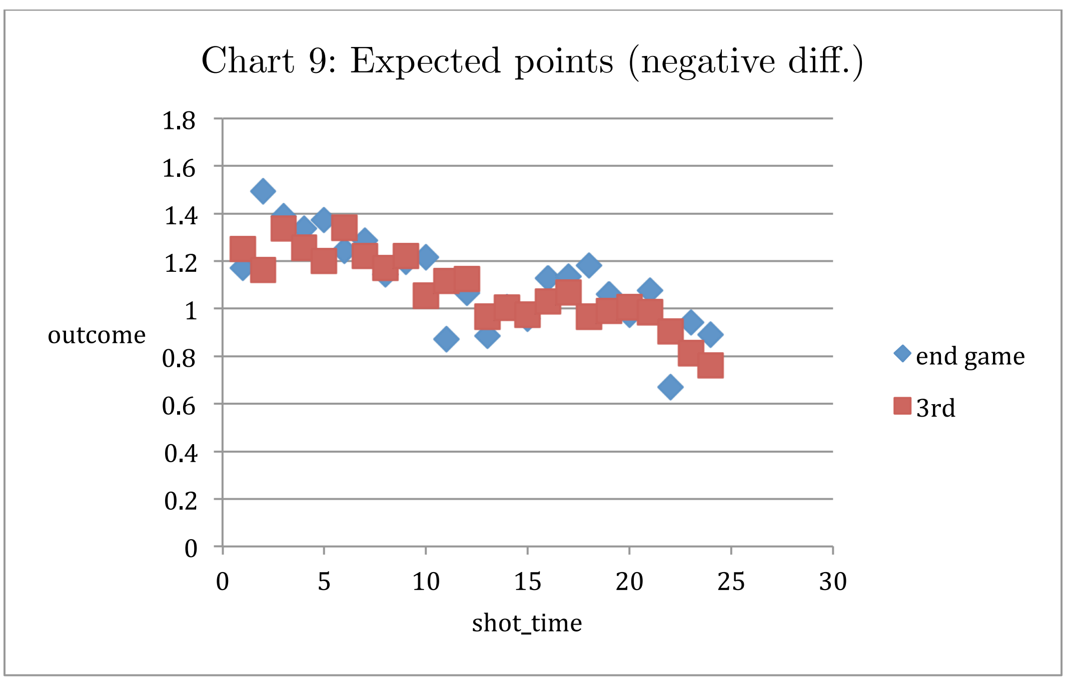

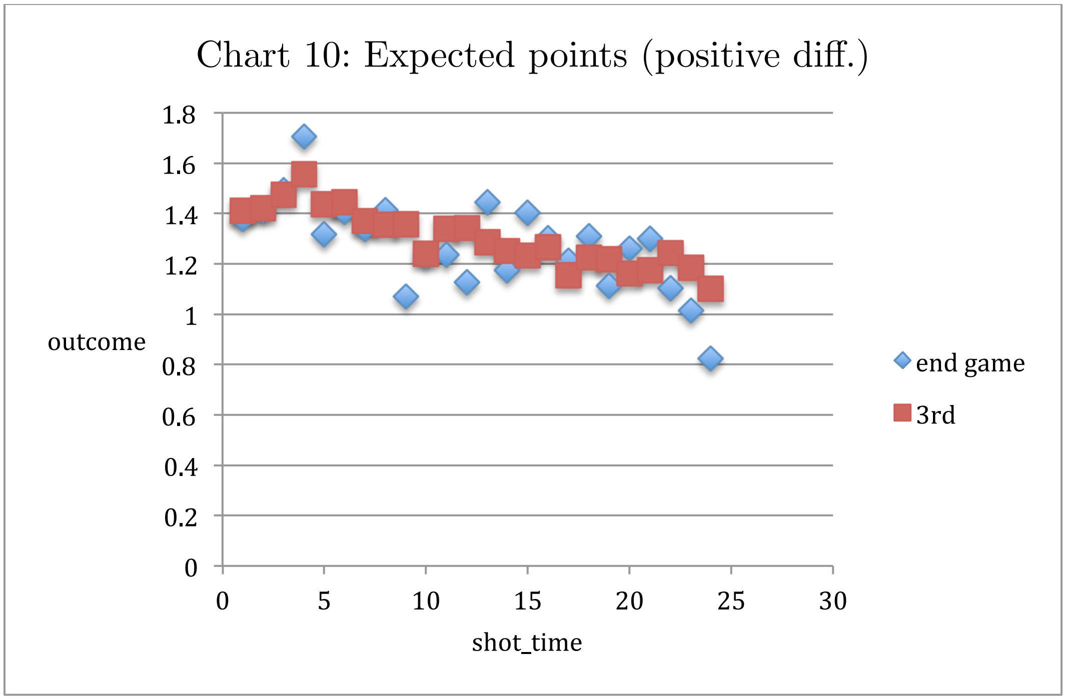

The data reveal that for negative score differential situations, there seems to be no obvious difference between expected points in the 3rd quarter and at the end of the game for any second on the shot clock. In particular, it does not seem like taking shots early in the shot clock (low shot time) yields worse outcomes later in the game.

For positive score differential situations, it appears that expected points in the 3rd quarter and at the end of the game are similar, except for shots taken at the end of the shot clock (high shot time). This conforms to our expectations that teams with positive score differential sacrifice points in order to take more shots at the end of the shot clock.

Finally, I compare expected points across all shot times for late game and 3rd quarter situations. The purpose of this comparison is to show the total effect of the trade-off. The first part of the analysis showed that teams tend to change their pace depending on the score differential and the time remaining in the game, and the second part suggested that such a change in pace can result in lower expected points. If, for example, positive differential teams are indeed holding the ball longer and taking worse shots when they hold the ball longer, I expect the total effect to show up in the average across all positive differential cases.

Table 11: Expected points in late game vs. 3rd quarter		
	Late game	3rd quarter
Positive diff.	1.209	1.267
Negative diff.	1.106	1.058

When teams are ahead, they tend to do worse on the average shooting decision during the final minutes compared to the 3rd quarter. This agrees with the argument that teams do worse on end-of-shot-clock shots and take more of those shots. When teams are behind, however, they tend to do better on the average shooting decision during the final minutes.

4. Limitations

One issue is the difficulty of isolating relevant shooting opportunities from the rest of the data. I was interested in examining decision making in the halfcourt, but my data include some fast break attempts and shots from offensive rebounds as well. On a fast break, a team will usually attempt to score regardless of the situation because the expected point value of such an attempt is overwhelmingly high relative to attempts in the halfcourt. I included fast break shot attempts because it is often difficult to distinguish between an easy fast break opportunity and a shot taken quickly against a set defense. Similarly, a team will often attempt to score quickly after an offensive rebound because the rebounder is very close to the basket, presenting an attractive opportunity (a tip-in being an obvious example). However, depending on the nature of the rebound and the situation, a team will sometimes pass the ball back out and reset the offense. For these reasons, I did not try to remove these shot attempts, but I believe they added noise to the data.

Using the 3rd quarter as a baseline comparison with the end of the game while examining the trade-off between expected points and pace adjustment presents some problems. First of all, it is expected that teams that are trailing by 5 to 15 points in the 3rd quarter, all else equal, are of lower quality than teams that are trailing by the same differential at the end of the game; the opposite holds for teams that are ahead. Therefore, I would expect shooting decisions to produce fewer points in the 3rd quarter than at the end of the game. Furthermore, coaches tend to play their strongest lineups at the end of the game in order to give the team the best chance to win (assuming the outcome of the game is still in doubt). On the other hand, coaches often start the 3rd quarter with a strong lineup and gradually substitute the starters out for reserves in order to rest the starters for the end of the game; the final few minutes of the 3rd quarter often features two reserve lineups. While reserves tend to be weaker on both offense and defense, they tend to be less efficient offensively than starters (that is, the decrease in offensive talent outweighs the effect of weaker defense). Because of these features of the 3rd quarter relative to the end of the game, I would expect scoring efficiency to be lower in the 3rd quarter; this is a possible explanation for why negative differential teams appear to be score more points per attempt late in the game relative to the 3rd quarter.

5. Conclusion

In this paper, I investigated the behavior of NBA teams in using the shot clock to strategically increase or decrease the number of possessions remaining in the game in order to increase their probability of winning. I conclude that teams tend to increase the pace when they are trailing and decrease the pace when they are leading at the end of relatively close games. Furthermore, it is inconclusive whether teams sacrifice expected points in order to adjust the pace, but there is some evidence that teams that are ahead score fewer points when they shoot later in the shot clock during late-game situations relative to the 3rd quarter. Further investigation could focus on testing whether players' decisions are optimal; that is, whether they ought to be further increasing or decreasing the pace, given the strategy of their opponent.

References:

Goldman, Matt and Rao, Justin M. “Tick Tock Shot Clock: Optimal Stopping in the NBA.” (2009)

Skinner, Brian. “The Problem of Shot Selection in Basketball.” (2012)